



ADMINISTRATION GUIDE | PUBLIC

SAP Adaptive Server Enterprise 16.0 SP03

Document Version: 1.0 – 2020-03-04

HADR Users Guide

Content

- 1 Overview of the HADR System. 9**
- 2 Installation Planning. 12**
 - 2.1 Requirements and Restrictions. 12
 - 2.2 System Resource Requirements. 13
 - 2.3 Capacity Planning and Sizing. 14
 - 2.4 Application Compatibility. 15
 - HA Aware. 16
 - 2.5 Replication Limitations. 17
 - 2.6 Accessing the SAP ASE Cockpit Help. 20
 - 2.7 Unsupported Features. 20
- 3 Installing HADR for Custom Application. 22**
 - 3.1 Installation Worksheet. 22
 - The SAP Host Agent. 28
 - 3.2 Installing a New System. 30
 - Unloading the SAP ASE Binaries. 30
 - Configuring SAP ASE. 42
 - Installing the HADR System with Response Files, Console, and Silent Mode. 67
 - Adding Databases to the HADR System. 85
 - Migrating from an Existing Replication Server. 85
 - 3.3 Installing HADR with an Existing System. 86
 - Migrating an SMP Server to an HADR System. 87
 - Collecting Migration Configuration Details. 102
 - Sample `setup_hadr.rs` Response File. 104
 - 3.4 Configuring the SAP ASE Cockpit for the HADR System. 109
 - Configuration for Monitoring. 114
 - 3.5 Installing The Fault Manager. 115
 - Installing and Configuring the Fault Manager. 115
 - Fault Manager Behavior During Automatic Failovers. 132
 - Customizing the Fault Manager. 134
 - Fault Manager Profile Parameters. 134
 - Upgrading the Fault Manager. 142
 - 3.6 Using the Fault Manager with Custom Application. 145
 - Configuring the Fault Manager from the Command Line. 147
 - Administering the Fault Manager. 150
 - Uninstalling the Fault Manager. 151

3.7	Performing a Rolling Upgrade.	156
3.8	Upgrading SAP ASE Version 15.7 DR to Version 16.0 HADR.	173
	Upgrade the SAP Host Agent.	174
	Upgrading to an HADR System.	175
4	Installing HADR for Business Suite.	179
4.1	Requirements and Limitations.	179
4.2	Prepare the Image.	180
4.3	Installing the HADR System on the Primary.	181
	Installing the Business Suite Application.	181
	Install the Data Movement Component.	182
	Run setuphadr to Configure HADR.	184
4.4	Installing the HADR System on the Companion.	185
	Install the Business Suite Application.	185
	Install the Data Movement Component.	186
	Run setuphadr to Configure HADR on the Companion.	187
4.5	Post-Installation Tasks for Primary and Companion Servers.	190
	Restart Replication Server as a Service on Windows.	190
	Add the DR_admin Entry to SecureStore.	190
	Configure Replication Server with sap_tune_rs	191
	Add the dbs_syb_ha and dbs_syb_server Users.	191
4.6	Sample setup_hadr.rs Response File for Business Suite.	192
4.7	Using the Fault Manager with Business Suite.	196
	Configuring the Fault Manager on Business Suite.	198
	Administering the Fault Manager on Business Suite.	199
	Uninstalling the Fault Manager from Business Suite.	200
4.8	Performing a Rolling Upgrade.	205
4.9	Upgrading SAP ASE Version 15.7 DR to Version 16.0 HADR.	222
	Upgrade the SAP Host Agent.	223
	Upgrading to an HADR System.	224
5	HADR Security.	228
5.1	Enabling SSL for the HADR System.	228
5.2	Configuring SSL for External Replication.	229
	Enable SSL for Replicating into an HADR System.	230
	Enable SSL for Replicating Out from an HADR System.	230
5.3	Configuring the Fault Manager in an SSL-Enabled HADR Environment.	231
	Create a Server PSE for the SAP Host Agent on Database Hosts.	232
	Install the Fault Manager.	234
	Create a Client PSE on the Fault Manager Host for the SAP Host Agent.	235
	Create a Server PSE on the Fault Manager Host for the Heartbeat Client.	237
	Create a Client PSE for the Heartbeat Client on Database Hosts.	238

	Set Profile Parameters in the Fault Manager.	241
5.4	Encrypting Databases in an HADR Environment.	241
6	Replicating Data In and Out of an Existing HADR System.	244
6.1	Requirements and Restrictions.	244
6.2	Architectural Overview.	244
6.3	Configuring HADR with an External System.	246
	Requirements for Integrating an HADR System with External Replication.	246
	Configuring Replication In for an HADR System.	247
	Configuring Replication Out From an HADR System.	251
	Migrating SAP ASE to an HADR System.	254
6.4	Removing Replication.	259
	Removing External Replication.	259
	Migrating an HADR System with Replication Out to Standalone SAP ASE.	259
	Migrating an HADR System with Replication In to a Standalone SAP ASE.	261
6.5	Managing an HADR System with External Replication.	261
	Disabling and Enabling Replication to an External Replication Server.	262
	Rematerializing Databases for External Replication.	262
	Monitoring the HADR System Using the SPQ Agent.	270
	Managing Failovers within HADR with External Replication.	271
	Configuring a Primary SAP ASE to Replicate Data to an External System When the Standby Host is Down.	282
6.6	Tuning the SPQ Agent for External Replication.	291
	System Hardware, Operating System, and Networking Configurations.	292
7	Testing Your HADR Configuration.	293
8	HADR Administration.	296
8.1	Using isql to Log in to SAP ASE, Replication Server, and RMA.	296
8.2	Logins and Roles.	297
	Creating the Maintenance Login.	297
8.3	Adding Databases from the Command Line After Installation.	299
8.4	Loading from an External Dump.	302
8.5	Materializing and Rematerializing Databases.	306
8.6	Using SAP ASE Cockpit to Manage and Monitor the HADR System.	311
8.7	Suspending, Resuming, Enabling, and Disabling Databases.	314
8.8	Starting and Stopping the HADR System.	315
8.9	Unplanned Failover.	318
	Performing an Unplanned Failover in the Synchronous Replication State.	319
	Performing an Unplanned Failover in the Asynchronous Replication State.	320
	Recovering the Primary Data Server If SAP Replication Server is Unavailable.	322
	(Optional) Restarting the Primary Data Server Without Synchronization.	323
	Recovery for MultiDatabase Transactions.	324

8.10	The Split-Brain Check: Preventing Multiple Primary Servers.	325
8.11	Planned and Unplanned Failovers.	325
	Verifying the Environment After a Failover.	330
8.12	Member Modes and States.	330
	Determining the Member Mode and State.	331
	Forcing a Server to Primary Mode.	333
8.13	Checks to Perform After an Unplanned Failover.	334
	Managing Data Loss in an Unplanned Failover Operation.	334
	Managing Fault Manager Alerts.	335
8.14	Checking the System Status.	335
	Tracing Latency.	336
	Determining the Backlog.	336
8.15	Evaluating the Health of an HADR Cluster	337
	Common Failure Scenarios.	338
8.16	Read-Only Support from Companion Node in an HADR Cluster.	338
8.17	Adding SAP Replication Server Device Space.	341
	Scenarios That May Require Additional Device Space.	341
8.18	Tuning the HADR Components.	342
	Tuning Replication Agent and Replication Server.	342
	Checking Latency with rs_ticket.	344
8.19	Customizing File Directories.	345
8.20	SQL Statement Replication.	349
	Enabling SQL Statement Replication.	349
	Setting SQL Statement Replication Threshold.	351
	Disabling SQL Statement Replication.	352
	Displaying SQL Statement Settings.	353
8.21	Managing Configuration Files.	355
9	Developing Applications for an HADR System.	356
9.1	SAP jConnect Features for HADR.	356
	Connection Properties for HADR in SAP jConnect.	356
	Retrieve HADR State Change Messages from SAP jConnect.	357
	Retrieve HADR_LIST_MAP Connection Property from SAP jConnect.	360
9.2	SAP ASE ODBC Driver Features for HADR.	361
	Connection Properties for HADR in the SAP ASE ODBC Driver.	361
	Use the SAP ASE ODBC Driver to Get HADR State Change Messages from SAP ASE.	363
	Retrieve the HADR List.	365
	Application Failover.	366
9.3	SAP CTLIB Features for HADR.	371
	Connection Properties for HADR in SAP CTLIB.	371
	Using the CTLIB Context/Connection level properties.	372
	Handling Failover Events.	373

9.4	Support for HA applications on HADR Clients and Servers.	374
9.5	Support for Cluster Edition Applications on HADR Clients and Servers.	375
9.6	HADR Messages from SAP ASE.	377
10	HADR Internals.	380
10.1	Connections.	381
11	Troubleshooting.	386
11.1	Troubleshooting the HADR System.	386
11.2	Recovering from Errors in an HADR System.	388
11.3	Recovering the Primary Data Server If SAP Replication Server is Unavailable.	398
11.4	Restarting the Primary Data Server Without Synchronization.	399
11.5	Installation.	400
	SAP Installer Issues.	400
	Recovering from a Failed Setup.	401
	Performing a Teardown.	403
	Removing an HADR Environment Using the removehadr Utility.	406
11.6	Monitoring.	407
	Utilities for Monitoring the HADR System.	408
	Monitoring the Replication Agent.	408
11.7	Replication.	410
	Troubleshooting the Replication System.	410
	Troubleshooting the RMA.	418
	Checking RMA Version from the Executable.	419
	Troubleshooting Data That is Not Replicating.	420
	Troubleshooting a Secondary Truncation Point That is Not Moving.	423
11.8	Performance.	424
11.9	Failover.	424
11.10	Access, and Login Redirection.	427
	Troubleshooting the SAP ASE Cockpit.	427
	Troubleshooting Replication Agent.	428
11.11	Troubleshooting the Fault Manager.	433
	Fault Manager and SAP Host Agent Commands.	442
11.12	Configuring the RMI Port.	442
12	HADR Reference.	444
12.1	RMA Commands.	444
	sap_add_device.	444
	sap_cancel.	446
	sap_collect_log.	446
	sap_configure_rat.	448
	sap_configure_rs.	451

sap_delay_replication.	454
sap_disable_external_replication.	458
sap_disable_replication.	459
sap_drop_device.	462
sap_drop_host.	463
sap_enable_external_replication.	465
sap_enable_replication.	466
sap_failover.	468
sap_failover_drain_to_er.	472
sap_help.	475
sap_host_available.	476
sap_list_device.	478
sap_materialize.	479
sap_pre_setup_check.	487
sap_purge_trace.	488
sap_resume_component.	490
sap_resume_replication.	492
sap_send_trace.	493
sap_set.	496
sap_set_databases.	505
sap_set_host.	506
sap_set_password.	507
sap_set_replication_service.	512
sap_setup_replication.	513
sap_sql_replication.	514
sap_status.	520
sap_suspend_component.	541
sap_suspend_replication.	544
sap_teardown.	545
sap_tune_rat.	546
sap_tune_rs.	548
sap_update_replication.	551
sap_upgrade_server.	562
sap_verify_replication.	564
sap_version.	567
12.2 SAP ASE Commands.	567
Use Cases for SAP ASE Commands.	567
sp_hadr_admin Syntax.	569
12.3 RMA Proxy Tables.	577
hadrGetLog.	578
hadrGetTicketHistory.	579

	hadrStatusActivePath.	580
	hadrStatusPath.	582
	hadrStatusResource.	584
	hadrStatusRoute.	586
	hadrStatusSynchronization.	588
12.4	Replication Server Commands.	589
	Replication Server Configuration Parameters.	589
	Monitoring State Information.	592
	Starting and Stopping Capture.	603
	Managing SPQ.	605
	Managing SPQ Agent.	609
	Managing External Replication Connections.	617
12.5	Replication Agent Thread for SAP ASE Commands.	624
	RepAgent Configuration Parameters.	624
	Output from sp_help_rep_agent.	627
	sp_configure 'stream rep msg channel timeout'.	630

1 Overview of the HADR System

The always-on option is a high-availability and disaster recovery (HADR) system that consists of two SAP ASE servers: one designated as the primary on which all transaction processing takes place, the other acts as a warm standby (referred to as a "standby server" in DR mode, and as a "companion" in HA mode) for the primary server, and which contains copies of designated databases from the primary server.

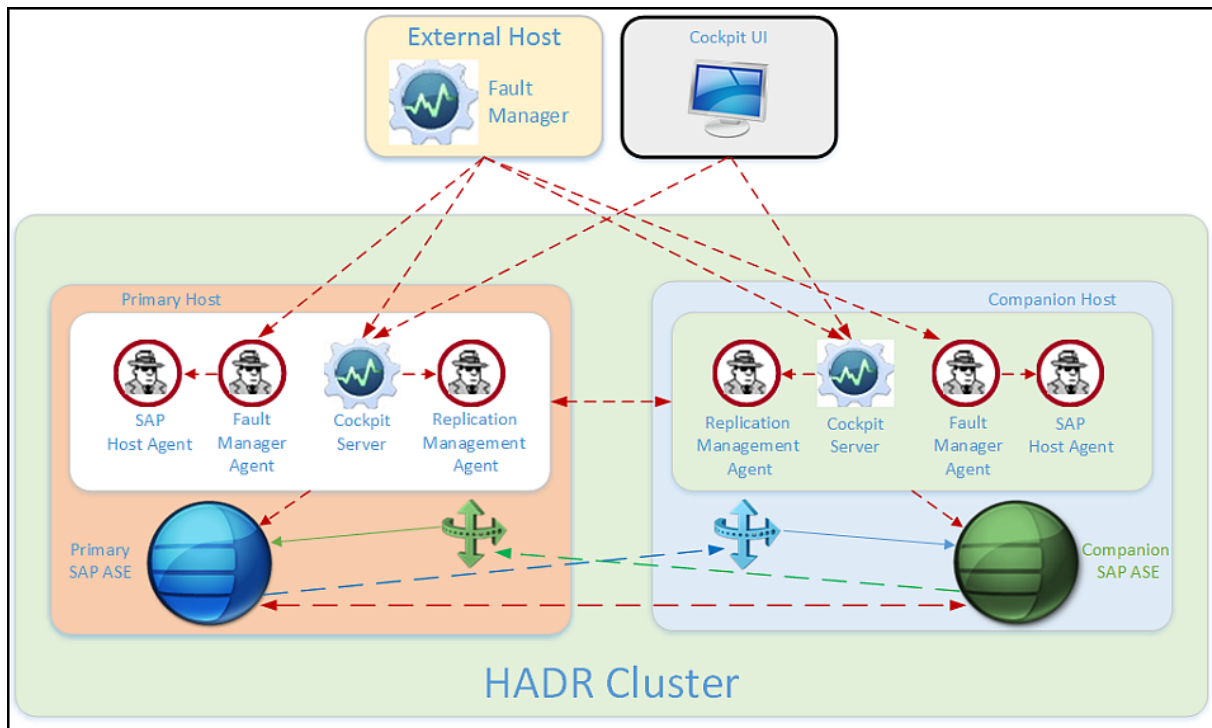
i Note

- The HADR feature included with SAP ASE supports only a single-companion server. Versions SPO3 PLO3 and later support a three node architecture, with two companion SAP ASE servers and a disaster recovery node.
- You can manage multiple databases (certificated with 20) at the same time in an HADR environment for custom applications using SAP Replication Server version 16.0 SPO3 PLO4 and later.

Some high-availability solutions (for example, the SAP Adaptive Server Enterprise Cluster Edition) share or use common resources between nodes. However, the HADR system is a "shared nothing" configuration; that is, each node has separate resources, including disks.

In an HADR system, servers are separate entities, and data is replicated from the primary server to the companion server. If the primary server fails, a companion server is promoted, either manually or automatically, to the role of primary server. Once the promotion is complete, clients can reconnect to the new primary server, and see all committed data, including data that was committed on the previous primary server.

This figure describes architecture for an HADR system.



Servers can be separated geographically, which makes an HADR system capable of withstanding the loss of an entire computing facility. For example, your system may include a primary server that is located in San Francisco and a companion server in San Jose, ensuring that if the primary server is destroyed, the companion server is safe and ready to assume control.

The HADR system includes an embedded SAP Replication Server, which synchronizes the databases between the primary and companion servers. SAP ASE uses the Replication Management Agent (RMA) to communicate with Replication Server, and SAP Replication Server uses Open Client connectivity to communicate with the companion SAP ASE.

i Note

The always-on option, which provides the HADR solution, requires the ASE_ALWAYS_ON license.

The Replication Agent detects any data changes made on the primary server, and sends them to the primary SAP Replication Server. In the figure above, the unidirectional arrows indicate that, although both SAP Replication Servers are configured, only one direction is enabled at a time.

The HADR system supports synchronous replication between the primary and standby servers for high availability so the two servers can keep in sync with zero data loss (ZDL). This requires a network link that is fast enough between the primary and standby server so that synchronous replication can keep up with the primary server's workload. This means that the network latency is approximately the same speed as the local disk I/O speed, generally 5 milliseconds or less. Anything longer than a few milliseconds may result in a slower response to write operations at the primary.

The HADR system supports asynchronous replication between the primary and standby servers for disaster recovery. The primary and standby servers using asynchronous replication can be geographically distant, meaning they can have a slower network link. With asynchronous replication, Replication Agent Thread captures the primary server's workload, which is delivered asynchronously to SAP Replication Server, which then applies these workload change to the companion server.

The most fundamental service offered by the HADR system is the failover—planned or unplanned—from the primary to the companion server, which allows maintenance activity to occur on the old primary server while applications continue on the new primary.

The HADR system provides protection in the event of a disaster: if the primary server is lost, the companion server can be used as a replacement. Client applications can switch to the companion server, and the companion server is quickly available for users. If the SAP Replication Server was in synchronous mode before the failure of the primary server, the Fault Manager automatically initiates failover, and HA-aware applications are transparently failed over with zero data loss. There may be some data loss if the SAP Replication Server was in asynchronous mode, in which case you use manual intervention to failover for disaster recover. If the HADR cluster is in asynchronous mode, applications are not transparently failed over because the cluster anticipates that it may need to perform some corrective actions concerning potential data loss prior to starting new transactions.

Connection attempts to the companion server without the necessary privileges are silently redirected to the primary companion via the login redirection mechanism, which is supported by Connectivity libraries (see [Developing Applications for an HADR System \[page 356\]](#)). If login redirection is not enabled, client connections fail and are disconnected.

The SAP ASE HADR option installs these components:

- SAP ASE
- SAP Replication Server

- Replication Management Agent (RMA)
- SAP Host Agent
- Fault Manager
- SAP ASE Cockpit

2 Installation Planning

There are a number of tasks you perform prior to installing and configuring the HADR system, including reviewing the recommendations and restrictions, planning the system's capacity and sizing, and verifying your application compatibility.

2.1 Requirements and Restrictions

There are a number of requirements and restrictions for the HADR system.

The requirements are

- This release of the HADR system is supported on the following platforms:
 - Solaris x64 on SPARC
 - Linux x64 on x86-64

i Note

Set the tcp parameter `net.ipv4.tcp_retries2` to 8 for all hosts in the HADR system including the Fault Manager host. See *Adjusting the Client Connection Timeout for Linux in Configuration Guide for UNIX* for more details.

- HP IA-64
- AIX 64-bit
- Windows x64
- Operating system version and patch levels – see the SAP ASE installation guide and release bulletin for your platform.
- (Windows) Install Microsoft Visual C++ redistributable package for VS2013. If this package is not available, the Fault Manager installer fails, and you see this message:

```
rsecssfx.exe - System Error:  
The program can't start because MSVCP120.dll is missing from your computer.
```

- Use the same platform for both the primary and the companion servers. For example, you cannot have a primary server on Solaris and a companion on Linux.
- Create an operating system user named "sybase" to install and maintain the software.
- The HADR system requires two hosts. Installing Fault Manager requires a separate third host. The Fault Manager host must have the same platform as the HADR nodes. For example, if you install the HADR nodes on Linux x64, the Fault Manager host must also be installed on Linux x64.
- Synchronous replication requires a solid state drive (SSD) or other fast storage device.
- (Linux) Fault Manager requires GLIBC version 2.7 or later.
- (HP) Fault Manager requires the C++ libCsup11.so.1 library.
- A cluster ID (CID) database. The cluster ID is a three-letter identifier for the cluster. The CID database is for internal use only.

- The page size, default language, character set, and sort order must match on both the primary and companion SAP ASE. Replication Server is configured by Replication Management Agent (RMA), which makes sure that the Replication Server character set matches that of the primary server. If you configure the HADR system using the SAP installer or the `setuphadr` utility, the tools takes care of ensuring consistency across the HADR system.
- Applications must be replication compatible. See [Application Compatibility \[page 15\]](#).
- Applications cannot rely on unsupported features. See [Unsupported Features \[page 20\]](#).
- As the `$SYBASE/REP-16_0/ASA17/bin/dbltm` is 32-bit, install the following 32-bit libraries in your Replication Server's host operating system:

For Linux RedHat:

```
○ sudo yum install libgcc.i686
```

```
○ sudo yum install glibc.i686
```

For Linux SuSE:

```
○ zypper install glibc-32bit
```

```
○ zypper install libgcc_s1-32bit
```

For more information on how to install 32 bit libraries, refer the Linux OS distribution vendor page.

- Requirements for SAP Host Agent:
 - A `sybadm` user with sudo privileges for SAP Host Agent. The `sybadm` user requires sudo privileges.
 - The person who installs SAP Host Agent (directly or through SAP ASE installer) must have root or sudo permission. SAP Host Agent installation requires root or sudo permission because it installs files to `/usr/sap/hostctrl` and creates the `sapadm` user.
 - You must install the SAP Host Agent on a machine that has a remote mount to the HADR release area so that the `saphostctrl` utility can read the SAP Replication Server log file. The `saphostctrl` utility cannot read the SAP Replication Server log file from a remote drive or partition.

The restrictions are:

- Do not use the `sp_dbextend` system procedure to extend the size of databases. Doing so can result in the active and primary databases not being the same size, and replication to the standby server could be blocked. Additionally, mismatched database and device sizes can cause access to applications and rematerialization efforts to block after a failover.
- HADR is not supported on the Developer and Express Edition licenses.
- Business Suite does not support SAP ASE Cockpit.
- Replication of a table without unique keys if there is a Java object datatype in the table.

2.2 System Resource Requirements

Installing the HADR system with disaster recovery includes a number of system resource requirements.

- Each database that participates in HADR, including the master database, requires a minimum of 2GB of space for the simple persistent queue (SPQ).
- Each database from any volume likely requires an additional CPU core for processing at the replicate system. High volume databases or databases with very wide tables may require additional CPU cores. If the replicate SAP ASE is used only for DR purposes, it is probably using a small amount of processing

power, so its CPU capacity is likely available to other servers. However, if the standby system is used for reporting, additional CPU capacity may be needed for the system.

- HADR components may need approximately 2GB of memory for each replicated database. Since SAP ASE typically uses pinned shared memory segments that are preallocated (unlike CPU, which is not), Replication Server cannot easily share memory with SAP ASE. As a result, this 2GB of memory is in addition to other SAP ASE requirements.
- Replication Server components requires 3 consecutive ports, beginning with the Replication Server port number (for example, 5005, 5006, and 5007) and the RMA requires 5 consecutive ports, ending with the specified port number (for example 4988, 4989, 4990, 4991, and 4992). These ports must be accessible from the other hosts involved in the HADR system, including the Fault Manager hosts.

2.3 Capacity Planning and Sizing

When you configure Replication Server, use the primary transaction log and the expected rate of primary transaction log generation as the key parameters to tune and optimize the replication environment to ensure optimal performance.

SAP recommends the following sizes for the server resources in an HADR system:

Server Size	Primary transaction log generation rate (GBs per hour)	Replication Server memory, in GBs	Replication Server CPU	Supports this number of large transactions in parallel
Small	7	4	2	1
Medium	15	8	4	2
Large	25	16	8	4
Extra Large	25	24	16	8

These examples describe the tuning parameters (in terms of storage and computing power, in GBs and CPUs, respectively) that are used to achieve the best performance for a given rate of output log generation:

- If the rate of primary log generation is 3.5 GB per hour and the acceptable latency is less than five seconds, use the `sap_tune_rs` command to configure Replication Server with 4 GB of storage and 2 CPUs. For example:

```
sap_tune_rs SFHADR1, 4,2
```

This example assumes there is a single user database being replicated; user sites can have many more than this. See [Tuning the HADR Components \[page 342\]](#) and [sap_tune_rs \[page 548\]](#).

- Include a single large transaction thread if there are occasional large transactions in your applications by setting `dsi_num_large_xact_threads=1`. See the *Replication Server Administration Guide - Volume 2 > Performance Tuning > Using parallel DSI threads > Configuring parallel DSI for optimal performance*.
- If the rate of primary log generation is considerably high (for example, 11 GB per hour) and the acceptable latency is less than five seconds, configure the Replication Server with 8 GB of storage and 4 CPUs as the computing power.

- Use `sp_spaceused syslogs` or `select loginfo '<dbname>', 'active_pages'` to measure the primary server's transaction log for a period of time.

Note

Make sure you disable `trunc on checkpoint` or any other commands that truncate the transaction log. The transaction log should not be truncated when measuring the transaction log generation rate.

The following example takes a sample of the log generation rate on the primary server for the `user_DB` database (first verifying that `trunc log on chkpt` is disabled on this database). This example uses a 10-minute period of time, which is short for a production system. Your site should use a longer time period to view a reasonable output:

```
sp_dboption 'user_DB', 'trunc log on chkpt', 'false'
go
use user_DB
go
sp_spaceused syslogs
go
waitfor delay '00:10:00'
go
sp_spaceused syslogs
go
name total_pages free_pages used_pages reserved_pages
-----
syslogs 6656 383 40 0
name total_pages free_pages used_pages reserved_pages
-----
syslogs 6656 383 5000 0
To calculate the Primary log generation rate for 10mins
= (After used pages - Before used pages) * server_page_size
= ((5000-40)*16)/1024
= 77.5MB in 10mins
= 465MB/hr
```

See the *Reference Manual: System Procedures* for more information about `sp_spaceused`.

The `sap_tune_rs` command accepts these input parameters: `<site_name>`, **memory limit**, and **number of CPUs**. Replication Server uses these parameters to achieve the best performance while keeping latency to a minimum.

When necessary, use `sap_tune_rs` command to tune Replication Server.

Use `sap_tune_rs` command and the appropriate parameters to display the modified Replication Server configuration for an overview of the modified Replication Server tunings (configuration parameters).

`sap_tune_rs` command prompts you to restart Replication Server for the `sap_tune_rs` modifications to take effect.

2.4 Application Compatibility

The HADR system includes four different clients for connections.

- HADR-aware clients – using connectivity, drivers receive HA failover event notification and utilize features that identify when, and to which server, to reconnect after a failover event. The application takes necessary

actions to re-establish the context (for example, the default database, `set` options, prepared statements, and so on), and resubmits incomplete transactions. If the application has created any intermediate data in non-replicated databases as part of re-establishing the context (for example, in `tempdb`), regenerate this intermediate data as well.

- HA-aware clients – use existing HA failover mechanisms to move to the new primary server after a failover event. However, they must re-establish context in the new primary server and resubmit incomplete transactions (with the configuration changes), similar to HADR-aware clients.
- Cluster Edition-aware clients – similar to HADR and HA-aware clients, connections from Cluster Edition-aware clients automatically reconnect to the new primary server after a failure. In addition, similar to HADR and HA-aware clients, the application should include functionality to re-establish context and resubmit failed transactions.
- Non-HA clients – receive socket disconnect errors, and must reconnect to the new primary server managed by the application. Configure the Fault Manager to use a floating IP address that points to the new primary site as part of failover action. When configured for floating IP, the Fault Manager starts the listener using `<floatingIP:port_number>` on the primary server, and upon failover, moves the IP to the companion host and starts the listener using that `<floatingIP:port_number>` on the new primary server. Non-HA and Non-HADR aware clients can reconnect to the primary server by using the `<floatingIP:port_number>` in the connection string.

2.4.1 HA Aware

Required changes to CTLIB applications running in an HADR system include setting the `CS_HAFAILOVER` property, modifying the `interfaces` file, writing application failover messages, and adding return codes.

Procedure

1. Set the `CS_HAFAILOVER` property using the `ct_config` and `ct_con_props` CTLIB API calls. Set this property at either the context or the connection level using the following syntax:

```
ct_config(context, action, CS_HAFAILOVER, buf, buflen, outlen)
ct_con_props(connection, action, CS_HAFAILOVER, buf, buflen, outlen)
```

2. Modify the `interfaces` file so clients fail over to the secondary companion.

The `interfaces` file includes a line labeled `hafailover` that enables clients to reconnect to the secondary companion when the primary companion crashes, or when you issue a `shutdown with nowait`, which triggers a failover.

3. Write application failover messages according to these parameters:
 - As soon as the primary begins to go down, clients receive an informational message that failover is about to occur. Treat this as an informational message in the client error handlers.
 - After you set the `CS_HAFAILOVER` failover property and the `interfaces` file has a valid entry for the `hafailover` server, the client connection is a failover connection, and clients reconnect to the secondary companion appropriately.
However, if the failover property is set but the `interfaces` file does not have an entry for the `hafailover` server (or vice-versa), there is no failover connection but rather, a normal connection

with the failover property turned off. Inform the user to check the failover property to determine whether the connection is a failover connection.

4. Add return codes.

When a successful failover occurs, the client issues a return value named `CS_RET_HAFAILOVER`, which is specific to the following CTLIB API calls:

```
ret = ct_results(cmd, result_type)
```

```
ret = ct_send(cmd)
```

`CS_RET_HAFAILOVER` is returned from the API call during a synchronous connection. In an asynchronous connection, these APIs issue `CS_PENDING`, and the `callback` function returns `CS_RET_HAFAILOVER`. Depending on the return code, the customer can perform the required processing, such as sending the next command to be executed.

Rebuild your applications, linking them with the libraries included with the failover software.

i Note

In a Custom Application environment, you cannot connect clients with the failover property until you issue `sp_companion resume`. If you do try to reconnect them after issuing `sp_companion prepare_failback`, the client stops responding until you issue `sp_companion resume`.

2.5 Replication Limitations

Restrictions for Primary Key Constraints or Unique Indexes

A `sp_primarykey` designation is insufficient for application tables in databases participating in the HADR system; they require a primary key constraint or a unique index. Although some tables may work without a primary key, operations such as normal updates and deletes as well as inserts, updates and deletes of large object (LOB) data may be extremely slow, and this dramatically increases latency, which in turn significantly increases the failover time.

Restrictions for Columns that Use Float Or Real Datatypes as Primary Key Columns

Application tables should not use columns that use float or real datatypes as primary key columns. The interpretation of approximate numerics such as `float` or `real` is often offloaded to floating point processors on the CPU. Different CPU versions may have different floating point unit (FPU) versions, so even the same CPU hardware (such as Intel x86/64 Haswell EX) may translate the floating point values differently. Because

the HADR system uses logical replication, `update` and `delete where` clauses constructed from primary keys with `float` or `real` datatypes may return 0 rows affected if the inserted `float` or `real` value was interpreted differently by the underlying hardware. Due to replication validation at the replicate, this condition would result in the HADR system suspending delivery to the standby server until the problem is fixed.

Restrictions for Primary Key Updates with Multiple Rows Affected

Application tables with primary keys based on sequential or monotonically increasing values should not perform multirow updates on the primary keys using an expression.

You should not allow client applications to update primary key columns in multiple rows of a replicated table such that keys duplicate keys of another row, as these updates may cause errors in the replicate database.

In this example, `pcol` is the primary key for the `reptbl` table in the primary database, and it includes three rows with values of 1, 2, and 3:

```
pcol
-----
      1
      2
      3
```

Running this command may cause errors or incorrect data in the replicate database:

```
update reptbl
set pcol = pcol + 1
```

The values for `pcol` after running the command at the primary database are:

```
pcol
-----
      2
      3
      4
```

Replication Agent retrieves the log records and submits the records to Replication Server using commands similar to:

```
update reptbl
set pcol = 2 where pcol = 1
update reptbl
set pcol = 3 where pcol = 2
update reptbl
set pcol = 4 where pcol = 3
```

However, because Replication Server treats each row as an independent update, the first row is updated three times and the second row is updated twice. If there is a unique index on the table, the additional updates cause errors in the replicate databases. If the replicate table does not contain a unique index, this table will have duplicate rows, as shown here.

Values Before	Values After
1	4
2	4
3	4

Restrictions on Distributed Transaction Management

Applications cannot depend on distributed transaction management (DTM, for example XA) transaction recovery. DTM is often coordinated across multiple data sources, such as a JMS message bus and a DBMS. In these systems, removing the message from the bus should not happen unless the database operation was committed. If the database operation commits, the message must be removed from the message bus.

DTM specifications include recovery procedures for when either of the sources fails prior to the commit. Generally, any DTM-capable data source must be able to recover a DTM transaction to the point of failure if the data source fails. As a result, the DTM coordinator can re-enlist the transaction and roll it back or commit it based on the last known transaction state at the DTM coordinator.

If the transaction was prepared on both nodes but the commit state is unknown, the DTM coordinator typically reconnects to the recovered data source and attempts to commit the transaction.

In an HADR system, the failover data server is a different instance and is unaware of the distributed transaction IDs. In addition, because the embedded replication only forwards fully committed transactions, the uncommitted database records are stranded inside the data replication components. While these records can be extracted and replayed manually, it is not a fast operation that can be achieved during a failover. Consequently, distributed transactions are not supported.

Restrictions Caused by Cross-Database Transactions

If the standby system is used for reporting, the reporting applications cannot be sensitive to timing differences caused by cross-database transactions. For example, if a transaction inserts data into `database_A` and `database_B` on the primary, because these inserts proceed independently and in parallel through the HADR system, they may be applied in a different order at the replicate databases. This may result in brief data inconsistencies for reports that query across both databases.

This restriction holds true for cross-database declarative constraints in which two independent transactions at the primary insert into `database_A` and `database_B`, respectively, and a foreign key exists from `database_B` to `database_A`. At the standby, due to independent and parallel processing of the different database log records, the child insert in `database_B` may happen ahead of the parent row in `database_A` (the HADR system is able to suspend DSI enforcement for `write` operations).

Configuration Restrictions

Do not set these configuration parameters to 0 at both the server and the connection levels:

- `sqm_cmd_cache_size`
- `sqt_max_cache_size`
- `dsi_sqt_max_cache_size`

Setting these configuration parameters to 0 results in this sequence of events:

1. The workload is aborted
2. RMA hits a stacktrace from an insufficient number of connections

3. The primary server hangs in a log suspend state on many of the DR_admin connections
4. The active Replication Server assumes hibernation mode

2.6 Accessing the SAP ASE Cockpit Help

SAP Adaptive Server Enterprise Cockpit (SAP ASE Cockpit) is a graphical administration tool for managing and monitoring SAP ASE and the HADR system. SAP ASE Cockpit supports SAP ASE version 16.0 SP02.

SAP ASE Cockpit provides availability monitoring, historical monitoring, and real-time monitoring in a scalable Web application. It offers real-time alerts of availability, performance, and capacity issues, intelligent tools for spotting performance and usage trends, as well as the general health of the HADR system. Availability, performance, and capacity alerts are configured and enabled by default. Unlike SAP Control Center, SAP ASE Cockpit is designed as an onboard management solution, where you install the cockpit on each SAP ASE host to manage and monitor that system.

test conditions

Related Information

[SAP Adaptive Server Enterprise Cockpit](#)

2.7 Unsupported Features

The HADR system does not support the full functionality of SAP ASE version 16.0 SP03.

The following is either not supported, or has restrictions:

- HADR is not supported in process kernel mode.
- Multi-path replication.
- In-memory databases.
- Shared disk cluster.
- High-availability platforms (for example, Veritas HA, Sun Cluster, HACMP, Service Guard, and so on).
- Distributed transaction management (DTM).
- LDAP network security mechanism. You can use LDAP to connect to SAP ASE from your user clients, but cannot use LDAP in the HADR system.
- Primary and companion servers participating in the HADR system using different platforms, page sizes, default language, character set, or sort orders (for example, a different character set on the primary and companion server).
- Localization of stream replication messages. Stream replication is also known as Component Interface (CI).
- `sp_setreplicate`, a deprecated system procedure.

- Stored procedures marked for replication using a table replication definition (using `sp_setrepproc` with the `table` option). You can, however, use `sp_setrepproc` with the `table` option when you are replicating stored procedures into or out of an HADR system.
- Request functions. The HADR system does, however, support the replication of stored procedures as applied functions from the primary as long as you mark the procedure for function replication by using `sp_setrepproc <procedure_name>, 'function'`. You can also replicate stored procedures as applied or request functions into or out of an HADR system.

3 Installing HADR for Custom Application

The steps for installing the HADR system differ depending on whether you are installing a completely new environment, using an existing SAP ASE for the primary, using the SAP installer with or without a response file, or installing using the `setuphadr` utility.

For information about:

- Installing a new system – see [Installing a New System \[page 30\]](#) and [Using setup.bin or setupConsole.exe with a Response File \[page 67\]](#)
- Using an existing SAP ASE for the primary companion – see [Installing HADR with an Existing System \[page 86\]](#)

You can configure, or migrate an existing, HADR system now or run the `$$SYBASE/$SYBASE_ASE/bin/setuphadr` utility to configure HADR at later time. See [Migrating an SMP Server to an HADR System \[page 87\]](#).

Installing and Configuring SAP ASE on a Two-Node System

SAP ASE supports the HADR system in a remotely distributed two-node topology with primary and companion sites. Run the installer separately on both sites.

i Note

Installing the software on the first node prepares SAP ASE and Backup Server for the HADR system. However, the actual setup of the HADR system occurs when you install the software on the second node.

The HADR system requires that the always-on nodes use SSD (solid-state drive) or another type of fast storage device for Replication Server if you configure the HADR system with synchronous replication mode.

See [Migrating an SMP Server to an HADR System \[page 87\]](#) for information about migrating an existing SAP ASE server to an HADR system.

3.1 Installation Worksheet

Installing the SAP HADR system requires you to enter and re-enter values for the primary, standby, and Fault Manager hosts. Record the values in a worksheet as you go through the installation to use as future reference. Sample (or default) values used in this guide are provided in square brackets.

i Note

Do not use this sheet to record passwords.

Value	Primary Site	Companion Site
1) SAP ASE installation directory		
	[/work/SAP1]	
2) Installed under different user account?		
	[No]	
3) Data directory		
	[/work/SAP1/data]	
4) SAP ASE name (default is the host name in capital letters)		
	[SFSAP1]	
	[SJSAP2]	
5) SAP ASE system administrator's password. Do not record your password.		
6) Enable SAP ASE for Cockpit monitoring		
7) Technical user		
	[tech_user]	
8) Technical user password. Do not record your password.		
9) Host name		
	[SFMACHINE1]	
	[SJMACHINE2]	
10) SAP ASE port number (default is 5000)		
	[5000]	
11) SAP ASE error log location		
	[/work/SAP1/ASE-16_0/ install/SFSAP1.log]	
12) Application type		
	Same on primary and companion sites	
	[Mixed (OLTP/DSS)]	

Value	Primary Site	Companion Site
13) SAP ASE page size		
Same on primary and companion sites		
[4k]		
14) Default language		
Same on primary and companion sites		
[us_english]		
15) Default character set		
Same on primary and companion sites		
[iso_1]		
16) Default sort order		
Same on primary and companion sites		
[bin_iso_1]		
17) Optimize SAP ASE configuration		
18) Create sample databases		
19) Master device location, size, and database size	•	•
Same on primary and companion sites	•	•
• [/work/SAP1/data/master.dat]	•	•
• [52MB]		
• [26MB]		
20) System procedure device location, size, and database size	•	•
• [/work/SAP1/data/sysprocs.dat]	•	•
• [196MB]		
• [196MB]		
21) System device location, size, and database size	•	•
• [/work/SAP1/data/sybsysdb.dat]	•	•
• [6MB]		
• [6MB]		

Value	Primary Site	Companion Site
22) Tempdb device location, size, and database size	•	•
• [/work/SAP1/data/tempdbdev.dat]	•	•
• [100MB]		
• [100MB]		
23) Enable PCI		
24) Backup Server name		
[SFSAP1_BS]		
25) Backup Server port (default is 5001)		
[5001]		
26) Backup Server error log location		
[/work/SAP1/ASE-16_0/install/SFSAP1_BS.log]		
27) Allow Hosts		
[None]		
28) Cluster ID (No default value. Requires three alphanumeric characters)		
Same on primary and companion sites		
[AS1]		
29) Replication mode		
Same on primary and companion sites. Synchronous or asynchronous.		
[sync]		
30) Site name (also known as HADR server name. No default value. Must be alphanumeric with no punctuation or special characters)		
[SFHADR1]		
[SJHADR2]		
31) Database dump directory		
[/work/SAP1/data]		

Value	Primary Site	Companion Site
32) RMA RMI port (default is 7000)		
[7000]		
<div style="background-color: #f0f0f0; padding: 10px;"> <p>i Note</p> <p>RMA RMI occupies five consecutive ports, with the configured port occupying the highest number. If the configured RMA RMI port number is 7000, for example, it also needs ports 6999, 6998, 6997, and 6996.</p> </div>		
33) RMA TDS port (default is 7001)		
[7001]		
34) Replication Server port (default is 5005)		
[5005]		
35) SRS device buffer directory		
[/work/SAP1/data]		
36) SRS device buffer size (three times the aggregate of all log devices recommended)		
[256MB]		
37) SRS simple persistent queue directory		
[/work/SAP1/data]		
38) SRS simple persistent queue (SPQ) size (should be the size of the largest log device)		
[2000MB]		
39) HADR maintenance user		
Same on primary and companion sites		
[DR_maint]		
40) HADR maintenance password		
Same on primary and companion sites.		
Do not record your password.		

Value	Primary Site	Companion Site
42) RMA administrator user		
Same on primary and companion sites		
[DR_admin]		
42) RMA administrator password		
Same on primary and companion sites.		
Do not record your password.		
43) Databases to replicate		
Same on primary and companion sites.		
200 MB or larger for the new, required, cluster ID database.		
44) Host Name for the primary and secondary sites.		
[SFMACHINE1]		
[SJMACHINE2]		
45) Cockpit HTTP port (default is 4282)		
[4282]		
46) Cockpit HTTPS port (default is 4283)		
[4283]		
47) Cockpit TDS port (default is 4998)		
[4998]		
48) Cockpit RMI port (default is 4992)		
[4992]		
49) Cockpit administrator user		
Same on primary and companion sites		
[sccadmin]		
50) Cockpit administrator password		
Same on primary and companion sites.		
Do not record your password.		
51) Cockpit Agent administrator user		
[uafadmin]		
Same on primary and companion sites		

Value	Primary Site	Companion Site
52) Cockpit Agent administrator password		
Same on primary and companion sites.		
Do not record your password.		
53) Cockpit Repository password.		
Do not record your password.		
54) SAP ASE Install User (the operating system user who installed SAP ASE)		
55) SAP ASE System administrator password (same on primary and companion sites).		
Do not record your password.		
56) SAP Host Agent user		
[sapadm]		
57) SAP Host Agent password		
Same on primary and companion sites.		
Do not record your password.		
58) Fault Manager heartbeat to heartbeat port (default is 13797)		
Same on primary and companion sites		
[13797]		
59) Primary Fault Manager heartbeat port (default is 13777)		N/A
[13777]		
60) Companion Fault Manager heartbeat port (default is 13787)	N/A	
[13787]		

3.1.1 The SAP Host Agent

The HADR system uses the SAP Host Agent to perform several lifecycle management tasks, such as operating-system monitoring, database monitoring, system instance control, and provisioning.

You can install the SAP Host Agent when you install the HADR system (see *Unloading the SAP ASE Binaries*) or manually at another time. SAP Host Agent requires root or "sudo" permission to install. The SAP Host Agent needs to run as root to perform these operations:

- Monitoring the operating system and the databases.

- Starting, or restarting SAP ASE, Replication Server, and RMA, and so on for the Fault Manager by switching to the required user.
- Avoiding a split-brain condition by using the `kill` command to stop the SAP ASE process and, if configured in the Fault Manager profile, shuts down host.

Contact your system administrator if you do not have this permission.

See [Installing SAP Host Agent Manually](#) to manually install the SAP Host Agent.

i Note

Before you install Fault Manager on a third host, install SAP Host Agent on the hosts running SAP ASE, and set the `sapadm` operating system password.

Installing the SAP Host Agent manually requires a `.SAR` file, which is located in `archives` directory of the location where you extracted the installation image.

Starting and Stopping the SAP Host Agent

Starting and stopping the SAP Host Agent requires `sudo` or root privileges.

The SAP Host Agent is usually started automatically when you restart the operating system. You can also manually control it using the `saphostexec` program.

To start and stop the SAP Host Agent:

- (Windows)

```
%ProgramFiles%\SAP\hostctrl\exe\saphostexec.exe -[option] [pf=<ProfilePath>]
```

- (UNIX)

```
/usr/sap/hostctrl/exe/saphostexec -[option] [pf=<ProfilePath>]
```

Where `<ProfilePath>` is path to the profile file (`host_profile`) of SAP Host Agent. On UNIX, issue the `ps` command to determine the profile path if Fault Manager is running (see the bold text):

```
ps -ef|grep sap
root    11727      1  0 Dec11 ?          00:00:06 /usr/sap/hostctrl/exe/
saphostexec pf=/usr/sap/hostctrl/exe/host_profile
sapadm  11730      1  0 Dec11 ?          00:00:24 /usr/sap/hostctrl/exe/
sapstartsrv pf=/usr/sap/hostctrl/exe/host_profile -D
root    11764      1  0 Dec11 ?          00:02:54 /usr/sap/hostctrl/exe/saposcol -
l -w60 pf=/usr/sap/hostctrl/exe/host_profile
sap     24316 22274   0 12:18 pts/5    00:00:00 grep sap
```

By default the `host_profile` file is located in the executable directory. Option is one of:

- `-restart` – Starts or restarts SAP Host Agent
- `-stop` – Stops a running SAP Host Agent

See [SAP Host Agent Reference - Command Line Options of the saphostexec Executable](#) for a complete list of the `saphostexec` options.

3.2 Installing a New System

Perform the steps in this section to install a new HADR system.

See [Installing HADR with an Existing System \[page 86\]](#) for information about installing and configuring an HADR system using an existing SAP ASE server.

3.2.1 Unloading the SAP ASE Binaries

If it does not yet exist, the installer creates the target directory and installs the selected components into that directory. At the end of the installation, verify that the product has installed correctly. You may also need to perform additional configuration procedures.

Procedure

1. Insert the installation media in the appropriate drive, or download and extract the SAP ASE installation image from *Software Downloads* of the *SAP Support Portal* at <http://support.sap.com/swdc>.
2. (UNIX) Verify the stack size limit is at least 8192. To check the stack size limit, enter:
 - On the Bourne shell – `ulimit -s`
 - On the C-shell – `limit stacksize`

To set the stack size limit:

- On the Bourne shell – `ulimit -s 8192`
- On the C-shell – `limit stacksize 8192`

Make sure the `<LANG>` is set to C or any other value that is valid on your system. By default, the `<LANG>` environment variable on UNIX is set to POSIX, which can cause the installation to fail on the secondary system.

3. (AIX only) Set the data size limit to "unlimited":
 - On the Bourne shell – `ulimit -d unlimited`
 - On the C-shell – `limit datasize unlimited`
4. If you downloaded the product from SAP Service Marketplace, login as the "sybase" user, or the user you added with installation and configuration privileges, and change to the directory where you extracted the installation image.
5. Start the installer:

```
./setup.bin
```

6. If you are installing using a CD or DVD, mount the disk.

The location of the mount command is site-specific and may differ from the instructions shown here. If you cannot mount the drive using the path shown, check your operating system documentation or contact your system administrator.

Log in as root and issue:

```
# mount -t iso9660 /dev/cdrom /mnt/cdrom
```

i Note

mount commands and arguments vary according to the platforms. See the installation guide for your platform for more information.

7. Start the installer:

```
cd /mnt/<device_name>  
./setup.bin
```

Where

- `<device_name>` is the directory (mount point) you specified when mounting the CD or DVD drive.
- `setup.bin` is the name of the executable file name for installing SAP ASE.

Use the `-r` parameter to record your inputs in a response file when you run the SAP installer:

```
./setup -r <path_to_response_file>
```

For example:

```
./setup -r /work/SAP1_response_file.txt
```

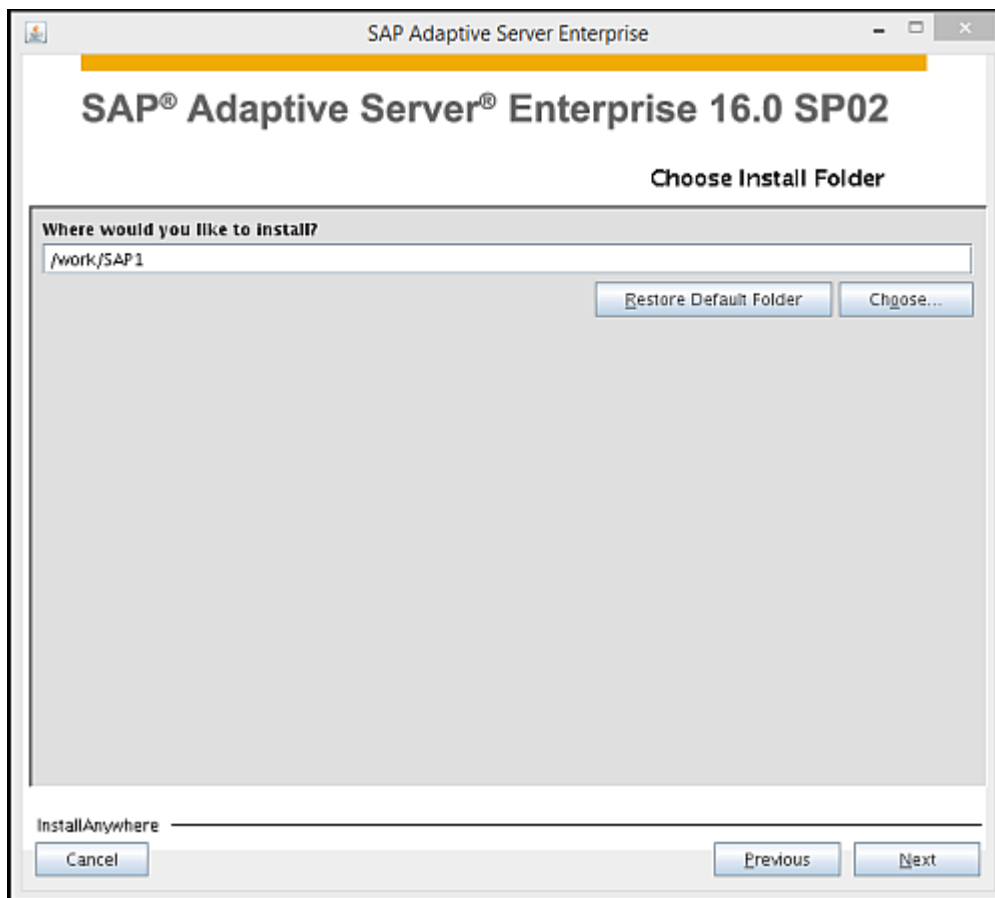
See [Installing the HADR System with Response Files, Console, and Silent Mode \[page 67\]](#).

If there is not enough disk space in the temporary disk space directory, set the `IATEMPDIR` environment variable to `<tmp_dir>` (`<TEMP>` on Windows) before running the installer again, where `<tmp_dir>` is where the installation program writes the temporary installation files. Include the full path to `<tmp_dir>`.

8. If you are prompted with a language selection list, specify the appropriate language.
9. On the Introduction screen, click [Next](#).



10. On the *Choose Install Folder* screen, accept the default directory or enter a new path to specify where to install SAP ASE.



This directory is your install directory, and the value for \$SYBASE. Make sure you have correct permissions—and sufficient space—to create the directories.


i Note

Do not use the following when specifying the installation path:

- Double-byte characters
- Single- or double-quote characters

The installer does not recognize these, and displays an error.

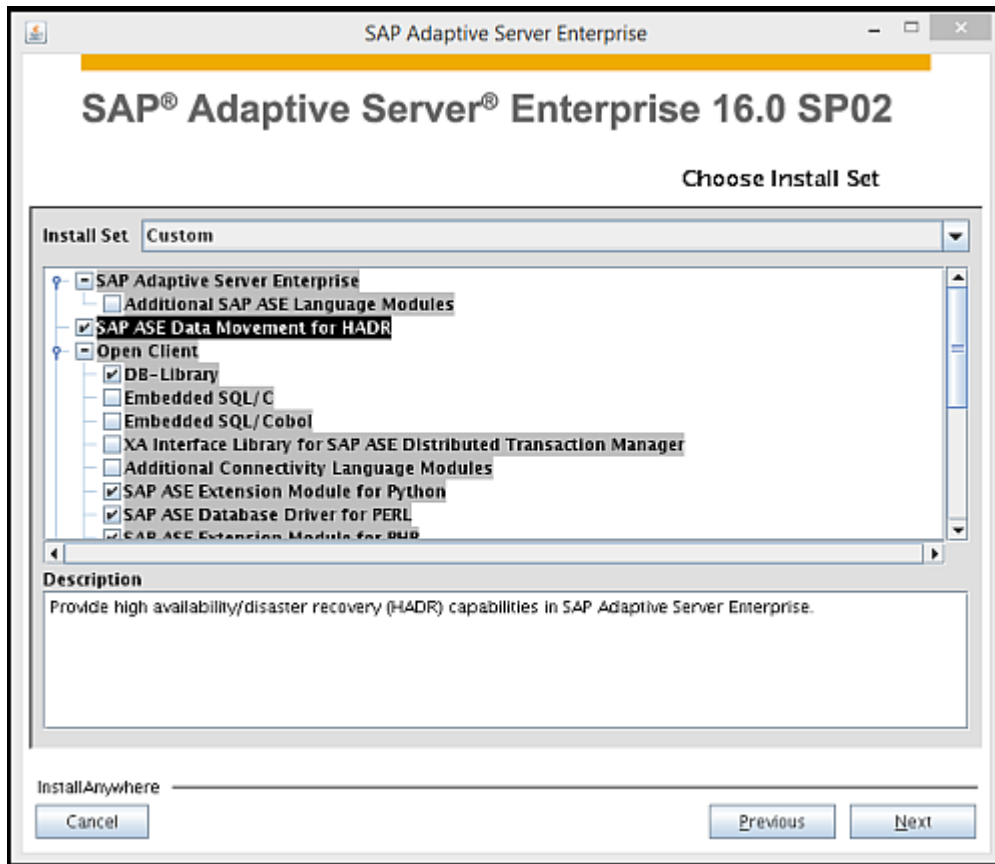
11. Click *Next*.
12. For upgrades only: If the installer displays the *Choose Update Installation* screen, there is an older version of SAP ASE in the directory you specified for the installation; perform an upgrade rather than a new installation.



13. In the *Choose Install Set* screen, select *Typical SAP ASE*, *database always-on option*, *Full*, or *Custom*. For an HADR installation, do not choose *Typical*.



14. If you selected *Custom*, include the *SAP ASE Data Movement for HADR* option with any additional items you select..



15. You can install the SAP Host Agent now or at another time (see [The SAP Host Agent \[page 28\]](#)).




On UNIX, enter and confirm the sudo password to install the SAP Host Agent.

i Note


The machine hosting the SAP Host Agent requires that you have sudo permission. The Fault Manager requires the SAP Host Agent and the sapadm operating system user created by the SAP Host Agent installation. If you do not have sudo permission, the system administrator can install SAP Host Agent later.

On Windows, enter a password for the sapadm operating system that adheres to the operating system password requirements, such as length, number of characters and digits, and so on.

16. Select the appropriate license type.



17. In the *End-user License Agreement* screen, select your geographical location from the drop-down list, review the license agreement, and accept it.



18. If you selected "Install licensed copy of SAP ASE", the installer prompts you to select the product license information:

- *Product Edition* – select the appropriate edition.
- *License Type* – select the appropriate license type.

SAP® Adaptive Server® Enterprise 16.0 SP02

Product Licenses

Please select the product edition and license type you would like to configure.

Product Edition: Enterprise Edition

License Type: CP: CPU License

InstallAnywhere _____

Cancel Previous Next

19. In the *SySAM Notification* screen, select Yes if you want to receive email alerts, and complete the following fields:

- *SMTP server host name* – name of the SMTP server.
- *SMTP server port number* – port number on which the SMTP server runs.
- *Sender email* – email address or account from which the alerts are sent. In the alert email, this is the “from” field.
- *Recipient emails* – comma-separated email address or distribution list of the people who will receive the alerts. In the alert email, this is the “to” field.
- *Message Severity for email alerts* – minimum severity that triggers email alerts.

SAP® Adaptive Server® Enterprise 16.0 SP02

SySAM Notification

Please configure the SySAM email alert mechanism. When configured, specified recipients will receive email notifications about SySAM events that may need administrator attention.

Do you want to configure email alerts?

Yes
 No

SMTP server host name

SMTP server port number

Sender email

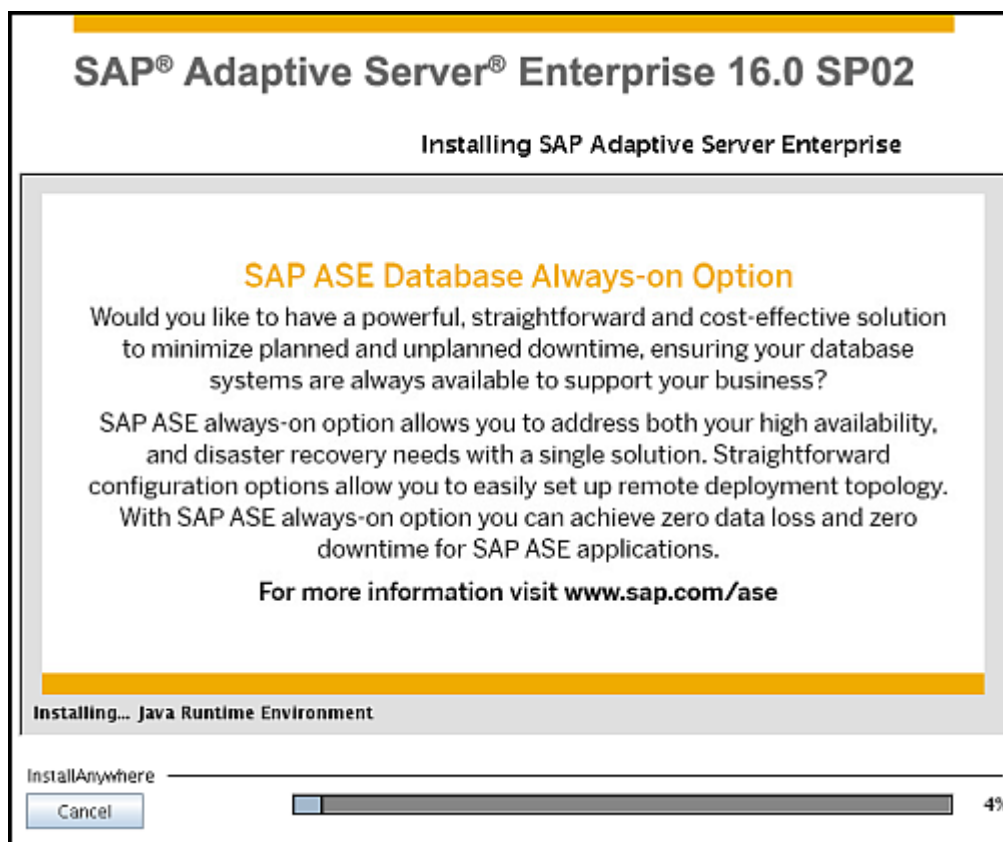
Recipient emails

Message Severity for email alerts ▼

InstallAnywhere _____

20. Review the preinstallation summary. Click [Previous](#) to make changes, or click [Next](#) for the installer to unload the files to the disk.

When the installation starts, you see a screen similar to the following, with a progress bar:



i Note

Installing the server can take several minutes.

To configure the HADR installation now, continue to [Configuring SAP ASE \[page 42\]](#).

To configure HADR at a later time, use the `$$SYBASE/$SYBASE_ASE/bin/setuphadr` utility. See [Installing HADR with an Existing System \[page 86\]](#).

3.2.2 Configuring SAP ASE

After you have installed the SAP ASE binaries, configure it for your HADR system (or, migrate an existing SAP ASE to your HADR system).

Prerequisites

These steps assume you have successfully completed the steps in the previous topic, [Unloading the SAP ASE Binaries \[page 30\]](#).

The HADR system requires Backup Server. The default language, character set, sort order, page size, system, user database sizes, and passwords for SAP ASE and Backup Server must be the same on the primary and

companion servers. See the *SAP ASE System Administration Guide, Volume 2 > Developing a Backup and Recovery Plan* for information about Backup Server.

Procedure

1. The Configure New Servers screen shows a list of all items you can minimally configure. By default, all products are selected. Verify the appropriate selections for your site are selected, and click [Next](#).



2. On the *Configure Servers with Different User Account* screen, indicate whether you are configuring the servers under a different user account.



The default is *No*. If you choose *Yes*, enter the user name and password:

SAP Adaptive Server Enterprise

SAP® Adaptive Server® Enterprise 16.0 SP02

Configure Servers with Different User Account

Do you want to configure the server(s) under different user account?

Yes

User name

User password

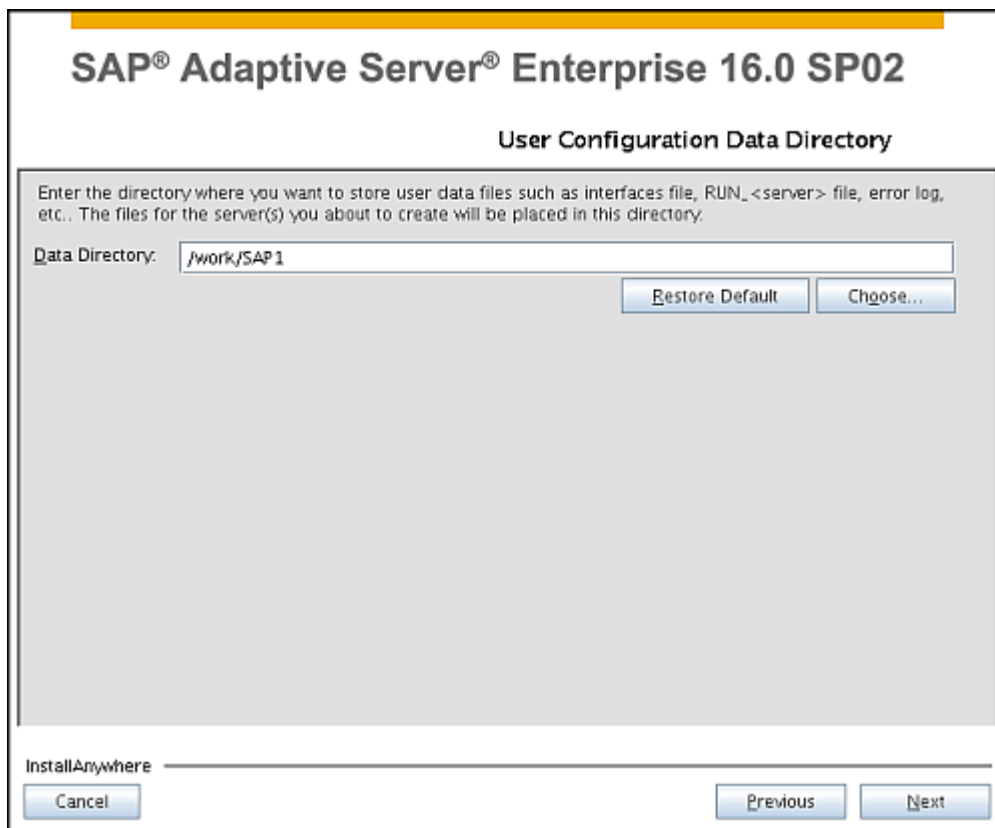
Confirm user password

No

InstallAnywhere _____

Cancel Previous Next

3. In the *User Configuration Data DirectoryNext*. This directory is your install directory, and the value mapped to \$SYBASE. Make sure you have correct permissions—and sufficient space—to create the directories.



4. Click *Next*.
5. In the Configure New SAP ASE screen, enter:

Option	Description
SAP ASE Name	Server name (do not include underscores in the name).
System Administrator's Password	Enter and confirm your password. Use the same value for both the primary and standby sites. screen, accept the default directory or enter a new path to specify where to install the SAP ASE binaries, then click
Enable SAP ASE for SAP ASE Cockpit Monitoring	Select to enable SAP ASE Cockpit to monitor SAP ASE.
Technical User	screen, accept the defaultSelect and confirm the technical user name and password if you are enabling SAP ASE Cockpit monitoring.
Port Number	SAP ASE port number. Default is 5000.
Error Log	Name and location of the error log file. Defaults to <code>servername.log</code> .
Application Type	(Must be the same on the primary and standby sites to be in sync with page size, default language, and so on) Select one: <ul style="list-style-type: none"> ○ (Default) MIXED – both OLTP and DSS. ○ OLTP – online transaction processing generally consists of smaller, less complex transactions. ○ DSS – decision-support systems generally have less update activity with large complex queries.

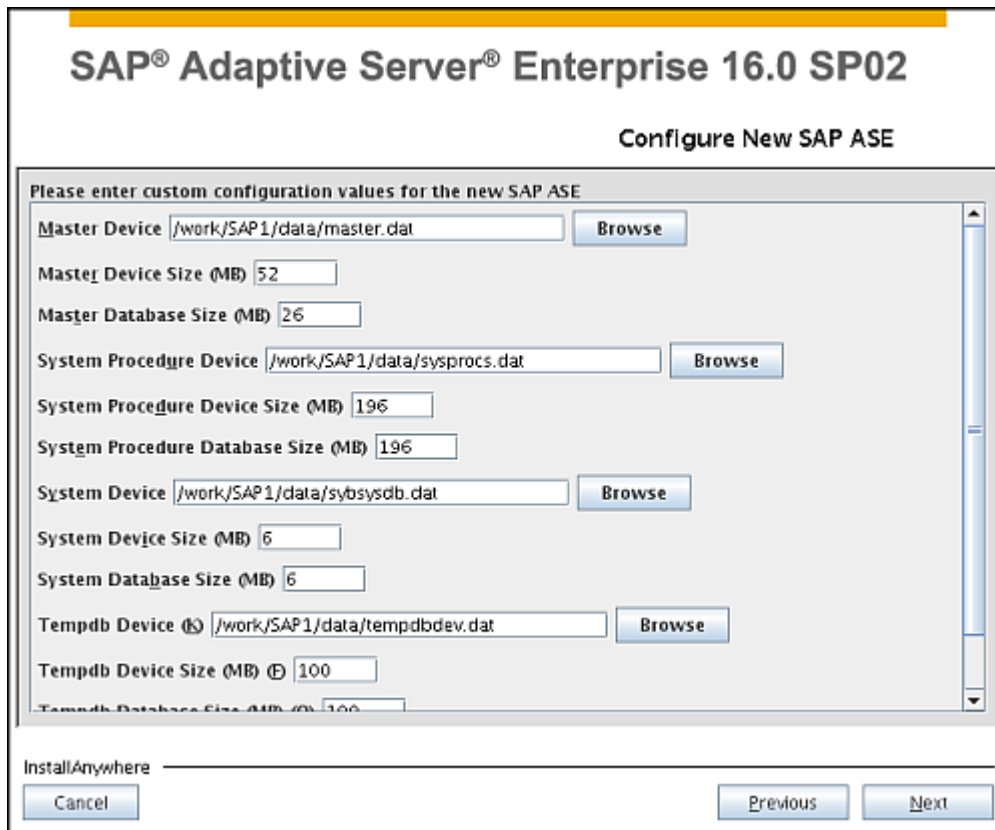
Option	Description
Page Size	<p>Must be the same on the primary and companion servers:</p> <ul style="list-style-type: none"> ○ 2 KB ○ (Default) 4 KB ○ 8 KB ○ 16 KB
Default Language	<p>Use the same value for both the primary and standby sites. The default is us-english.</p> <p>In Typical installations, only us-english is available. If you did not install additional language modules when you unloaded the binaries, only "us_english" is available.</p>
Default Character Set	<p>(Must be the same on the primary and standby sites) The default values are:</p> <ul style="list-style-type: none"> ○ roman8 – HP Itanium ○ cp850 – Windows ○ iso_1 – for other platforms
Default Sort Order	<p>(Must be the same on the primary and standby sites) The default values are:</p> <ul style="list-style-type: none"> ○ bin_roman8 – HP Itanium ○ bin_cp850 – Windows ○ bin_iso_1 – For other platforms
Optimize SAP ASE Configuration	<p>Check the box to optimize the configuration for your system.</p> <div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p>i Note</p> <p>If you specify a value that is larger than the available resource for allocation to the server, the optimize configuration may fail, causing the server to not start.</p> </div>
Create Sample Databases	<p>Select this option for the installer to install sample databases. The installer automatically calculates any additional space needed for your master device.</p>

Click [Next](#).

6. Enter values, or accept the defaults, for:

Configuration	Value
Master Device	The full path to the master device.
Master Device Size	The default is: <ul style="list-style-type: none"> ○ 2 KB – 36 MB ○ (Default) 4 KB – 52 MB ○ 8 KB – 104 MB ○ 16 KB – 208 MB
Master Database Size	The size of the master database: <ul style="list-style-type: none"> ○ 2 KB – 18 MB ○ (Default) 4 KB – 26 MB ○ 8 KB – 52 MB ○ 16 KB – 104 MB
System Procedure Device	The full path to the system procedure device.
System Procedure Device Size (MB)	The default is 196 MB, regardless of logical page size.
System Procedure Database Size	The default is 196 MB, regardless of logical page size.
System Device	The full path to the system device.
System Device Size (MB)	The default is:

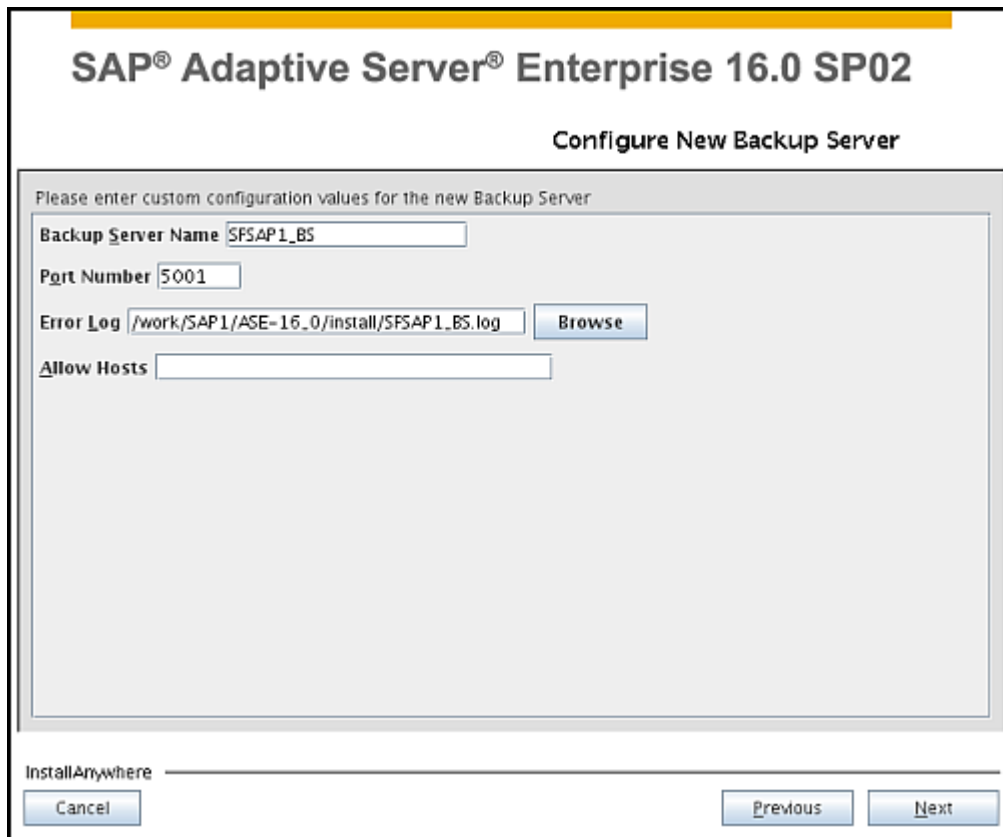
Configuration	Value
	<ul style="list-style-type: none"> ○ 2 KB – 3 MB ○ (Default) 4 KB – 6 MB ○ 8 KB – 12 MB ○ 16 KB – 24 MB
System Database Size	The default is: <ul style="list-style-type: none"> ○ 2 KB – 3 MB ○ (Default) 4 KB – 6 MB ○ 8 KB – 12 MB ○ 16 KB – 24 MB
Tempdb Device	The full path to the tempdb device.
Tempdb Device Size (MB)	The default is 100 MB, regardless of logical page size.
Tempdb Database Size	The default is 100 MB, regardless of logical page size.
Enable PCI	Select to enable PCI.



7. Enter values, or accept the defaults, for:

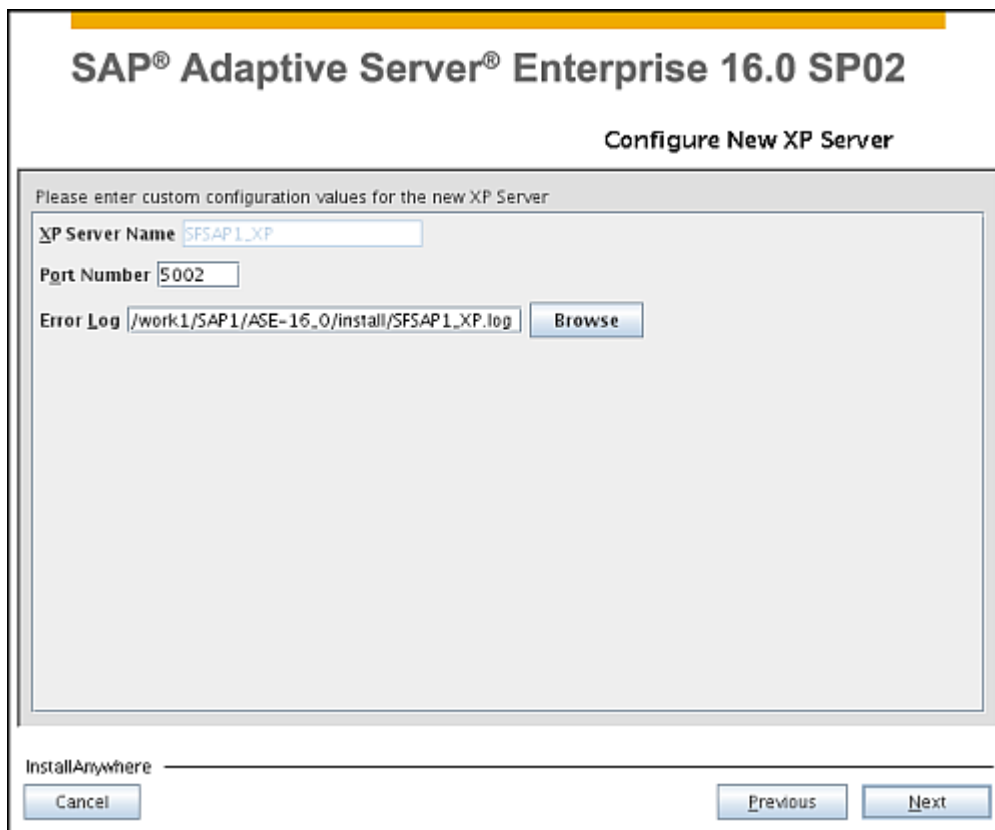
Configuration	Value
Backup Server Name	The default is <SAP_ASE_name>_BS.
Port Number	The port number of the Backup Server. A unique port number between 1025 and 65535. The default is 5001.

Configuration	Value
Error Log	The full path to the error log file.
Allow Hosts	(Unnecessary for HADR) Specify any remote hosts you want to use or to connect to this Backup Server. You can add primary and standby hosts, but this is not required because the installer updates the information when it builds the HADR system.



8. If you installed XP Server, enter values, or accept the defaults, for:

Configuration	Value
XP Server Name	The default is <SAP_ASE_name>_XP.
Port Number	The port number of the XP Server. The default is 5002.
Error Log	The full path to the error log file.



9. If you installed Job Scheduler, enter values, or accept the defaults, for:

Configuration	Value
Job Scheduler Agent Name	The default is <SAP_ase_name>_JSAGENT.
Port Number	The port number of the Job Scheduler Agent. The default is 4900.
Management Device	Full path to the device. The default is \$SYBASE/data/sybmgmt.db.dat.
Management Device Size (MB)	The default is 76 MB.
Management Database Size (MB)	The default is 76 MB.

SAP® Adaptive Server® Enterprise 16.0 SP02

Configure New Job Scheduler

Please enter custom configuration values for the new Job Scheduler

Job Scheduler Agent Name

Port Number

Management Device

Management Device Size (MB)

Management Database Size (MB)

InstallAnywhere _____

10. If you installed Job Scheduler, it requires the Self Management utility. Enter values, or accept the defaults, for:

Configuration	Value
Self Management User Name	The default is sa.
Self Management User Password	Enter and confirm.

SAP® Adaptive Server® Enterprise 16.0 SP02

Configure Self Management

Self Management allows the user to easily create a scheduled job that will run update statistics on any table, in any server database in the SAP ASE. Please enter the user name and password to be used in Self Management. If the user name does not already exist in the SAP ASE, one will be created.

Self Management User Name


Self Management User Password

Confirm Self Management User Password

InstallAnywhere

11. Enter values for:

Configuration	Value
Cluster ID	ID of the cluster (must be 3 characters long, starting with a letter).
Setup site	Choose Primary or companion.
Replication mode	Synchronous or asynchronous.



12. Enter values, or accept the defaults, for:

- *Site name* – Name of the site for HADR. Value must be unique on the primary and standby sites.
- *Database dump directory* – The directory Backup Server uses to perform dumps or loads when you materialize databases from the primary to the companion server.
- *RMA RMI port* – The port number for RMA RMI. The default is 7000.

i Note

RMA RMI occupies five consecutive ports, with the configured port occupying the highest number. If the configured RMA RMI port number is 7000, for example, it also needs ports 6999, 6998, 6997, and 6996.

- *RMA TDS port* – The port number for RMA TDS. The default is 7001.
- *Replication Server port* – The port number on which Replication Server talks to SAP ASE. The default is 5005.
- *SRS device buffer directory* – The directory in which you create the Replication Server buffer devices. The device buffer is comprised of inbound and outbound queues, which should be located on different file systems.
- *SRS device buffer size (MB)* – The size of the buffer device. The default is 256 MB (recommended three times the aggregate of all log devices).
- *SRS simple persistent queue directory* – The full path to the persistent queue.

Note

If you are configuring the HADR system with synchronous replication, SAP recommends that you specify a directory on an SSD (solid state drive) or other type of fast storage device for the Replication Server simple persistent queue directory.

- *SRS simple persistent queue size (MB)* – The size of the persistent queue. The default is 2000 MB.

SAP Adaptive Server Enterprise

SAP® Adaptive Server® Enterprise 16.0 SP02

ASE HADR on Primary Site

Enter SAP ASE, RMA, and SRS configuration values for primary site.

For simple persistent queue directory, it is strongly recommended you enter the directory to an SSD (solid state drive) or other type of fast read/write storage device.

Site name

Database dump directory:

RMA RMI port

RMA IDS port

Replication Server port

SRS device buffer directory:

SRS device buffer size (MB)

SRS simple persistent queue directory:

SRS simple persistent queue size (MB)

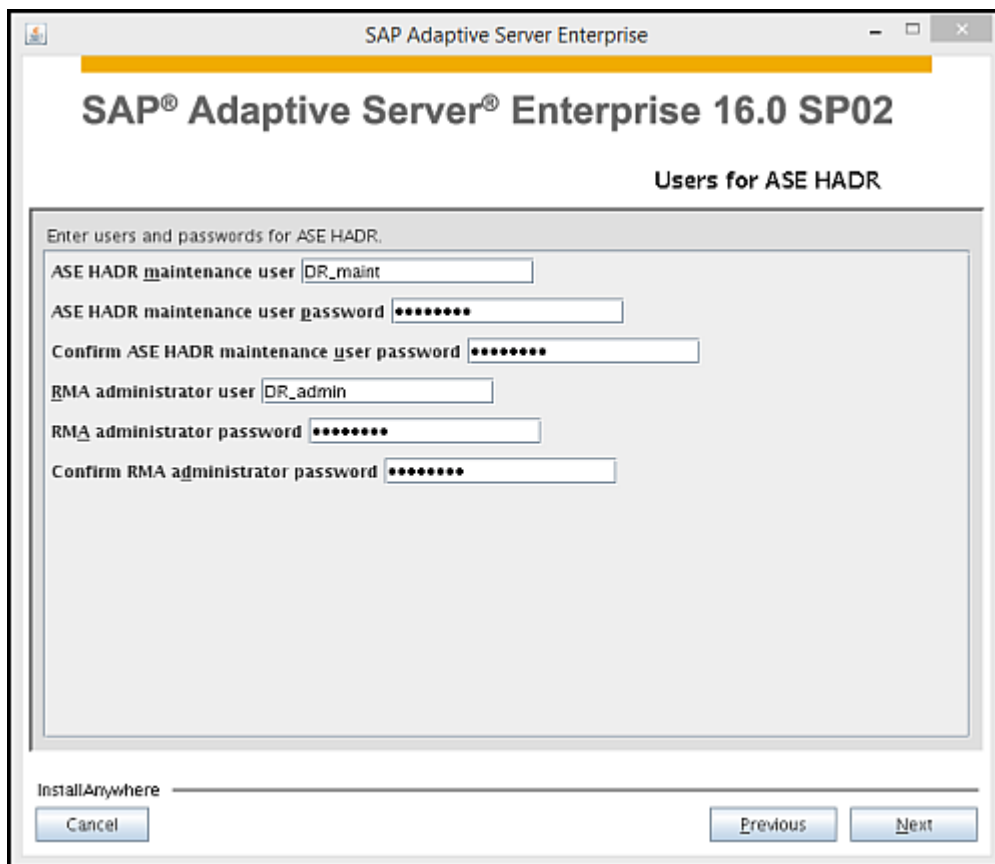
InstallAnywhere

13. Enter values for:

Note

User names and passwords must be the same on both primary and secondary companion servers. The user name must start with an alphabetic character and cannot exceed 30 characters in length. Passwords must have at least 6 characters.

- *ASE HADR maintenance user* – Name of the user replicating DML and DDL commands. The default is DR_maint.
- *ASE HADR maintenance user password* – Enter and confirm the user's password.
- *RMA administrator* – Name of the Replication Server administrator. DR_admin by default.
- *RMA administrator password* – Enter and confirm the administrator password.



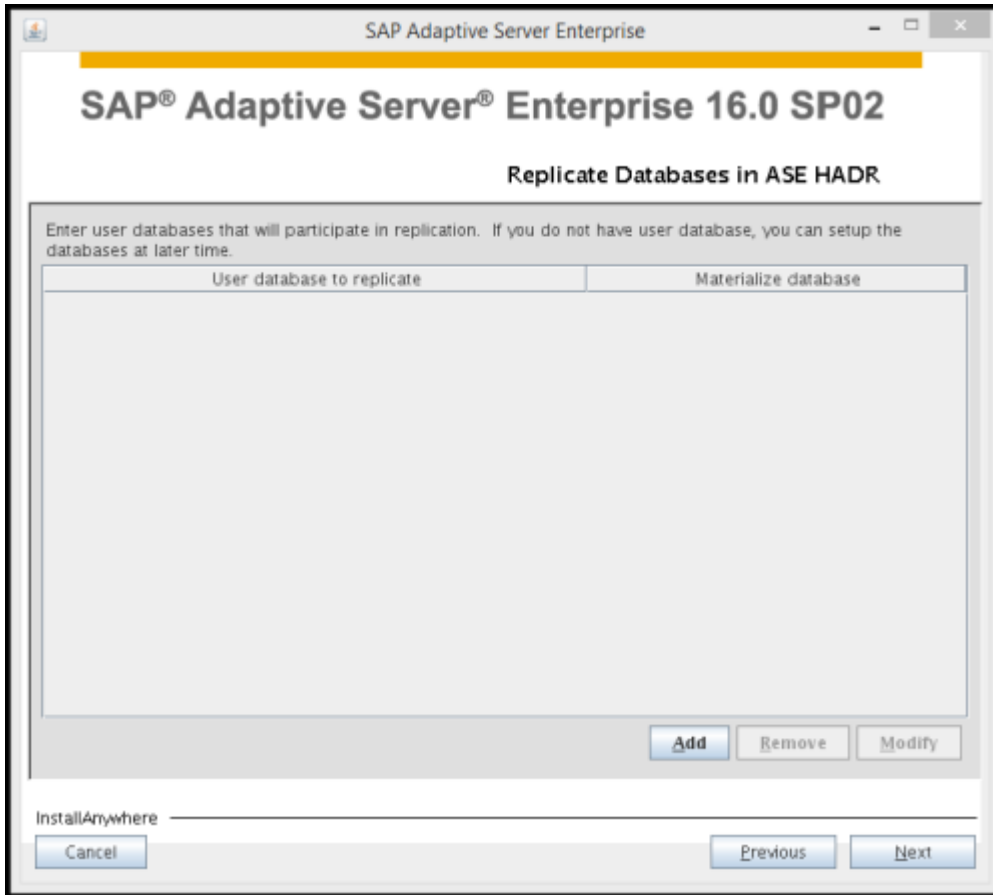
14. In the *Replicate Databases in ASE HADR* screen, enter new databases to be created and replicated, and select whether to materialize the databases as part of the HADR configuration. Use the same values for both primary and companion sites.

Click *Next* to create the databases manually, and then replicate and load the data later, or to add the databases to replicate and materialize later from SAP ASE Cockpit.

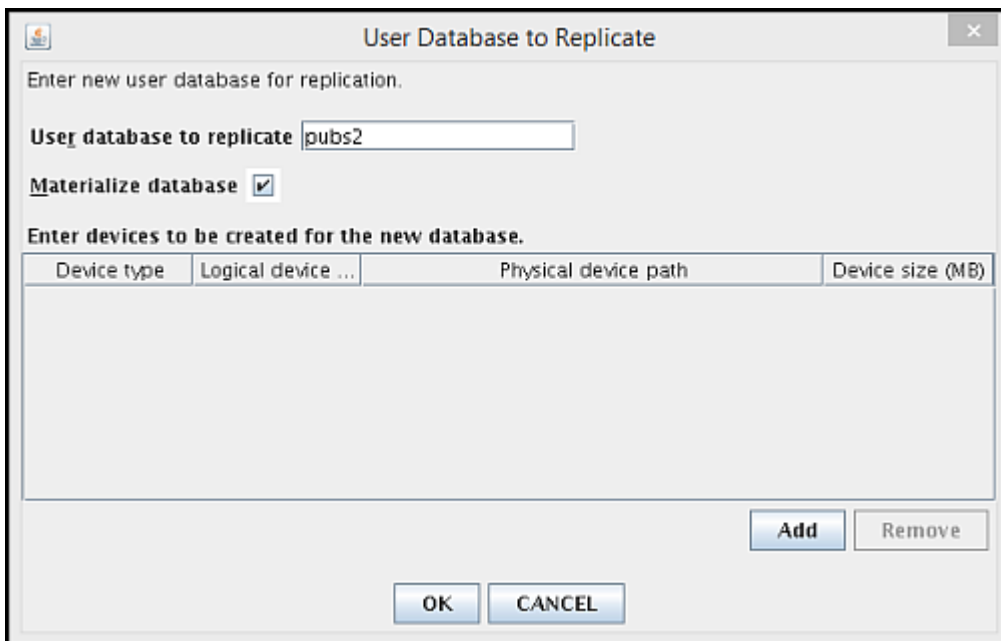
Note

The HADR system requires a cluster ID database. If you did not enter this database, the installer creates a 200 MB database by default. You should, however, create a database that is larger than this. The default size (based on the model database) is too small for an HADR system because some Fault Manager activities use the CID and generate a lot of log records. However, if you create a user database before creating the CID database, the size of the CID database uses an appropriate size for user data replication.

You can increase the size of any database after you finish the HADR setup with the `alter database` command (you may need to first create a new disk device with the `disk init` command), and you can add the databases to replicate and materialize later from SAP ASE Cockpit.



Do not enter the master database; it is automatically added as part of the HADR configuration. Clicking [Add](#) opens the User Database to Replicate screen.



15. In the User Database to Replicate screen, enter the values for:

- *User database to replicate* – The name of database you are replicating. This example uses the `sales` database.
- *Materialize database* – Select to include this database in materialization.

i Note

Do not select `sybmgmt.db` database. It is used for the Job Scheduler and should not be replicated.

- *Enter Devices to be created for the new database* – Click the *Add* button and enter:
 - *Device type* – Data or log device.
 - *Logical device* – Name of the logical device.
 - *Physical device path* – Full path to the device.
 - *Device size* – Size of device, in megabytes.

Click OK to return to the Replicate Databases in ASE HADR screen, then click *Next*.

Enter new user database for replication.

User database to replicate

Materialize database

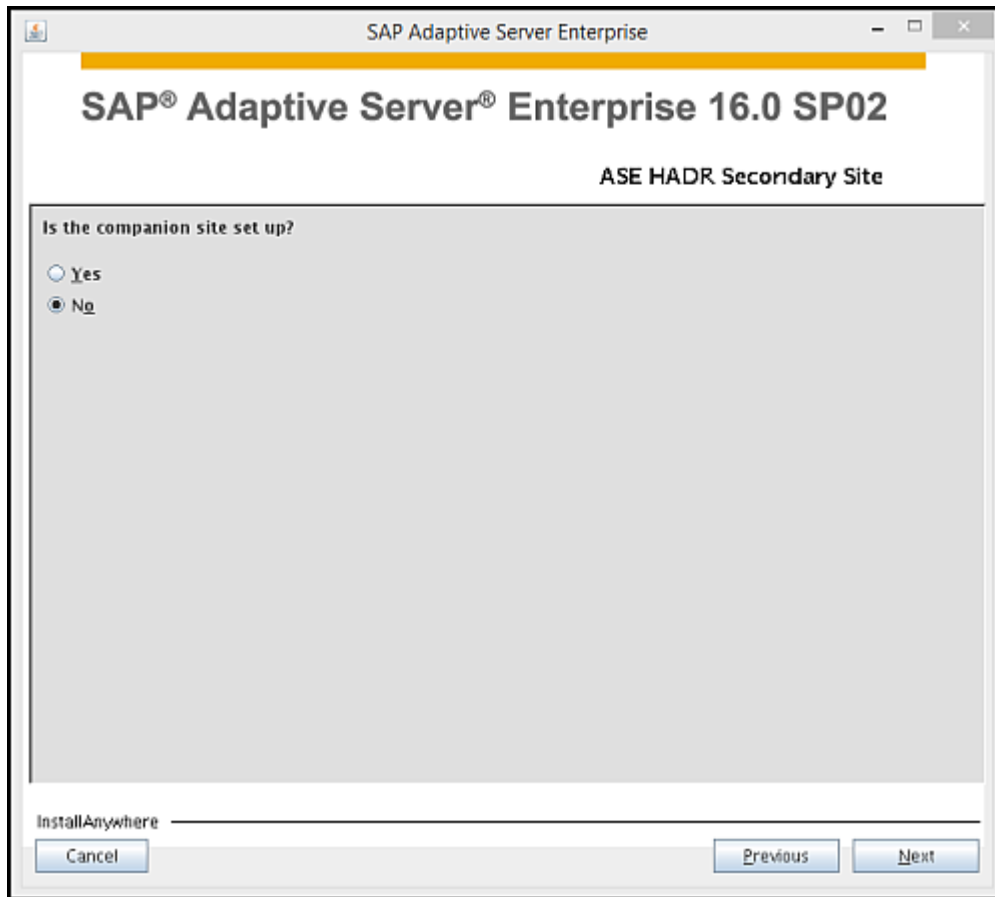
Enter devices to be created for the new database.

Device type	Logical device ...	Physical device path	Device size (MB)
data	salesdata	/work/SAP1/data/salesdata	500
log	saleslog	/work/SAP1/data/saleslog	50

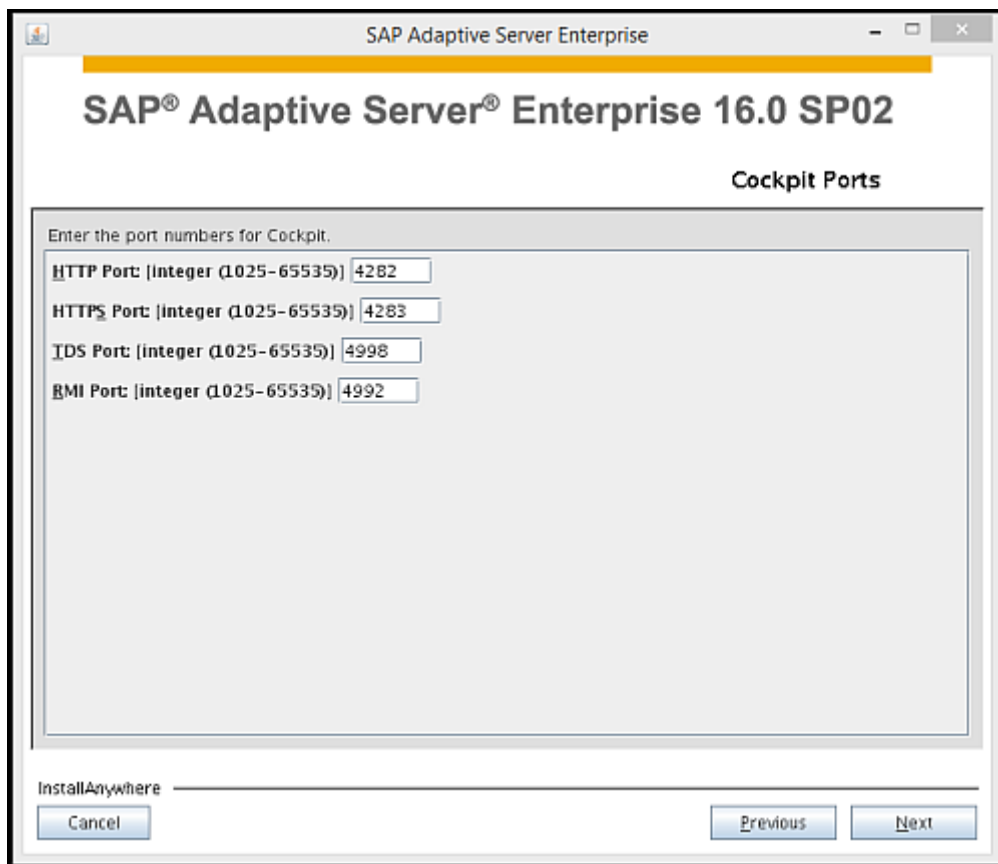
Add Remove

OK CANCEL

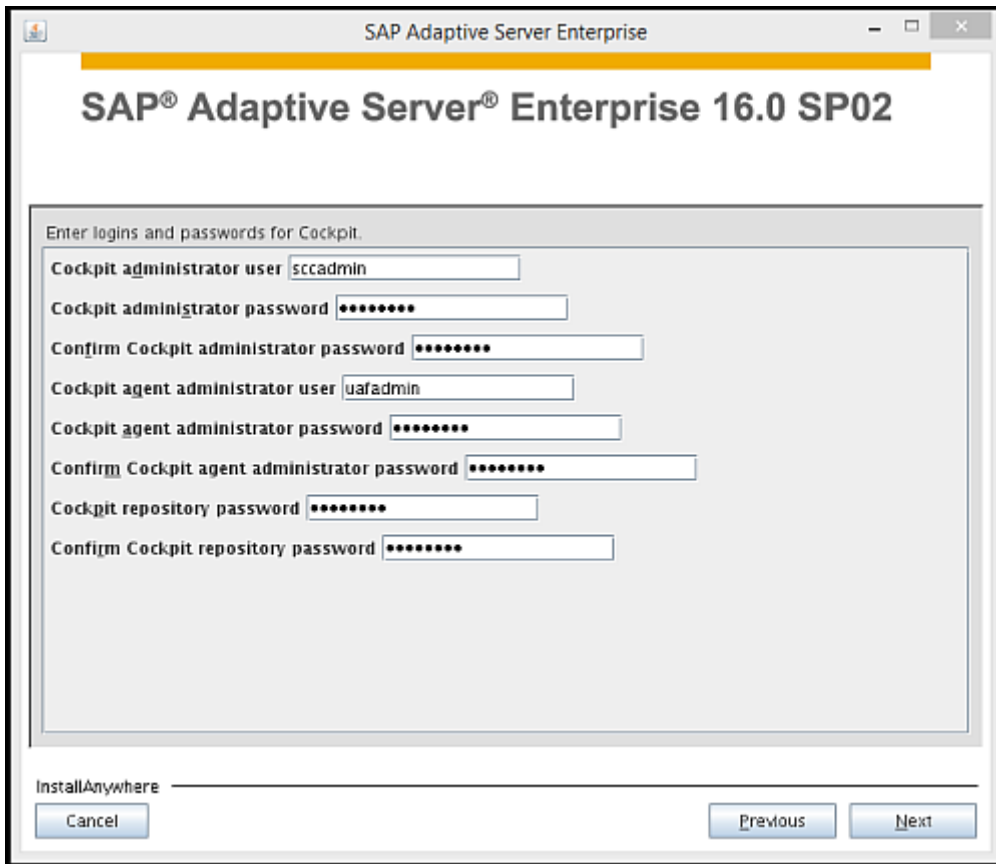
16. In the ASE HADR Secondary Site screen choose *Yes* if the companion site is up, or *No* if it is not. Click *Next*.



17. (If you select *No*) Provide values for cockpit configuration.
 - a. If you installed the SAP ASE Cockpit, enter the port numbers, or accept the defaults (port numbers cannot conflict with those used by other applications or services on your system):
 - *HTTP Port* – choose an integer between 1025 and 65535 (the default is 4282).
 - *HTTPS Port* – choose an integer between 1025 and 65535 (the default is 4283).
 - *TDS Port* – choose an integer between 1025 and 65535 (the default is 4998)
 - *RMI Port* – choose an integer between 1025 and 65535 (the default is 4992)



- b. Set and confirm SAP ASE Cockpit logins and passwords (use the same user names and passwords on the primary and standby servers):
- Cockpit administrator user and password
 - Cockpit agent administrator user and password
 - Cockpit repository password



18. (If you select *Yes*) Click *Next*:

- a. In the ASE HADR on Companion Site screen, enter the values for the servers you have set up on the secondary site:
 - *Site name* – The name of the site for the HADR system (value must be different from the name of the primary site).
 - *SAP ASE host name* – The name of the machine on which the secondary site SAP ASE is running.
 - *SAP ASE installation directory* – The directory where SAP ASE was installed.
 - *SAP ASE Name* – The name of the secondary server.
 - *SAP ASE port* – The number of the port on which the secondary Backup Server is listening.
 - *Backup Server Name* – The name of the Backup Server.
 - *Backup Server port* – The number of the port on which the Backup Server for the secondary server is listening.
 - *Database dump directory* – The default directory in which secondary server performs dumps.
 - *RMA RMI port* – The port number for the RMA RMI. The default is 7000.

i Note

RMA RMI occupies five consecutive ports, with the configured port occupying the highest number. If the configured RMA RMI port number is 7000, for example, it also needs ports 6999, 6998, 6997, and 6996.

- *RMA TDS port* – The port number for the RMA. The default is 7001.
- *Replication Server port* – The port number on which the secondary Replication Server talks to SAP ASE.

- *SRS device buffer directory* – The full path to the SRS device.
- *SRS device buffer size* –The size of the buffer device. The default is 256 MB.
- *SRS simple persistent queue directory* – The full path to the SRS persistent queue (you should specify a directory on an SSD (solid-state drive) or other type of fast storage device).
- *SRS simple persistent queue size* – The size of the SRS simple persistent queue. The default is 2000 MB.

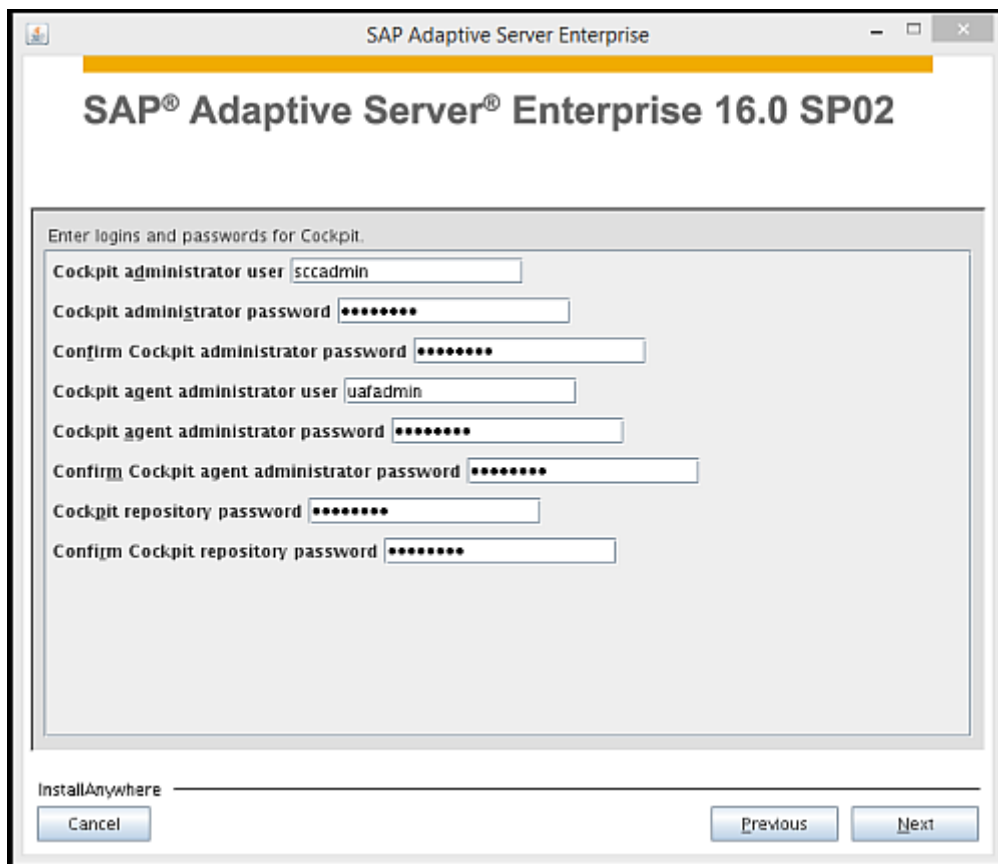
- b. If you installed the SAP ASE Cockpit, set the *Cockpit Hosts and Ports* option.

You can accept the default options, or specify other, unused ports, to ensure that the port numbers do not conflict with those used by other applications or services on your system, then click *Next*:

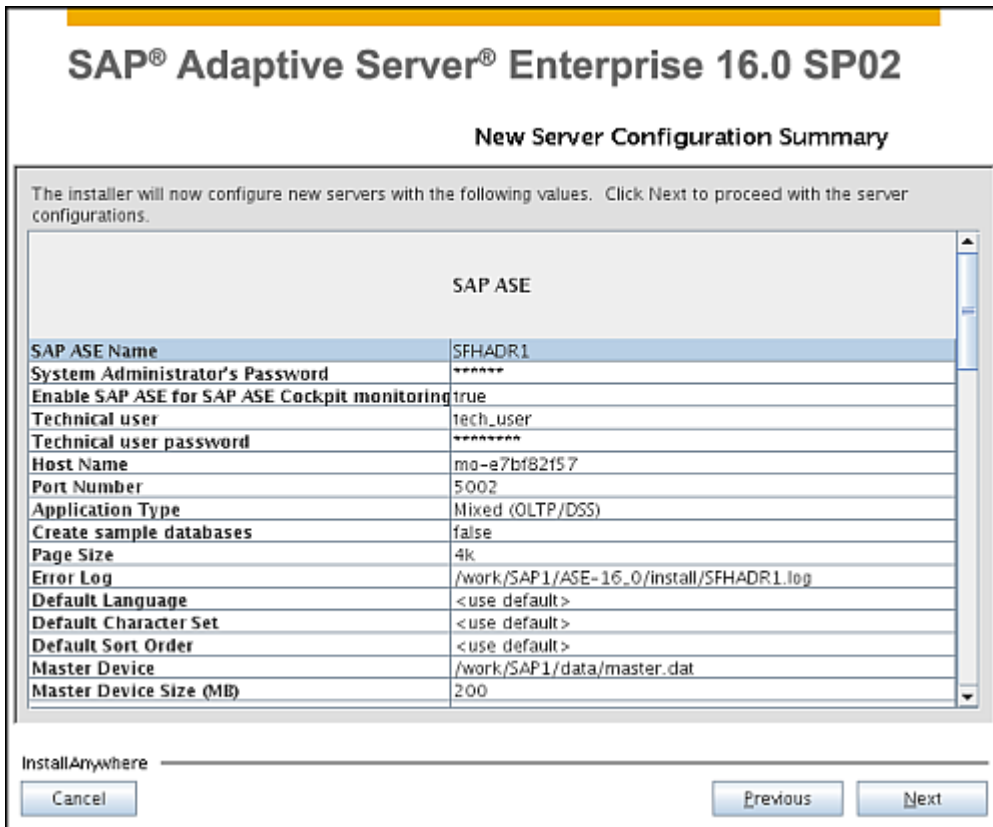
- *Host Name* – is the name of the machine on which you are installing cockpit.
- *HTTP Port* – choose an integer between 1025 and 65535 (the default is 4282).
- *HTTPS Port* – choose an integer between 1025 and 65535 (the default is 4283).
- *TDS Port* – choose an integer between 1025 and 65535 (the default is 4998).
- *RMI Port* – choose an integer between 1025 and 65535 (the default is 4992).

The screenshot shows a configuration window titled "SAP® Adaptive Server® Enterprise 16.0 SP02" with a subtitle "Cockpit Host and Ports". The window contains a text area with the instruction "Enter the host and port numbers for Cockpit." Below this are five input fields: "Host Name" with the value "SFMACHINE1", "HTTP Port: [integer (1025-65535)]" with the value "4282", "HTTPS Port: [integer (1025-65535)]" with the value "4283", "IDS Port: [integer (1025-65535)]" with the value "4999", and "RMI Port: [integer (1025-65535)]" with the value "4992". At the bottom of the window, there is a progress bar labeled "InstallAnywhere" and three buttons: "Cancel", "Previous", and "Next".

- c. Set and confirm SAP ASE Cockpit logins and passwords, then click *Next*:
- Cockpit administrator user and password
 - Cockpit agent administrator user and password
 - Cockpit repository password



19. Review the configuration summary. Click *Previous* to make changes or *Next* when you are ready to proceed with the configuration.



20. Click [Next](#) again to proceed with the configuration

When the installation is finished, you see the Installation Completed screen. Click [Done](#).



i Note

If you have not yet configured the secondary site, repeat the installation steps there to complete the SAP ASE database HADR configuration. The installer generated a response file named after the site you last created (`$SYBASE/log/companion_responses.txt` or `$SYBASE/log/primary_responses.txt`). Copy this to the other site and use it to configure the next server using the `-f` parameter.

The order does not matter for new installations, so you can install a companion or primary server first.

If both the primary and companion sites are configured, run the Fault Manager installer on a third host. The Fault Manager installer is located in `<ASE_installer>/FaultManager`. See [Installing and Configuring the Fault Manager \[page 115\]](#). The installer generated a response file named `$SYBASE/log/fault_manager_responses.txt`, which you use to install and configure the Fault Manager.

i Note

(UNIX only) The installation process installs the SAP Host Agent and creates the `sapadm` operating system login, but its login's password is not set. Run `sudo passwd sapadm` on the primary and companion sites to set this login's password before you install the Fault Manager.

3.2.3 Installing the HADR System with Response Files, Console, and Silent Mode

In addition to the SAP installer, you can use response files, and console and silent mode to install the HADR system.

3.2.3.1 Using setup.bin or setupConsole.exe with a Response File

Using `setup.bin` (`setupConsole.exe` on Windows) with a response file allows you to automate the HADR installation or install the Fault Manager.

The HADR installation `setup.bin` binary is located in `<installer_directory>/setup.bin`.

The Fault Manager installation `setup.bin` binary is located in `<installer_directory>/FaultManager/setup.bin`.

The syntax for using the response files is:

```
./setup.bin -f <response_file>
```

The `-f` parameter requires an absolute path.

Sample response files for the primary server, companion server, and the Fault Manager are located in:

- Primary server – `<installer_directory>/sample_response_primary.txt`
- Companion server – `<installer_directory>/sample_response_companion.txt`
- Fault Manager – `<installer_directory>/FaultManager/sample_response.txt`

In addition to these response files, you can generate your own by:

- Running the SAP installer with the `-r <response_file>` parameter to record your selections to a response file. The `-r` parameter requires an absolute path.
- Generating responses file in the `SYBASE/log` directory based on your input from the GUI installation for the primary and companion servers. However, these response files will be incomplete, and you cannot use these in a silent mode installation (the installation prompts you for missing information).

Edit each of these response files for your site, and run them with the `setup.bin` installer to install and configure the HADR system.

If you do not include the passwords in the response files, `setup.bin` prompts you for them during the console installation.

For example, if you edit the response file for HADR and rename it `SFASE1_response.txt`, the output looks similar to:

```
./setup.bin -f /work/SAP1/SFASE1_response.txt
Preparing to install...
Extracting the JRE from the installer archive...
Unpacking the JRE...
Extracting the installation resources from the installer archive...
Configuring the installer for this system's environment...
Launching installer...
```

```

Graphical installers are not supported by the VM. The console mode will be used
instead...
=====
SAP Adaptive Server Enterprise                               (created with InstallAnywhere)
-----
Preparing CONSOLE Mode Installation...
=====
Introduction
-----
InstallAnywhere will guide you through the installation of SAP Adaptive Server
Enterprise 16.0 SP02 GA.
It is strongly recommended that you quit all programs before continuing with
this installation.
Respond to each prompt to proceed to the next step in the installation.  If you
want to change something on a previous step, type 'back'.
You may cancel this installation at any time by typing 'quit'.
PRESS <ENTER> TO CONTINUE:

```

If you edit the Fault Manager response file for your site and rename it `FM_response.txt`, the output looks similar to:

```

./setup.bin -f /work/SAP1/FM_response.txt
Preparing to install...
Extracting the JRE from the installer archive...
Unpacking the JRE...
Extracting the installation resources from the installer archive...
Configuring the installer for this system's environment...
Launching installer...
Graphical installers are not supported by the VM. The console mode will be used
instead...
=====
Fault Manager                                             (created with InstallAnywhere)
-----
Preparing CONSOLE Mode Installation...
=====
Introduction
-----
InstallAnywhere will guide you through the installation of Fault Manager 1.0
GA.
It is strongly recommended that you quit all programs before continuing with
this installation.
Before you proceed, make sure that:
* SAP ASE, Replication Management Agent (RMA), Replication Server, and SAP Host
Agent are set up and running on the primary and companion sites.
* "sapadm" operating system user has a valid password on the primary and
companion sites.
Respond to each prompt to proceed to the next step in the installation.  If you
want to change something on a previous step, type 'back'.
You may cancel this installation at any time by typing 'quit'.
PRESS <ENTER> TO CONTINUE:

```

3.2.3.1.1 Sample Response Files

Edit the `sample_response_primary.txt` and `sample_response_companion.txt` response files to install HADR systems in your environment.

Below is a sample response file that uses the inputs from [Configuring SAP ASE \[page 42\]](#) to install the HADR system described there. The changed responses are in **bold**.

```

#####
# HADR sample responses file for SAP Adaptive Server Enterprise.

```

```

#
# This sample responses file installs and setup primary site for HADR.
#
# Prerequisite:
#   HADR is installed and setup on companion site.
#
# Following servers are configured:
#
#   SAP ASE (with HADR)
#   Backupserver
#   XP server
#   Job Scheduler (with Self Management)
#   ASE Cockpit
#
#####
#Validate Response File
#-----
#
RUN_SILENT=true
#Choose Install Folder
#-----
# USER_INSTALL_DIR=<Destination directory in absolute path>
#
USER_INSTALL_DIR=/work/SAP1
#Install older version
#-----
# INSTALL_OLDER_VERSION=<true|false>.
#   This determines whether installer can overwrite newer version.
#
INSTALL_OLDER_VERSION=false
#Choose Update Installation
#-----
# DO_UPDATE_INSTALL=<true|false>.
#   This determines if the installer selects and applies
#   updates to the installed product/features.
# DO_UPDATE_INSTALL_HADR_COMPONENT=<ASE|DM|ALL>
#   Which component to update for ASE HADR.
#   This is only valid if DO_UPDATE_INSTALL=true and
#   installed directory has ASE HADR.
#   Valid values are:
#   -----
#   ASE      --> Update only the SAP ASE components in rolling upgrade
#   DM       --> Update only the Data Movement component in rolling upgrade
#   ALL      --> Update all components
#
DO_UPDATE_INSTALL=false
DO_UPDATE_INSTALL_HADR_COMPONENT=DM
#Choose Install Set
#-----
# CHOSEN_INSTALL_SET=<Typical|TypicalASEHADR|Full|Custom>
# CHOSEN_FEATURE_LIST=<Features you want to install>
#   Valid values are:
#   -----
#   fase_srv          --> SAP Adaptive Server Enterprise
#   fase_add_lm       --> Additional SAP ASE Language Modules
#   fase_hadr         --> SAP ASE Data Movement for HADR
#
#       Available on:
#           Itanium/HP-UX 64-bit
#           IBM/AIX 64-bit
#           x86-64/Linux 64-bit
#           SPARC/Solaris 64-bit
#   fopen_client     --> Open Client
#   fdblib           --> DB-Library
#   fesql_c_lang     --> Embedded SQL/C
#   fesql_cobol_lang --> Embedded SQL/Cobol
#   fxa              --> XA Interface Library for SAP ASE Distributed
Transaction Manager
#       Available on:

```

```

# Itanium/HP-UX 64-bit
# IBM/AIX 64-bit
# x86-64/Linux 64-bit
# SPARC/Solaris 64-bit
# x86-64/Solaris 64-bit
# x86-64/Windows 64-bit
# fconn_add_lm --> Additional Connectivity Language Modules
# fjconnect160 --> jConnect 16.0 for JDBC
# fodbcl --> SAP ASE ODBC Driver
# fodata_ase --> OData Server for SAP ASE
# Available on:
# x86-64/Linux 64-bit
# x86-64/Windows 64-bit
# fdbisql --> Interactive SQL
# fqptune --> QPTune
# fsysam_util --> SySAM License Utilities
# fsysam_server --> SySAM License Server
# fscc_server --> Cockpit
# fasecmap --> SAP ASE Cockpit
# fase_cagent --> Remote Command and Control Agent for SAP ASE
# fconn_python --> SAP ASE extension module for Python
# fconn_perl --> SAP ASE database driver for PERL
# Available on:
# x86-64/Linux 64-bit
# x86-64/Windows 64-bit
# fconn_php --> SAP ASE extension module for PHP
#
# Notes:
# - If DO_UPDATE_INSTALL=true, CHOSEN_INSTALL_SET and CHOSEN_FEATURE_LIST
# are ignored.
# - If CHOSEN_INSTALL_SET is set to "Typical", "TypicalASEHADR", or "Full", do
not set
# CHOSEN_FEATURE_LIST.
#
#
CHOSEN_FEATURE_LIST=fase_srv,fopen_client,fdlib,fconn_python,fconn_perl,fconn_ph
p,fjconnect160,fodbcl,fdbisql,fqptune,fsysam_util,fscc_server,fasecmap,fase_cagen
t,fase_hadr
CHOSEN_INSTALL_SET=TypicalASEHADR
#SAP Host Agent
#-----
# Install SAP Host Agent for ASE HADR.
#
# You need root permission to install
# SAP Host Agent. Enter your password for
# installer to execute "sudo" command to
# install SAP Host Agent. If you do not
# have "sudo" permission, set
# INSTALL_SAP_HOST_AGENT=FALSE and ask
# your system administrator to manually
# install SAP Host Agent at later time.
#
# Notes:
# - You also can set SUDO_PASSWORD property value
# through SUDO_PASSWORD environment variable.
#
INSTALL_SAP_HOST_AGENT=TRUE
SUDO_PASSWORD=Secret396
#Choose Product License Type
#-----
# SYBASE_PRODUCT_LICENSE_TYPE=<license|evaluate|express>
# This is the End User License Agreement (EULA) you agreed to when run
# installer with "-DAGREE_TO_SAP_LICENSE=true" argument.
#
# Note:
# - 'evaluate' and 'express' only available on some platforms.
#
SYBASE_PRODUCT_LICENSE_TYPE=license

```

```

#Choose Sybase Software Asset Management License
#-----
# SYSAM_LICENSE_SOURCE=<license_file|existing_license_server|
proceed_without_license>
# SYSAM_LICENSE_FILE_PATHNAME=<license key file path>
#     Required when SYSAM_LICENSE_SOURCE is set to 'license_file'.
# SYSAM_EXISTING_LICENSE_SERVER_HOSTNAME=<license key server name>
#     Required when SYSAM_LICENSE_SOURCE is set to 'existing_license_server'.
# SYSAM_EXISTING_LICENSE_SERVER_PORTNUMBER=<license key server port number>
#     Set this to null for default port number.
#
SYSAM_LICENSE_SOURCE=license_file
SYSAM_LICENSE_FILE_PATHNAME=/work/SYSAM
SYSAM_EXISTING_LICENSE_SERVER_HOSTNAME=SFMACHINE1
SYSAM_EXISTING_LICENSE_SERVER_PORTNUMBER=27001
#Choose SYSAM Product Edition and License Type
#-----
# SYSAM_PRODUCT_EDITION=<Enterprise Edition|Small Business Edition|Unknown>
# SYSAM_LICENSE_TYPE=<License type>
#     Valid SYSAM_LICENSE_TYPE value for SYSAM_PRODUCT_EDITION='Enterprise
Edition':
#
#         CP : CPU License
#         SF : Standby CPU License
#         SR : Server License
#         SV : Standby Server License
#         DT : Development and Test License
#         EV : Evaluation License
#         OT : Other License
#         SS : Standalone Seat License
#         NA : Not Applicable or Other License
#         AC : OEM Application Deployment CPU License
#         BC : OEM Application Deployment Standby CPU License
#         AR : OEM Application Deployment Server License
#         BR : OEM Application Deployment Standby Server License
#         AO : OEM Application Deployment Other License
#         LP : Application Specific CPU License
#         LF : Application Specific Standby CPU License
#         LR : Application Specific Server License
#         LV : Application Specific Standby Server License
#         Unknown
#     Valid SYSAM_LICENSE_TYPE value for SYSAM_PRODUCT_EDITION='Small Business
Edition':
#
#         CP : CPU License
#         SF : Standby CPU License
#         SR : Server License
#         SV : Standby Server License
#         DT : Development and Test License
#         EV : Evaluation License
#         OT : Other License
#         SS : Standalone Seat License
#         NA : Not Applicable or Other License
#         AC : OEM Application Deployment CPU License
#         BC : OEM Application Deployment Standby CPU License
#         AR : OEM Application Deployment Server License
#         BR : OEM Application Deployment Standby Server License
#         AO : OEM Application Deployment Other License
#         LP : Application Specific CPU License
#         LF : Application Specific Standby CPU License
#         LR : Application Specific Server License
#         LV : Application Specific Standby Server License
#         CH : Chip License
#         SH : Standby Chip License
#         AH : Application Deployment Chip License
#         BH : Application Deployment Standby Chip License
#         LH : Application Specific Chip License
#         LI : Application Specific Standby Chip License
#         Unknown

```

```

#       Valid SYSAM_LICENSE_TYPE value for
SYSAM_PRODUCT_EDITION=SYSAM_PRODUCT_EDITION=Unknown
#       None
#
SYSAM_PRODUCT_EDITION=Enterprise Edition
SYSAM_LICENSE_TYPE=CP : CPU License
#Software Asset Management Notification Setting
#-----
# SYSAM_NOTIFICATION_ENABLE=<true|false>
#       Enable SysAM email notification
# SYSAM_NOTIFICATION_SMTP_HOSTNAME=<SMTP server host name>
#       Required if SYSAM_NOTIFICATION_ENABLE=true
# SYSAM_NOTIFICATION_SMTP_PORTNUMBER=<SMTP server port number>
#       Required if SYSAM_NOTIFICATION_ENABLE=true
# SYSAM_NOTIFICATION_SENDER_EMAIL=<Sender email>
#       Required if SYSAM_NOTIFICATION_ENABLE=true
# SYSAM_NOTIFICATION_RECIPIENT_EMAIL=<Recipient emails>
#       Required if SYSAM_NOTIFICATION_ENABLE=true
# SYSAM_NOTIFICATION_EMAIL_SEVERITY=<INFORMATIONAL|WARNING|ERROR>
#       Required if SYSAM_NOTIFICATION_ENABLE=true
#
SYSAM_NOTIFICATION_ENABLE=true
SYSAM_NOTIFICATION_SMTP_HOSTNAME=smtp
SYSAM_NOTIFICATION_SMTP_PORTNUMBER=25
SYSAM_NOTIFICATION_SENDER_EMAIL=SAP
SYSAM_NOTIFICATION_RECIPIENT_EMAIL=hadr_sa@bigcompany.com
SYSAM_NOTIFICATION_EMAIL_SEVERITY=WARNING
#Choose Update SAP ASE
#-----
# DO_UPDATE_ASE_SERVER=<true|false>
#       This property determines whether to update the existing SAP ASE.
#       It is only valid if DO_UPDATE_INSTALL=true.
# UPDATE_ASE_SERVER_NAME [n]=<SAP ASE name to update>
# UPDATE_ASE_PASSWORD [n]=<SAP ASE SA password>
#
# Notes:
# - You also can set UPDATE_ASE_SERVER_NAME [n] and UPDATE_ASE_PASSWORD [n]
#   property values through the environment variables UPDATE_ASE_SERVER_NAME [n]
#   and UPDATE_ASE_PASSWORD [n], respectively.
# - If the ASE password is null, set UPDATE_ASE_SERVER_NAME [n] value to "NA".
#
DO_UPDATE_ASE_SERVER=false
UPDATE_ASE_SERVER_NAME_1=
UPDATE_ASE_PASSWORD_1=
UPDATE_ASE_SERVER_NAME_2=
UPDATE_ASE_PASSWORD_2=
#Configure New Servers
#-----
# SY_CONFIG_ASE_SERVER=<true|false>
#       This property determines whether to configure SAP ASE.
# SY_CONFIG_HADR_SERVER=<true|false>
#       This property determines whether to setup ASE HADR.
#       Available on:
#           Itanium/HP-UX 64-bit
#           Power/AIX 64-bit
#           x86-64/Linux 64-bit
#           SPARC/Solaris 64-bit
# SY_CONFIG_BS_SERVER=<true|false>
#       This property determines whether to configure Backup Server.
# SY_CONFIG_XP_SERVER=<true|false>
#       This property determines whether to configure XP Server.
# SY_CONFIG_JS_SERVER=<true|false>
#       This property determines whether to configure Job Scheduler Agent.
# SY_CONFIG_SM_SERVER=<true|false>
#       This property determines whether to enable Self Management.
# SY_CONFIG_SCC_SERVER=<true|false>
#       This property determines whether to configure Cockpit.
#

```



```

# Notes:
# - These properties are ignored if you set DO_UPDATE_INSTALL=true.
#   See above for updating the existing ASE servers.
# - If SY_CONFIG_BS_SERVER, SY_CONFIG_XP_SERVER, and/or SY_CONFIG_JS_SERVER,
#   are set to "true", SY_CONFIG_ASE_SERVER must also set to "true".
# - If SY_CONFIG_HADR_SERVER=true, SY_CONFIG_ASE_SERVER and SY_CONFIG_BS_SERVER
#   must also set to "true".
# - If SY_CONFIG_SM_SERVER=true, SY_CONFIG_JS_SERVER must also set to "true".
#
SY_CONFIG_ASE_SERVER=true
SY_CONFIG_HADR_SERVER=true
SY_CONFIG_BS_SERVER=true
SY_CONFIG_XP_SERVER=false
SY_CONFIG_JS_SERVER=false
SY_CONFIG_SM_SERVER=false
SY_CONFIG_SCC_SERVER=true
#Configure Servers with Different User Account
#-----
# If SY_CFG_USER_ACCOUNT_CHANGE=yes, below properties are required:
#
# SY_CFG_USER_ACCOUNT_NAME=<user name>
# SY_CFG_USER_ACCOUNT_PASSWORD=<user password>
#
SY_CFG_USER_ACCOUNT_CHANGE=no
SY_CFG_USER_ACCOUNT_NAME=
SY_CFG_USER_ACCOUNT_PASSWORD=
#User Configuration Data Directory
#-----
SY_CFG_USER_DATA_DIRECTORY=
#Configure New SAP ASE
#-----
# If SY_CONFIG_ASE_SERVER=true, below properties are required:
#
# SY_CFG_ASE_SERVER_NAME=<SAP ASE name>
# SY_CFG_ASE_PORT_NUMBER=<SAP ASE port number>
# SY_CFG_ASE_APPL_TYPE=<MIXED|OLTP|DSS>
#   This property describes the intended use of the ASE server:
#       OLTP    --> Online Transaction Processing
#       DSS     --> Decision Support Systems
#       Mixed   --> OLTP/DSS
# SY_CFG_ASE_PAGESIZE=<2k|4k|8k|16k>
#   SAP ASE page size
# ASE ADDL_CMD_ARG=<Additional command line arguments to ASE>
# SY_CFG_ASE_PASSWORD=<SA's password>
# SY_CFG_ASE_MASTER_DEV_NAME=<Master device name>
# SY_CFG_ASE_MASTER_DEV_SIZE=<Master device size in MB>
# SY_CFG_ASE_MASTER_DB_SIZE=<Master database size in MB>
# SY_CFG_ASE_SYBPROC_DEV_NAME=<System procedure device name>
# SY_CFG_ASE_SYBPROC_DEV_SIZE=<System procedure device size in MB>
# SY_CFG_ASE_SYBPROC_DB_SIZE=<System procedure database size in MB>
# SY_CFG_ASE_SYBTEMP_DEV_NAME=<System device name for two phase commit>
# SY_CFG_ASE_SYBTEMP_DEV_SIZE=<System device size in MB>
# SY_CFG_ASE_SYBTEMP_DB_SIZE=<System database size in MB>
# SY_CFG_ASE_ERROR_LOG=<Error log>
# CFG_REMOTE_AND_CONTROL_AGENT=<true|false>
#   Configures Remote Command and Control Agent for SAP ASE.
#   This is only valid if Remote Command and Control Agent
#   for SAP ASE is installed.
# ENABLE_COCKPIT_MONITORING=<true|false>
#   Enable SAP ASE for SAP ASE Cockpit monitoring.
#   This is only valid if SAP ASE Cockpit is installed.
# COCKPIT_TECH_USER=<SAP ASE Cockpit technical user>
#   Required if ENABLE_COCKPIT_MONITORING=true.
# COCKPIT_TECH_USER_PASSWORD=<SAP ASE Cockpit technical user password>
#   Required if ENABLE_COCKPIT_MONITORING=true.
# SY_CFG_ASE_PCI_ENABLE=<true|false>
#   This property determines if Java in database feature is enabled.
# SY_CFG_ASE_PCI_DEV_NAME=<Device name for "sybpcidb" database>

```

```

#         Required if SY_CFG_ASE_PCI_ENABLE=true
# SY_CFG_ASE_PCI_DEV_SIZE=<PCI device size in MB>
#         Required if SY_CFG_ASE_PCI_ENABLE=true
# SY_CFG_ASE_PCI_DB_SIZE=<PCI database size in MB>
#         Required if SY_CFG_ASE_PCI_ENABLE=true
# SY_CFG_ASE_TEMP_DEV_NAME=<Device for temporary database>
# SY_CFG_ASE_TEMP_DEV_SIZE=<Temporary device size in MB>
# SY_CFG_ASE_TEMP_DB_SIZE=<Temporary database size in MB>
# SY_CFG_ASE_OPT_ENABLE=<true|false>
#         This property determines if it should optimized the SAP ASE.
# SY_CFG_ASE_CPU_NUMBER=<Available physical memory for ASE in MB>
#         Required if SY_CFG_ASE_OPT_ENABLE=true
# SY_CFG_ASE_MEMORY=<Available number of CPU for ASE>
#         Required if SY_CFG_ASE_OPT_ENABLE=true
# SY_CFG_ASE_LANG=<Default language in ASE>
#         Set this to null to use default language.
# SY_CFG_ASE_CHARSET=<Default character set in ASE>
#         Set this to null to use default character set.
# SY_CFG_ASE_SORTORDER=<Default sort order in ASE>
#         Set this to null to use default sort order.
# SY_CFG_ASE_SAMPLE_DB=<true|false>
#         Determines whether to create pubs2 and pubs3 sample databases.
#
SY_CFG_ASE_SERVER_NAME=SFSAP1
SY_CFG_ASE_PORT_NUMBER=5000
SY_CFG_ASE_APPL_TYPE=MIXED
SY_CFG_ASE_PAGESIZE=4k
ASE_ADDL_CMD_ARG=
SY_CFG_ASE_PASSWORD=Secret396
SY_CFG_ASE_MASTER_DEV_NAME=/work/SAP1/data/master.dat
SY_CFG_ASE_MASTER_DEV_SIZE=52
SY_CFG_ASE_MASTER_DB_SIZE=26
SY_CFG_ASE_SYBPROC_DEV_NAME=/work/SAP1/data/sysprocs.dat
SY_CFG_ASE_SYBPROC_DEV_SIZE=196
SY_CFG_ASE_SYBPROC_DB_SIZE=196
SY_CFG_ASE_SYBTEMP_DEV_NAME=/work/SAP1/data/sybsysdb.dat
SY_CFG_ASE_SYBTEMP_DEV_SIZE=6
SY_CFG_ASE_SYBTEMP_DB_SIZE=6
SY_CFG_ASE_ERROR_LOG=/work/SAP1/ASE-16_0/install/SYBASE.log
CFG_REMOTE_AND_CONTROL_AGENT=false
ENABLE_COCKPIT_MONITORING=true
COCKPIT_TECH_USER=tech_user
COCKPIT_TECH_USER_PASSWORD=Secret396
SY_CFG_ASE_PCI_ENABLE=false
SY_CFG_ASE_PCI_DEV_NAME=
SY_CFG_ASE_PCI_DEV_SIZE=
SY_CFG_ASE_PCI_DB_SIZE=
SY_CFG_ASE_TEMP_DEV_NAME=/work/SAP1/data/tempdbdev.dat
SY_CFG_ASE_TEMP_DEV_SIZE=100
SY_CFG_ASE_TEMP_DB_SIZE=100
SY_CFG_ASE_OPT_ENABLE=false
SY_CFG_ASE_CPU_NUMBER=
SY_CFG_ASE_MEMORY=
SY_CFG_ASE_LANG=
SY_CFG_ASE_CHARSET=
SY_CFG_ASE_SORTORDER=
SY_CFG_ASE_SAMPLE_DB=false
#Configure New Backup Server
#-----
# If SY_CONFIG_BS_SERVER=true, below properties are required:
#
# SY_CFG_BS_SERVER_NAME=<Backup Server name>
# SY_CFG_BS_PORT_NUMBER=<Backup Server port number>
# SY_CFG_BS_ERROR_LOG=<Backup Server error log>
# SY_CFG_BS_ALLOW_HOSTS=<Comma separated list of hosts allow to connect to this
Backup Server>
# BS_ADDL_CMD_ARG=<Additional command line arguments to Backup Server>
#

```

```

SY_CFG_BS_SERVER_NAME=SFSAP1_BS
SY_CFG_BS_PORT_NUMBER=5001
SY_CFG_BS_ERROR_LOG=/work/SAP1/ASE-16_0/install/SFSAP1_BS.log
SY_CFG_BS_ALLOW_HOSTS=
BS_ADDL_CMD_ARG=
#Configure New XP Server
#-----
# If SY_CONFIG_XP_SERVER=true, below properties are required:
#
# SY_CFG_XP_SERVER_NAME=<XP Server name>
#     The name must be "<ASE server name (SY_CFG_ASE_SERVER_NAME) in
uppercase>_XP" format.
# SY_CFG_XP_PORT_NUMBER=<XP Server port number>
# SY_CFG_XP_ERROR_LOG=<XP Server error log>
#
SY_CFG_XP_SERVER_NAME=SFSAP1_XP
SY_CFG_XP_PORT_NUMBER=5002
SY_CFG_XP_ERROR_LOG=/opt/sap/ASE-16_0/install/SFSAP1_XP.log
#Configure New Job Scheduler
#-----
# If SY_CONFIG_JS_SERVER=true, below properties are required:
#
# SY_CFG_JS_SERVER_NAME=<Job Scheduler Agent name>
# SY_CFG_JS_PORT_NUMBER=<Job Scheduler Agent port number>
# SY_CFG_JS_MANAG_DEV_NAME=<SAP ASE management device name>
# SY_CFG_JS_MANAG_DEV_SIZE=<SAP ASE management device size in MB>
# SY_CFG_JS_MANAG_DB_SIZE=<SAP ASE management database size in MB>
#
SY_CFG_JS_SERVER_NAME=SYBASE_JS
SY_CFG_JS_PORT_NUMBER=4900
SY_CFG_JS_MANAG_DEV_NAME=/opt/sap/data/sybmgmtdb.dat
SY_CFG_JS_MANAG_DEV_SIZE=76
SY_CFG_JS_MANAG_DB_SIZE=76
#Configure Self Management
#-----
# If SY_CONFIG_SM_SERVER=true, below properties are required:
#
#
# SY_CFG_SM_USER_NAME=<Self Management user name>
# SY_CFG_SM_PASSWORD=<Self Management user password>
#SY_CFG_ASE_PCI_ENABLE
SY_CFG_SM_USER_NAME=sa
SY_CFG_SM_PASSWORD=Secret396
#ASE HADR Setup Site
#-----
# If SY_CONFIG_HADR_SERVER=TRUE, below properties are required:
#
# HADR_CLUSTER_ID=<ID that identifies this cluster. Value must be 3 characters
in length>
# HADR_SETUP_SITE=<primary|companion>
#     Which site being configured
# HADR_SYNC_MODE=<sync|async>
#     How data is replicated
#
HADR_CLUSTER_ID=AS1
HADR_SETUP_SITE=primary
HADR_SYNC_MODE=sync
#ASE HADR on Primary Site
#-----
# ASE and SRS servers on this (primary) site.
#
HADR_PRIMARY_ASE_HOST_NAME=SFMACHINE1
HADR_PRIMARY_SITE_NAME=SFHADR1
HADR_PRIMARY_ASE_SERVER_NAME=SFSAP1
HADR_PRIMARY_ASE_SERVER_PORT=5000
HADR_PRIMARY_BS_SERVER_NAME=SFSAP1_BS
HADR_PRIMARY_BS_SERVER_PORT=5001
HADR_PRIMARY_BS_DUMP_DIR=/work/SAP1/data

```

```

HADR_PRIMARY_RMA_RMI_PORT=7000
# RMA RMI occupies five consecutive ports, with the configured port occupying
the highest number.
HADR_PRIMARY_RMA_TDS_PORT=7001
HADR_PRIMARY_SRS_PORT=5005
HADR_PRIMARY_SRS_DEVICE_BUFFER_DIR=/work/SAP1/data
HADR_PRIMARY_SRS_DEVICE_BUFFER_SIZE=256
# For synchronous replication, recommend path to an SSD
# (solid state drive) or other type of fast read/write storage device.
HADR_PRIMARY_SRS_PERSISTENT_QUEUE_DIR=/ssd/data
HADR_PRIMARY_SRS_PERSISTENT_QUEUE_SIZE=2000
#Users for ASE HADR
#-----
# Logins for HADR. These properties values must be same
# on both primary and companions sites.
#
# HADR_MAINTENANCE_USER=<ASE HADR maintenance user>
# HADR_MAINTENANCE_PASSWORD=<ASE HADR maintenance user password>
# RMA_ADMIN_USER=<RMA administrator user>
# RMA_ADMIN_PASSWORD=<RMA administrator password>
#
ASE_SA_USER=sa
ASE_SA_PASSWORD=Secret396
HADR_MAINTENANCE_USER=DR_maint
HADR_MAINTENANCE_PASSWORD=Secret396
RMA_ADMIN_USER=DR_admin
RMA_ADMIN_PASSWORD=Secret396
#Replicate Databases in ASE HADR
#-----
# User databases to create and setup for replication.
# The database names and sizes must be same on both
# primary and companions sites.
#
# HADR_USER_DB [x]=<Databases that will participate in replication>
# HADR_USER_DB_MATERIALIZE [x]=<"auto" materialize database to companion>
#
# HADR_COMPANION_DATA_DEVICE_CREATE [x] [y]=<logical device>,<physical device
path>,<device size in MB>
# HADR_COMPANION_LOG_DEVICE_CREATE [x] [y]=<logical device>,<physical device
path>,<device size in MB>
#
# HADR_USER_DB_1=sales
# HADR_USER_DB_MATERIALIZE_1=true
# HADR_COMPANION_DATA_DEVICE_CREATE_1_1=sales_data_dev1, /work/SAP1/data/
sales1.dat, 50
# HADR_COMPANION_DATA_DEVICE_CREATE_1_2=sales_data_dev2, /work/SAP1/data/
sales2.dat, 50
# HADR_COMPANION_LOG_DEVICE_CREATE_1_1=sales_log_dev1, /work/SAP1/data/
sales1.log, 50
# HADR_USER_DB_2=mydata
# HADR_USER_DB_MATERIALIZE_2=true
# HADR_COMPANION_DATA_DEVICE_CREATE_2_1=mydata_data_dev, /work/SAP1/data/
mydata.dat, 50
# HADR_COMPANION_LOG_DEVICE_CREATE_2_1=mydata_log_dev, /work/SAP1/data/
mydata.log, 25
#ASE HADR Secondary Site
#-----
#
# HADR_IS_SECONDARY_SITE_SETUP=<Is secondary site setup? true or false>
#
HADR_IS_SECONDARY_SITE_SETUP=false
#ASE HADR on Companion (Secondary) Site
#-----
# ASE and SRS servers on secondary (companion) site.
#
# If HADR_IS_SECONDARY_SITE_SETUP=true, below properties are required:
#
HADR_COMPANION_ASE_HOST_NAME=eiffel

```

```

HADR_COMPANION_SITE_NAME=Paris
HADR_COMPANION_ASE_SERVER_NAME=SYBASE
HADR_COMPANION_ASE_SERVER_PORT=5000
HADR_COMPANION_BS_SERVER_NAME=SYBASE_BS
HADR_COMPANION_BS_SERVER_PORT=5001
HADR_COMPANION_BS_DUMP_DIR=/opt/sap/data
HADR_COMPANION_COCKPIT_RMI_PORT=4992
HADR_COMPANION_COCKPIT_TDS_PORT=4998
HADR_COMPANION_RMA_RMI_PORT=7000
# RMA RMI occupies five consecutive ports, with the configured port occupying
the highest number.
HADR_COMPANION_RMA_TDS_PORT=7001
HADR_COMPANION_SRS_PORT=5005
HADR_COMPANION_SRS_DEVICE_BUFFER_DIR=/opt/sap/data
HADR_COMPANION_SRS_DEVICE_BUFFER_SIZE=256
# For synchronous replication, recommend path to an SSD
# (solid state drive) or other type of fast read/write storage device.
HADR_COMPANION_SRS_PERSISTENT_QUEUE_DIR=/ssd/data
HADR_COMPANION_SRS_PERSISTENT_QUEUE_SIZE=2000
#Cockpit Ports
#-----
#
# CONFIG_SCC_HTTP_PORT=<HTTP port number for Cockpit>
# CONFIG_SCC_HTTPS_PORT=<HTTPS port number for Cockpit>
# SCC_TDS_PORT_NUMBER=<TDS port number for Cockpit>
# SCC_RMI_PORT_NUMBER=<RMI port number for Cockpit>
CONFIG_SCC_HTTP_PORT=4282
CONFIG_SCC_HTTPS_PORT=4283
SCC_TDS_PORT_NUMBER=4998
SCC_RMI_PORT_NUMBER=4992
#Cockpit Logins
#-----
# These properties values must be same
# on both primary and companions sites.
#
# CONFIG_SCC_CSI_SCCADMIN_USER=<Cockpit administrator username>
# CONFIG_SCC_CSI_SCCADMIN_PWD=<Cockpit administrator password>
# CONFIG_SCC_CSI_UAFADMIN_USER=<Cockpit agent administrator username>
# CONFIG_SCC_CSI_UAFADMIN_PWD=<Cockpit agent administrator password>
#
# CONFIG_SCC_REPOSITORY_PWD=<Cockpit repository password>
#
CONFIG_SCC_CSI_SCCADMIN_USER=sccadmin
CONFIG_SCC_CSI_SCCADMIN_PWD=Secret396
CONFIG_SCC_CSI_UAFADMIN_USER=uafadmin
CONFIG_SCC_CSI_UAFADMIN_PWD=Secret396
CONFIG_SCC_REPOSITORY_PWD=Secret396

```

3.2.3.2 Installing the HADR System in Console Mode

Choose command line installation if you prefer a nonwindowing interface, or if you do not have graphical user interface setup on the machine.

Context

The steps for installing components in an interactive text mode are the same as when installing in GUI mode, except you use the following command to execute the installer from the command line, and you enter text to specify installation options:

(UNIX)

```
./setup.bin -i console
```

(Windows)

```
setupConsole.exe -i console
```

Procedure

1. To display Western European fonts on Windows:
 - a. Run `chcp 1252` to change the code page to 1252, which is the code page that Java uses to display Western European fonts for Windows.
 - b. Open the properties of `cmd.exe` and set Font to *Lucida Console*.
2. At the command line, start the installation program:

(UNIX)

```
./setup.bin -i console
```

(Windows)

```
setupConsole.exe -i console
```

The console output looks similar to:

```
./setup.bin -i console
Preparing to install...
Extracting the JRE from the installer archive...
Unpacking the JRE...
Extracting the installation resources from the installer archive...
Configuring the installer for this system's environment...
Launching installer...
=====
=
SAP Adaptive Server Enterprise                (created with
InstallAnywhere)
-----
-
Preparing CONSOLE Mode Installation...
=====
=
Introduction
-----
InstallAnywhere will guide you through the installation of SAP Adaptive Server
Enterprise 16.0 SP02 GA.
It is strongly recommended that you quit all programs before continuing with
this installation.
Respond to each prompt to proceed to the next step in the installation.  If
you
want to change something on a previous step, type 'back'.
You may cancel this installation at any time by typing 'quit'.
PRESS <ENTER> TO CONTINUE:
```

If `setupConsole.exe -i console` hangs on Windows 2008 or Windows 7, perform either of these workarounds:

- From **Start > All Programs > Accessories**, right-click *Command Prompt* and select *Run as administrator*, then enter `setupConsole.exe -i console`.
 - From the *Command Prompt* window, issue `setupConsole.exe -i console`, then press and hold the `Ctrl` key until you see the output.
3. Follow the remaining prompts to install the HADR system, then specify the basic settings for your installation. The flow of the installation is identical to a GUI installation, except output is written to a terminal window and responses are entered using the keyboard.

3.2.3.3 Installing the HADR System in Silent Mode

To perform a silent, or unattended installation, run the installer and provide a response file that contains your preferred installation configuration.

Prerequisites

Create a response file based on the instructions from [Using setup.bin or setupConsole.exe with a Response File \[page 67\]](#).

Context

The install site includes three response files:

- Primary server `install_directory/sample_response_primary.txt`
- Companion server `install_directory/sample_response_companion.txt`
- Fault Manager `install_directory/FaultManager/sample_response.txt`

Edit these response files for your site, and run them with the `setup.bin` installer to install and configure the HADR system. See [Sample Response Files \[page 68\]](#) for an example of an edited response file based on the choices in [Configuring SAP ASE \[page 42\]](#).

Procedure

1. Run the following, where `<responseFileName>` is the absolute path of the file name containing the installation options you chose:

```
setup.bin -f <response_file> -i silent -DAGREE_TO_SAP_LICENSE=true -
DRUN_SILENT=TRUE
```

Where `<response_file>` is the absolute path of the file name to the response file.

i Note

You must agree to the SAP License Agreement when installing in silent mode. You can either:

- Include the option `-DAGREE_TO_SAP_LICENSE=true` in the command line argument, or
- Edit the response file to include the property `AGREE_TO_SAP_LICENSE=true`.

Except for the absence of the GUI screens, all installation actions are the same, and the result of an installation in silent mode is exactly the same as one performed in GUI mode using the same responses.

2. The installer for SAP ASE requires non-null passwords for the sa login in SAP ASE, and uafadmin and sccadmin logins in SAP Control Center. For this reason, add these rows to the response file:
 - `SY_CFG_ASE_PASSWORD=<SAP ASE sa password>`
 - `CONFIG_SCC_CSI_SCCADMIN_PWD=<SCC admin password>`
 - `CONFIG_SCC_CSI_UAFADMIN_PWD=<SCC agent admin password>`
 - `CONFIG_SCC_REPOSITORY_PWD=<SCC repository password>`

Each password must be at least six characters long. The sccadmin and uafadmin logins need not be the same as the sa password. You can also set these passwords using these environment variables:

- `SY_CFG_ASE_PASSWORD`
- `CONFIG_SCC_CSI_SCCADMIN_PWD`
- `CONFIG_SCC_CSI_UAFADMIN_PWD`
- `CONFIG_SCC_REPOSITORY_PWD`

i Note

Secure these response files appropriately if they contain passwords.

3.2.3.4 Installing SAP ASE Cockpit from the Command Line

The software installer installs and configures the SAP ASE Cockpit by default. Performs these steps on each host if you installed the HADR system from the command line.

Context

These steps assume your HADR system (including the primary and standby servers, Replication Server, and RMA) is running.

Procedure

1. Verify that SAP ASE is running:

```
ps -ef | grep dataserver
sap      14426  7371  0 21:30 pts/0      00:00:00 grep dataserver
```



```
sap      19908 19907  1 Dec01 ?           02:47:34 /work/SAP1/ASE-16_0/bin/
dataserver -d/work/SAP1/data/master.dat -e/work/SAP1/ASE-16_0/install/
SFSAP1.log -c/work/SAP1/ASE-16_0/SFSAP1.cfg -M/work/SAP1/ASE-16_0 -N/work/
SAP1/ASE-16_0/sysam/SFSAP1.properties -i/work/SAP1 -sSFSAP1
```

2. Move to `$$SYBASE/COCKPIT-4/plugins/<server_name>` directory (in this example, `<$$SYBASE>` is `/work/SFSAP1`):

```
cd /work/SAP1/COCKPIT-4/plugins/SFSAP1
```

If this folder does not exist, copy the contents of the `$$SYBASE/COCKPIT-4/templates/com.sybase.ase` to `$$SYBASE/COCKPIT-4/plugins/`. For example:

```
cp -r $$SYBASE/COCKPIT-4/templates/com.sybase.ase $$SYBASE/COCKPIT-4/plugins/
SFSAP1
```

3. Verify the `<server_name>` directory contains these files and directories:
 - o `agent-plugin.dtd`
 - o `agent-plugin.xml`
 - o `lib`
 - o `mbean-descriptor.xml`
4. Verify the contents of the `agent-plugin.xml` file. If any values are incorrect, modify the contents to match your installation. The passwords are encrypted. These are the values for the sample HADR system (the Replication Server and RMA attributes are not required if you are not using an HADR system):

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<agent-plugin class="com.sybase.ua.plugins.ase.ASEAgentPlugin"
id="com.sybase.ase"
mbean-descriptor="mbean-descriptor.xml" mbean-type="Model MBean" name="ASE
Agent Plugin"
provider-name="SAP AG or an SAP affiliate company" register-on-startup="true"
version="16.0.0">
  <dependencies/>
  <properties>
    <!-- NOTE:
    Changes to this file must be made without the agent
    running. Upon agent shutdown all configuration
    information is written to this file and will overwrite
    any changes made while the agent was running.
    -->
    <set-property property="ase.server.name" value="SFSAP1"/>
    <set-property property="com.sybase.home" value="/work/SAP1"/>
    <set-property property="ase.home" value="/work/SAP1/ASE-16_0"/>
    <set-property property="ase.server.log" value="/work/SAP1/ASE-16_0/
install/SFSAP1.log"/>
    <set-property property="ase.interfaces.pathspec" value="/work/SAP1/
interfaces"/>
    <set-property property="ase.start.command" value="/work/SAP1/ASE-16_0/
install/RUN_SFSAP1"/>
    <set-property property="ase.port" value="5000"/>
    <set-property property="ase.user" value="sa"/>
    <set-property property="ase.password" value="1-AAAAEgQQpomOhkqiVVB/
tgOVg1wXo8JG/bXPo4LKnCSc2eVVWux4mzChdlXQ5Kyt62RuBshGVpEfLknBgDw8x27i
+eVY2f9k7qDhQW+mKP9oi4mFAos="/>
    <set-property property="ase.heartbeat.timer" value="60"/>
    <set-property property="ase.heartbeat.update.time" value=""/>
    <set-property property="ase.maintain.connection" value="true"/>
    <set-property property="rma.home" value="/work/SAP1/DM/RMA-16_0"/>
    <set-property property="rma.port" value="7001"/>
    <set-property property="rma.start.command" value="/work/SAP1/DM/
RMA-16_0/bin/RunContainer.sh"/>
```

```

    <set-property property="rma.log.dir" value="/work/SAP1/DM/RMA-16_0/
instances/AgentContainer/logs"/>
    <set-property property="rma.user" value="DR_admin"/>
    <set-property property="rma.password" value="1-
AAAAEgQQaUmiBafgUGKXF8iYfmhdj3Xtw8vDJ1fpWALHhmoCY7TIQbDuqP95dfxL9LkovWb0sdCZww
Xd9Q0ymVpRGwc5VsDPdt502Ua6z+H63Bw4Pcc="/>
    <set-property property="rs.sybase.home" value="/work/SAP1/DM"/>
    <set-property property="rs.home" value="/work/SAP1/DM/REP-16_0"/>
    <set-property property="rs.interfaces.pathspec" value="/work/SAP1/DM/
interfaces"/>
    <set-property property="rs.server.name" value="AS1_REP_SFHADR1"/>
    <set-property property="rs.port" value="5005"/>
    <set-property property="rs.start.command" value="/work/SAP1/DM/
AS1_REP_SFHADR1/RUN_AS1_REP_SFHADR1.sh"/>
    <set-property property="rs.server.log" value="/work/SAP1/DM/
AS1_REP_SFHADR1/AS1_REP_SFHADR1.log"/>
    <set-property property="rs.user" value="sa"/>
    <set-property property="rs.password" value="1-
AAAAEgQQ5DDP56g1OutUtXfmXX0BwU9ei70JoSHSxEz1gqIvDcLqf35ilHDx13fkMFwYac
+xtIv4b1OKSajeEumZ1hQgg="/>
  </properties>
</agent-plugin>

```

5. Re-encrypt the passwords for the SAP ASE, Replication Server, and RMA users. You perform these steps on each node because the generated encrypted passwords are different on each host:

- a. Move to `$SYBASE/COCKPIT-4/bin`.
- b. Execute the `passencrypt` password encryption utility:

```
./passcrypt
```

- c. Enter the password for the SAP ASE sa user:

```
[sybase@hadr-1 bin]$ ./passencrypt
Password:
1-
AAAAEgQQCaS88V4/9VrxH100Sli8vBQ72NbWCh66crULOrMX2GKEsKXn2CQeJokGitH5IhkXX6G
Zm5JNmlYJmzBVVppQjw==
```

- d. Replace the `agent-plugin.xml` password with the newly generated password.

6. Source the `$SYBASE/SYBASE.csh` (`SYBASE.sh` for the Korn shell) to configure the environment variables:

```
source SYBASE.csh
```

7. Move to the `$SYBASE/COCKPIT-4/bin` directory:

```
cd $SYBASE/COCKPIT-4/bin
```

8. Start the SAP ASE Cockpit:

```
./cockpit.sh --start
Starting Cockpit Server...
----- SYSTEM INFO -----
Home Directory:      /work/SAP1/COCKPIT-4
Version:             Cockpit Server 4 SP11
Node:                SFMACHINE1(10.173.1.109)
Log Message level:   WARN
Platform:            linux
Bitwidth:            64
OS Name:              Linux
OS Version:          2.6.32-504.8.1.el6.x86_64
OS Architecture:     amd64
Available Processors: 4
Total Physical Memory: 5974 MB
```

```

Free Physical Memory: 2529 MB
java.library.path: /work/SAP1/COCKPIT-4/rtlib:/work/SAP1/COCKPIT-4/bin/
sal6/linux64/lib64
Java Home: /work/SAP1/shared/SAPJRE-8_1_008_64BIT
Java Version: 1.8.0_45
Java VM Version: 8.1.008 25.45-b02
VM Process ID 18686@SFMAHCINE1
VM Arguments: -Xms128m, -XX:+HeapDumpOnOutOfMemoryError, -
XX:HeapDumpPath=/work/SAP1/COCKPIT-4/log,
-Dcom.sybase.home=/work/SAP1, -Dcom.sybase.ua.toplevel=/work/SAP1/COCKPIT-4, -
Dcom.sybase.ua.home=/work/SAP1/COCKPIT-4,
-Dcom.sybase.platform=linux, -Djava.library.path=/work/SAP1/COCKPIT-4/rtlib:/
work/SAP1/COCKPIT-4/bin/sal6/linux64/lib64,
-Djava.security.policy=/work/SAP1/COCKPIT-4/conf/java.policy, -
Djava.util.logging.manager=com.sybase.ua.util.logging.bridge.JavaToLog4jLogMan
ager,
-Dcom.sybase.security.BootstrapConfigurationURL=file:///work/SAP1/COCKPIT-4/
conf/csibootstrap.properties
Cockpit Web Console:
https://SFMACHINE1:4283/cockpit
Type help for more options
cockpit>

```

9. The SAP ASE Cockpit is running when the prompt appears. Issue `shutdown` if you see any errors, fix the appropriate lines in `agent-config.xml`, then restart the SAP ASE Cockpit.

`cockpit.sh` displays the connection URL for the client just before the `cockpit` prompt (in bold text above).

10. Launch a browser and open the URL specified by the `cockpit.sh` output (`https://SFMACHINE1:4283/cockpit` in the output above). When the page opens, you may get a warning about this being an untrusted connection. Click "I understand the risks" and "Add exception."
11. If necessary, install Adobe Flash.
12. Enter the user name and password to log into SAP ASE Cockpit.
13. When you first connect, SAP ASE Cockpit may issue a statement about inadequate values for configuration parameters. Enter the appropriate values and click *OK*.

i Note

You need not restart SAP ASE.

SAP ASE Cockpit displays the Monitoring tab. You can safely ignore the errors at the bottom of the screen.

14. If you did not configure the cockpit technical user in the SAP ASE Cockpit, configure it now.
 - a. Select the *Explore* tab.
 - b. Select ► *ASE Servers* ► *<server_name>* ► *Create Cockpit Technical User* ►

i Note

If the cockpit technical user is already created (for example, during an install using the SAP installer), the menu item reads *Update Cockpit Technical User*, which includes a wizard for updating a technical user in SAP ASE for cockpit service.

- c. Follow the prompts for the *Create Technical User* wizard.
15. Register the Agent.
 - a. Select the *Explore* tab.
 - b. Select ► *ASE Servers* ► *<server_name>* ► *Properties* ►.

- c. Select the *Agent* tab.
- d. Enter the agent port number (default is 4992) and click *Register*.
- e. Enter `uafadmin` as the user name and enter a password.
- f. Select *Authenticate*.

16. If you have forgotten and now need to reset the password:

- a. Stop cockpit:

```
$SYBASE/COCKPIT-4/bin/cockpit --stop
```

- b. Move to `$SYBASE/COCKPIT-4/bin`:

```
cd $SYBASE/COCKPIT-4/bin
```

- c. Execute the `passencrypt` script:

```
passencrypt -csi
```

- d. Enter the new password:

```
passencrypt -csi
Password:
```

- e. Copy the output for the new password (see the bold text below):

```
passencrypt -csi
Password:
{SHA-256:aszZYC3i5Ms=}vvytnz5U95b7UyTMxrRxq7TizJY8R088Ri8IimnAFXU=
```

- f. Move to `$SYBASE/COCKPIT-4/conf`:

```
cd $SYBASE/COCKPIT-4/conf
```

- g. Edit the `csi_config.xml` file. Search for the `uafadmin` section of the file, and paste the value in the `uafadmin password` property between the two double quotes (see the bold text):

```
<!-- Default SCC agent admin account: uafadmin -->
<authenticationProvider controlFlag="sufficient"
  name="com.sybase.security.core.PreConfiguredUserLoginModule">
  <options name="username" value="uafadmin"/>
  <!-- options name="password"
value="{SHA-256:WNATpqw76zA=}GYeAKdTRiIh1VcqmWv1k/
A2pcXSHfLUBr9boP03ArKE=" / -->
  <options name="roles" value="SCC Agent Administrator"/>
</authenticationProvider>
```

- h. Restart SAP ASE Cockpit.

17. Register the Replication Management Agent (RMA). Perform the following on both servers:

- a. Select the *Explore* tab.
- b. Select **► ASE Servers ► <server_name> ► Properties ►**.
- c. Select the *HADR* tab. The installer displays the Register Agent screen.
- d. Enter user name and password for the RMA administrator user (typically `DR_admin`).
- e. Select *Authenticate*.

3.2.4 Adding Databases to the HADR System

There are a number of ways to add databases to the HADR system.

You can:

- Create them during installation. See the step describing the ASE HADR Setup Site screen in [Configuring SAP ASE \[page 42\]](#).
- Add them after installation. See [Adding Databases from the Command Line After Installation \[page 299\]](#).
- Use the RMA commands. See [sap_update_replication \[page 551\]](#).
- Use SAP ASE Cockpit. See *Manage SAP ASE > Always-On (HADR) Option > Adding an Existing Database to the HADR System* in the SAP ASE Cockpit documentation

Regardless of the method you use, the databases you create are initially empty. Enter data in the database by loading a dump taken from previous installation or create new data using your site's resources.

3.2.5 Migrating from an Existing Replication Server

Remove all replication configurations before upgrading to the latest version of SAP ASE, and rematerialize all databases as part of the HADR configuration.

Context

Remove the existing Replication Server connections, subscriptions, and so on from the environment before configuring it for HADR.

Procedure

1. Remove subscriptions to database, table, and function replication definitions, articles, or publications. This example drops the `authors_sub` subscription for the `authors_rep` table replication definition:

```
drop subscription authors_sub
for authors_rep
with replicate at SJSAP2.pubs2
without purge
```

2. Drop replication definitions, articles, or publications:

```
drop replication definition authors_rep
```

3. Remove databases from the replication system. This example drops the connection to the `pubs2` database in the SJSAP2 companion:

```
drop connection to SJSAP2.pubs2
```

4. Disable the secondary truncation point. When you migrate a primary database, the Replication Agent cannot be running. Turn off the secondary truncation point for the duration of the migration. Shut down the Replication Agent prior to the migration:

a. Stop the Replication Agent:

```
use <primary_dbname>
go
sp_stop_rep_agent <primary_dbname>
go
```

b. Disable the secondary truncation point in the SAP ASE database that is being migrated:

```
use <database_name>
go
dbcc settrunc('ltm', 'ignore')
go
```

3.2.5.1 Upgrading to SAP ASE Version 16.0 SP02

See the installation guide for your platform for instructions on upgrading SAP ASE to version 16.0 SP02.

3.3 Installing HADR with an Existing System

You can configure HADR from your existing system.

When you install the HADR system using an existing server:

- Use the existing SAP ASE as the primary server
- The primary and companion servers uses the same version number
- The companion and primary server uses the same values for the:
 - "sa" password
 - Default language, character set, sort order
 - Page sizes
 - Database sizes
- If you are migrating an SAP ASE configured as a high-availability (HA) cluster or for disaster recovery (DR) with Replication Server, tear down these systems before migrating to HADR. If you are migrating an SAP ASE Cluster Edition, convert this to a non-clustered instance before upgrading to the latest SAP ASE version with HADR.

i Note

After you create a new SAP ASE instance, copy the configuration file from the primary server to capture the configuration values, cache configurations, and to re-create any additional user temporary databases.

The user bindings for `tempdb` are carried over when the HADR system synchronizes `syslogins`.

3.3.1 Migrating an SMP Server to an HADR System

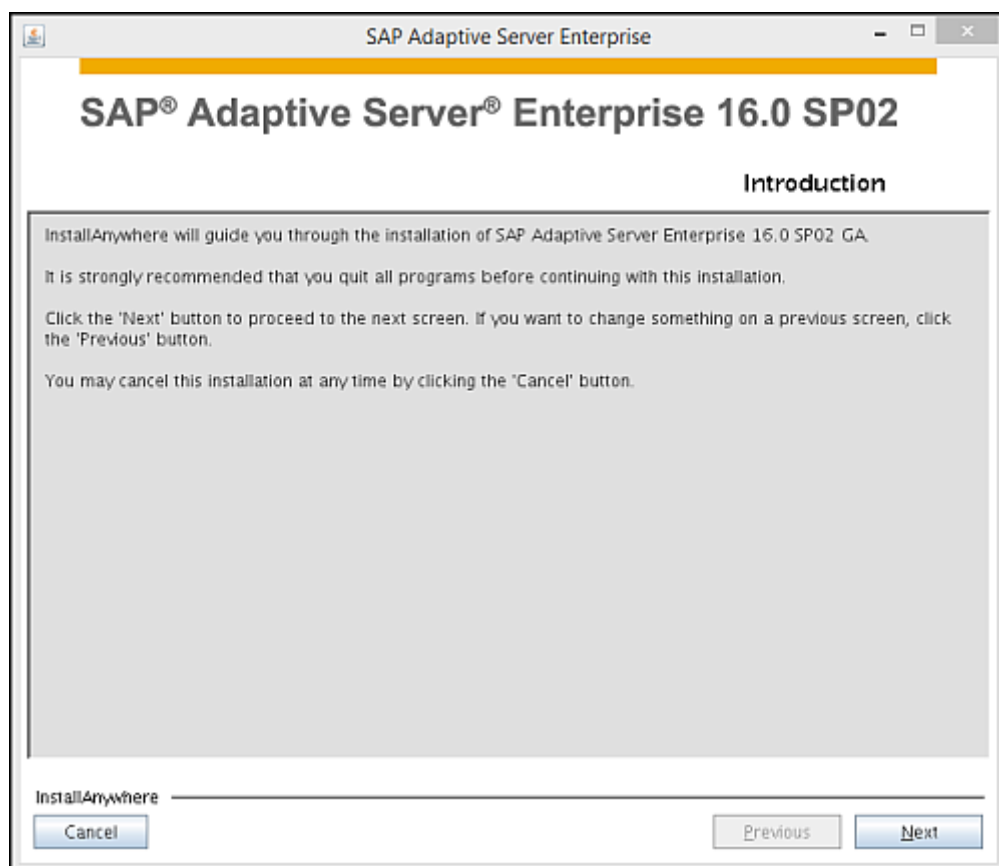
Use the `setuphadr` utility to migrate an existing SAP ASE server to the HADR system.

Procedure

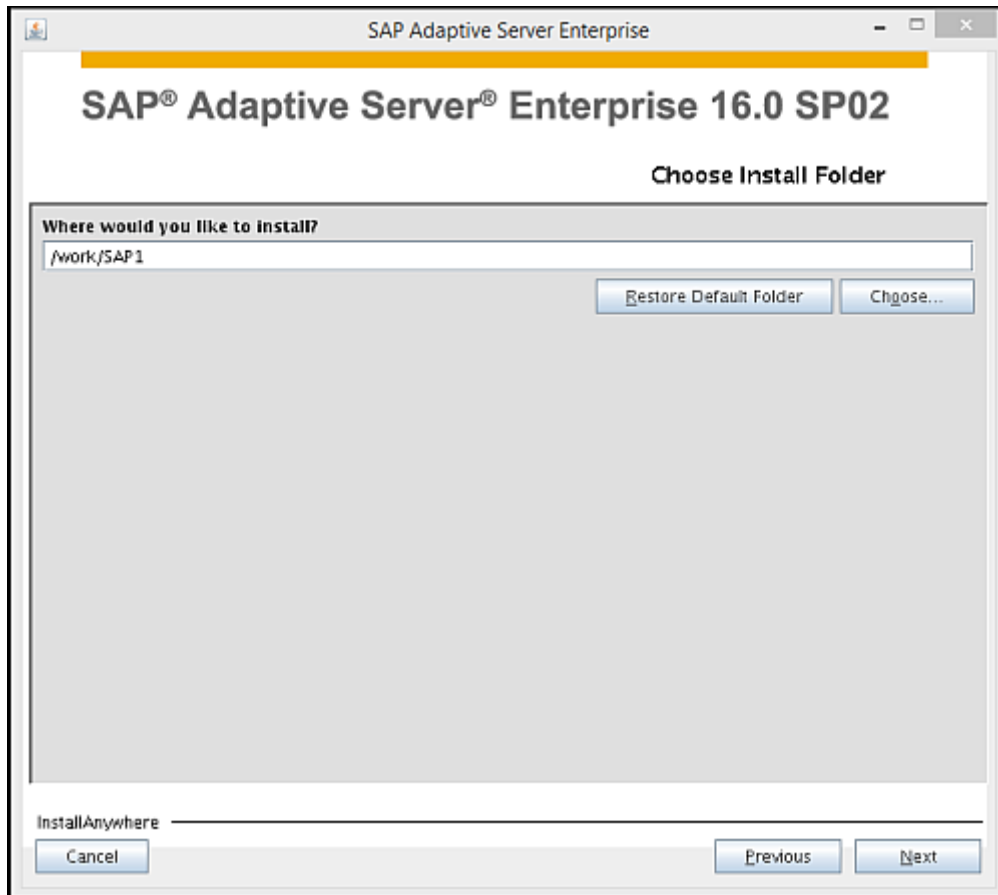
On the Primary Site

1. Upgrade SAP ASE to the current version. See the installation guide for information. The upgrade process will also upgrade other components such as Backup Server, XP Server, and Job Scheduler. This procedure contains no HADR-related steps. This server will become the primary server. The upgrade process uses the `updatease` (or `updatease.exe` on Windows) utility to upgrade from SAP ASE 16.0 and later, and uses `sqlupgrade` or `sqlupgraderes` (`syconfig.exe` or `sybatch.exe` on Windows) to upgrade from versions earlier than 16.0.
2. Install the Data Movement feature to `$$SYBASE`:
 - a. Launch the SAP installer from the installation directory:


```
./setup.bin
```




- b. In the Choose Install Folder screen, enter the current installation directory for the primary server (that is, `$$SYBASE`):



- c. In the Choose Install Set screen, select the *Custom* install set, then click *Next*.



- d. From the list of components you can install, select *SAP ASE Data Movement for HADR*. Also select Cockpit, if necessary.



- e. The Fault Manager requires the SAP Host Agent to be installed on the hosts running the database (SAP ASE). You need root authorization or sudo permission to install the SAP Host agent. If you choose to install the SAP Host Agent now rather than later, enter the password for sudo.



See [The SAP Host Agent \[page 28\]](#) for more information.


- f. In the Software License Type Selection screen, select the appropriate license type.

i Note

HADR is not supported on the Developer and Express Edition licenses.



- g. Select your geographical location from the drop-down list. Review the license agreement, and click "I agree..." then click *Next*.



h. Specify your license edition and type.

SAP® Adaptive Server® Enterprise 16.0 SP02

Product Licenses

Please select the product edition and license type you would like to configure.

Product Edition:

License Type:

InstallAnywhere _____

- i. If you select **Yes** in the SySAM Notification screen, configure any email alerts.

SAP® Adaptive Server® Enterprise 16.0 SP02

SySAM Notification

Please configure the SySAM email alert mechanism. When configured, specified recipients will receive email notifications about SySAM events that may need administrator attention.

Do you want to configure email alerts?

Yes

No

SMTP server host name:

SMTP server port number:


Sender email:

Recipient emails:

Message Severity for email alerts:

InstallAnywhere _____

- j. Review the preinstallation summary. Click [Previous](#) to make changes. When you are ready, click [Install](#) for the installer to unload the files to the disk.



- k. In the Configure New Servers screen, deselect all options as the servers are already configured. Click [Next](#).



When finished, the installer displays a success screen. Select *Done* to exit the installer.



3. (UNIX only) set the environment variables. Source the `SYBASE.csh` or `SYBASE.sh` files:

```
source $SYBASE/SYBASE.csh
```

4. Restart SAP ASE and Backup Server. Move to the `$SYBASE/$SYBASE_ASE/install` directory and issue:

```
./RUN_<server_name>
./RUN_<backup_server_name>
```

5. Start RMA:
 - UNIX – `$SYBASE/$SYBASE_ASE/bin/rma`
 - Windows – (recommended) start the *Sybase DR Agent* Windows service. Alternatively, you can issue `%SYBASE%\%SYBASE_ASE%\bin\rma.bat` at the command line.
6. Identify and record these values on the primary server (both the primary and companion servers require the same values):
 - Device sizes – use `sp_helpdevice`
 - Database sizes – use `sp_helpdb`
 - Page sizes – Use this to determine the logical page size:

```
select @@maxpagesize
```

- Default language, character set, sort order – use `sp_helplanguage` and `sp_helpsort`

See [Collecting Migration Configuration Details \[page 102\]](#) for information and examples.

7. Save the `$SYBASE/ASE-16_0/<server_name>.cfg` file.

8. The HADR system requires the cluster ID database. You can create this database now or have the `setuphadr` utility create it later. To create it now:
 - a. Log in to SAP ASE.
 - a. If necessary, use the `disk init` command to create the device for the cluster ID database.
 - b. Use the `create database` command to create the cluster ID database, making sure that it is at least 200 MB in size.

On the Companion Site

9. Use the SAP installer to unload the binaries for SAP ASE. If this includes the Data Movement feature, you must also install the Host Agent, if you installed that on the primary companion. Do **not**, however, configure the server at this time. See [Unloading the SAP ASE Binaries \[page 30\]](#).

After you unload the binaries and the installer displays the Configure New Servers screen, deselect all the options and click *Next*.



10. Set the environment variables. Source the `SYBASE.csh` or `SYBASE.sh` files:

```
source $SYBASE/SYBASE.csh
```

11. Use the `srvbuild` or `srvbuildres` utility (`syconfig.exe` or `sybatch.exe` on Windows) to create the companion server and Backup Server (and, if necessary, XP Server and Job Scheduler). `srvbuild` or `srvbuildres` also allow you to create the technical user and enable SAP ASE Cockpit monitoring. See the *Configuration Guide > Configuring New Servers with srvbuild*. Make sure that the primary and companion servers use the same:
 - Application type
 - Logical page size

- sa password

You need not optimize your server.

12. Use the `sqlloc` or `sqllocres` utility (or `syconfig.exe` or `sybatch.exe` on Windows) to configure the default language, character set, and sort order on the companion server to match those on the primary server. See the *Configuration Guide > Localization Support > Changing SAP ASE Localization Configuration* and the *System Administration Guide: Volume 1 > Configuring Character Sets, Sort Orders, and Languages > Changing the Character Set, Sort Order, or Message Language*.
13. Use `disk init` and `create database` to create empty user databases with sizes that match those on the primary SAP ASE.
14. The HADR system requires the cluster ID database. You can create this database now or have the `setuphadr` utility create it later. To create it now:
 - a. Log in to SAP ASE.
 - b. If necessary, use the `disk init` command to create the device for the cluster ID database.
 - c. Use the `create database` command to create the cluster ID database, making sure that it is at least 200 MB in size.
15. Shut down SAP ASE by logging into SAP ASE using `isql`, and issuing:

```
shutdown
```

16. Copy the `<server_name>.cfg` from the primary server to the companion server (the default location of the `<server_name>.cfg` file is in `<$SYBASE/$SYBASE_ASE>`).
17. Start SAP ASE using the newly copied `<server_name>.cfg` by moving to the `$SYBASE/ASE-16_0/install` directory and issuing:

```
./RUN_<server_name>
```

18. Start RMA:
 - UNIX – `-$SYBASE/$SYBASE_ASE/bin/rma`
 - Windows – (recommended) start the *Sybase DR Agent* Windows service. Alternatively, issue `%SYBASE%\%SYBASE_ASE%\bin\rma.bat` at the command line.

Run setuphadr on the Primary and Companion Sites to Configure HADR

19. Create a `setuphadr` response file. A sample `setup_hadr.rs` is available in the `$SYBASE/$SYBASE_ASE/init/sample_resource_files` directory (`%SYBASE%\%SYBASE_ASE%\sample\server` directory on Windows). Make a copy of this file and make changes according to the requirements of your site.

i Note

If you have not already done so, create the cluster ID database. Enter a value of 200 MB or larger if you specified the `ase_data_device_create_*` and `ase_log_device_create_*` properties for the cluster ID database.

20. On primary site:
 - a. Set these properties in the `setuphadr setup_hadr.rs` response file:

```
setup_site=<primary_site>  
is_secondary_site_setup=false
```

See [Sample setup_hadr.rs Response File \[page 104\]](#) for examples of the changes required.

- b. Modify the `setup_hadr.rs` file with data such as device sizes, database size, and so on, which you collected in the section titled "On the Primary Site" in *Migrating an SMP Server to an HADR System*.
- c. Run `setuphadr` with the response file:
 - (UNIX) – `SYBASE/$SYBASE_ASE/bin/setuphadr <path_to_response_file>`
 - (Windows) – `%SYBASE%\%SYBASE_ASE%\bin\setuphadr.bat <path_to_response_file>`

The output looks similar to:

```
./ASE-16_0/bin/setuphadr setup_SFHDR.rs
Setup ASE server configurations
    Set server configuration "max network packet size" to "16384"...
    Reboot SAP ASE "SFSAP1"...
Setup ASE server configurations...Success
Setup user databases
    Create user database AS1...
    Set "pubs2" database "trunc log on chkpt" option to "false"...
Setup user databases...Success
Setup ASE HADR maintenance user
    Create maintenance login "DR_maint"...
    Grant "sa_role" role to "DR_maint"...
    Grant "replication_role" role to "DR_maint"...
    Grant "replication_maint_role_gp" role to "DR_maint"...
    Create "sap_maint_user_role" role...
    Grant set session authorization to "sap_maint_user_role"...
    Grant "sap_maint_user_role" role to "DR_maint"...
    Add auto activated roles "sap_maint_user_role" to user
"DR_maint"...
    Allow "DR_maint" to be known as dbo in "master" database...
    Allow "DR_maint" to be known as dbo in "AS1" database...
    Allow "DR_maint" to be known as dbo in "pubs2" database...
Setup ASE HADR maintenance user...Success
Setup administrator user
    Create administrator login "DR_admin"...
    Grant "sa_role" role to "DR_admin"...
    Grant "ssr_role" role to "DR_admin"...
    Grant "replication_role" role to "DR_admin"...
    Grant "hadr_admin_role_gp" role to "DR_admin"...
    Grant "sybase_ts_role" role to "DR_admin"...
Setup administrator user...Success
Setup Backup server allow hosts
    Backup server on "site1" site: Add host "mo-
bf1dc68822.mo.sap.corp" to allow dump and load...
Setup Backup server allow hosts...Success
Setup complete on "site1" site. Please run Setup HADR on "site2" site to
complete the setup.
```

21. On the companion site:

- a. Set these properties in the `setup_hadr.rs` response file:

```
setup_site=<companion_site>
is_secondary_site_setup=true
```

- b. Run `setuphadr` with the response file:
 - (UNIX) – `SYBASE/$SYBASE_ASE/bin/setuphadr <path_to_response_file>`
 - (Windows) – `%SYBASE%\%SYBASE_ASE%\bin\setuphadr.bat <path_to_response_file>`

The output looks similar to:

```
./ASE-16_0/bin/setuphadr setup_SJHDR.rs
Setup user databases
    Set "pubs2" database "trunc log on chkpt" option to "false"...
Setup user databases...Success
```

```

Setup Backup server allow hosts
    Backup server on "site2" site: Add host "SFMACHINE1.BIG.corp" to
allow dump and load...
    Backup server on "site1" site: Add host "SJMACHINE2.BIG.corp" to
allow dump and load...
Setup Backup server allow hosts...Success
Setup RMA
    Set SAP ID to "AS1"...
    Set installation mode to "nonBS"...
    Set site name "SFHADR1" with SAP ASE host:port to
"SFMACHINE1.BIG.corp:5001" and Replication Server host:port to
"SFMACHINE1.BIG.corp:5005"...
    Set site name "SJHADR2" with SAP ASE host:port to
"SJMACHINE2.BIG.corp:5001" and Replication Server host:port to
"SJMACHINE2.BIG.corp:5005"...
    Set site name "SFHADR1" with Backup server port to "5002"...
    Set site name "SJHADR2" with Backup server port to "5002"...
    Set site name "SFHADR1" databases dump directory to "/work/SAP1/
data"...
    Set site name "SJHADR2" databases dump directory to "/work/SAP2/
data"...
    Set site name "SFHADR1" synchronization mode to "sync"...
    Set site name "SJHADR2" synchronization mode to "sync"...
    Set site name "SFHADR1" distribution mode to "remote"...
    Set site name "SJHADR2" distribution mode to "remote"...
    Set site name "SFHADR1" distribution target to site name
"SJHADR2"...
    Set site name "SJHADR2" distribution target to site name
"SFHADR1"...
    Set maintenance user to "DR_maint"...
    Set site name "SFHADR1" device buffer directory to "/work/SAP1/
data"...
    Set site name "SJHADR2" device buffer directory to "/work/SAP2/
data"...
    Set site name "SFHADR1" device buffer size to "512"...
    Set site name "SJHADR2" device buffer size to "512"...
    Set site name "SFHADR1" simple persistent queue directory to "/"
work/SAP1/data"...
    Set site name "SJHADR2" simple persistent queue directory to "/"
work/SAP2/data"...
    Set site name "SFHADR1" simple persistent queue size to "2000"...
    Set site name "SJHADR2" simple persistent queue size to "2000"...
    Set master, pubs2, AS1 databases to participate in replication...
Setup RMA...Success
Setup Replication
    Setup replication from "SFHADR1" to "SJHADR2"...
    Configuring remote replication server.....
    Configuring local replication
server.....
    Setting up replication on 'standby' host for local database
'master'.....
    Setting up replication on 'standby' host for local database
'pubs2'.....
    Setting up replication on 'standby' host for local database
'AS1'.....
Setup Replication...Success
Materialize Databases
    Materialize database "master"...
        Starting materialization of the master database from
source 'SFHADR1' to target 'SJHADR2'...
        Completed materialization of the master database from
source 'SFHADR1' to target 'SJHADR2'..
        Waiting 10 seconds: Before checking if Replication
Connection 'S1_SJHADR2.master' is suspended.....
    Materialize database "pubs2"...
        Executing ASE dump and load task for database
'pubs2'.....

```

```

        Waiting 10 seconds: Before checking if Replication
Connection 'S1_SJHADR2.pubs2' is suspended.....
        Materialize database "AS1"...
        Executing ASE dump and load task for database 'AS1'.....
        Successfully verified materialization on database 'AS1'..
        Stop the Replication Agent for database 'master' on host
'SFMACHINE1.BIG.corp:5001' and data server 'AS1_SFHDR1'..
        Start the Replication Agent for database 'master' on host
'SFMACHINE1.BIG.corp:5001' and data server 'AS1_SFHDR1'..
        Stop the Replication Agent for database 'AS1' on host
'SFMACHINE1.BIG.corp:5001' and data server 'AS1_SFHDR1'..
        Configuring Replication Server: set 'hide_maintuser_pwd'
to 'o'...
        Waiting 10 seconds: Before checking if Replication
Connection 'S1_SJHADR2.AS1' is suspended.....
        Completed automatic materialization of database 'AS1' from
source 'SFHDR1' to target 'SJHADR2'...
Materialize Databases...Success

```

22. If you installed SAP ASE Cockpit, you perform the registration steps on both sites. See [Configuring the SAP ASE Cockpit for the HADR System \[page 109\]](#).

3.3.2 Collecting Migration Configuration Details

Use the information you gather in this section to complete the steps in the section titled "On the Primary Site."

If you already have scripts for creating the devices and databases for the primary companion, use those scripts to configure the companion server. Otherwise you can use the `sp_helpdb`, `sp_helpdevice`, `sp_helpsort`, and `ddlgen` system procedures to make sure the companion server perfectly mimics the device and database make up of the primary companion.

For example, this devices associated with the a primary server:

```

sp_helpdevice
go
 device_name physical_name
description
                status cntrltype vdevno vpn_low vpn_high
-----
-----
master          /work/SAP1/data/master.dat          file system device, special,
dsync on, directio off, default disk, physical disk, 52.00 MB, Free: 8.00
MB              3              0              0              0              26623
salesdev1       /work/SAP1/data/salesdev1.dat        file system device, special,
dsync off, directio on, physical disk, 15.00 MB, Free: 5.00
MB              2              0              5              0              7679
salesdev2       /work/SAP1/data/salesdev2.dat        file system device, special,
dsync off, directio on, physical disk, 20.00 MB, Free: 0.00
MB              2              0              7              0              10239
saleslog1       /work/SAP1/data/saleslog1.dat        file system device, special,
dsync off, directio on, physical disk, 10.00 MB, Free: 2.00
MB              2              0              6              0              5119
sybmgmtdev     /work/SAP1/data/sybmngmtb.dat        file system device, special,
dsync off, directio on, physical disk, 76.00 MB, Free: 0.00
MB              2              0              4              0              38911
sysprocsdev    /work/SAP1/data/sysprocs.dat         file system device, special,
dsync off, directio on, physical disk, 196.00 MB, Free: 0.00
MB              2              0              1              0              100351
systemdbdev    /work/SAP1/data/sybsysdb.dat         file system device, special,
dsync off, directio on, physical disk, 6.00 MB, Free: 0.00
MB              2              0              2              0              3071

```

```

tapedump1 /dev/nst0 unknown device type, disk, dump
device
16 2 0 0 20000
tapedump2 /dev/nst1 unknown device type,
tape, 625 MB, dump
device 16 3
0 0 20000
tempdbdev /work/SAP1/data/tempdbdev.dat file system device, special,
dsync off, directio on, physical disk, 100.00 MB, Free: 0.00
MB 2 0 3 0 51199

```

The following shows that server SFSAP1 has the sales database installed:

```

sp_helpdb
go
name db_size owner dbid created durability lobcomplvl
inrowlen
status
-----
-----
-----
master 26.0 MB sa 1 Jan 18, 2016 full 0
NULL mixed log and
data
model 6.0 MB sa 3 Jan 18, 2016 full 0
NULL mixed log and
data
sales 38.0 MB sa 4 Jan 19, 2016 full 0
NULL no options
set
sybmgmtdb 76.0 MB sa 31515 Jan 18, 2016 full 0
NULL select into/bulkcopy/pllsort, trunc log on chkpt, mixed log and
data
sybssystemdb 12.0 MB sa 31513 Jan 18, 2016 full 0
NULL mixed log and
data
sybssystemprocs 196.0 MB sa 31514 Jan 18, 2016 full 0
NULL trunc log on chkpt, mixed log and
data
tempdb 106.0 MB sa 2 Jan 18, 2016 no_recovery 0
NULL select into/bulkcopy/pllsort, trunc log on chkpt, mixed log and data, allow
wide dol rows

```

This shows information specifically about a database names sales:

```

sp_helpdb sales
go
name db_size owner dbid created durability lobcomplvl inrowlen
status
-----
-----
-----
sales 38.0 MB sa 4 Jan 19, 2016 full 0 NULL no
options set
(1 row affected)

device_fragments size usage created free_kbytes
-----
-----
-----
salesdev1 10.0 MB data only Jan 19 2016 12:22PM 6824
saleslog1 8.0 MB log only Jan 19 2016 12:22PM not applicable
salesdev2 20.0 MB data only Jan 19 2016 12:25PM 20400
-----
-----

```

```
log only free kbytes =
8128
```

If your site includes scripts to install databases and devices, you can use these to make sure the primary and companion servers use the same configuration. For example, you can use the `installpubs2` script to install the `pubs2` database, described above, on the companion server.

If your site does not use scripts to configure the databases and devices, you can use the `ddlgen` utility to populate the `setup_hadr.res` file, which you will use to configure the companion server. For example, this displays the object definitions for the `pubs2` database:

```

$SYBASE/$SYBASE_ASE/bin/ddlgen -Usa -P -SSFMACHINE:5000 -TDB -Npubs2 -F% -Shdwh
USE master
go
PRINT "<<<< CREATE DATABASE pubs2>>>>"
go
CREATE DATABASE pubs2
           ON master = '4M' -- 2048 pages
WITH DURABILITY = FULL
go
use pubs2
go
exec sp_changedbowner 'sa', true
go
exec master.dbo.sp_dboption pubs2, 'trunc log on chkpt', true
go
checkpoint
go

```

3.3.3 Sample `setup_hadr.res` Response File

The codeblock below illustrates a sample `setup_hadr.res` response file based on the primary server, as described in the installation chapters of this user guide (changed responses are in **bold**).

```

#####
# Setup HADR sample responses file
#
# This sample response file sets-up SAP ASE HADR on
# hosts "host1" (primary) and "host2" (companion).
#
# Prerequisite :
# - New SAP ASE and Backup servers are already setup and started on "host1" and
# "host2".
#   See HADR User Guide for requirements on SAP ASE servers.
# - Replication Management Agent (RMA) is already started on "host1" and "host2".
#
# Usage :
# 1. On host1 (primary), run:
#   $SYBASE/$SYBASE_ASE/bin/setuphadr <this_responses_file>
#
# 2. Change this responses file properties:
#   setup_site=site2
#   is_secondary_site_setup=true
#
# 3. On host2 (companion), run:
#   $SYBASE/$SYBASE_ASE/bin/setuphadr <responses_file_from_step_2>
#
#####
# ID that identifies this cluster

```



```

#
# Value must be unique,
# begin with a letter and
# 3 characters in length.
cluster_id=AS1
# Which site being configured
#
# Note:
# You need to set "<setup_site_value>.*"
# properties in this responses file.
setup_site=site1
# Has the secondary site prepared for ASE HADR
#
# Valid values: true, false
#
# If set to true, "<secondary_setup_site_value>.*"
# properties must set in this responses file.
is_secondary_site_setup=false
# How data is replicated
#
# Valid values: sync, async
synchronization_mode=sync
# SAP ASE system administrator user/password
#
# setuphadr will prompt from standard input if not specified
ase_sa_user=sa
ase_sa_password=<sa_password>
# ASE HADR maintenance user/password
#
# Password must have at least 6 characters
# setuphadr will prompt from standard input if not specified
hadr_maintenance_user=DR_maint
hadr_maintenance_password=<maint_password>
# Replication Management Agent administrator user/password
#
# Password must have at least 6 characters
# setuphadr will prompt from standard input if not specified
rma_admin_user=DR_admin
rma_admin_password=<rma_admin_password>

#If need to disable the checks for reference constraints
#
#Valid values: true, false
disable_referential_constraints=false
# Databases that will participate in replication
# and "auto" materialize.
#
# If database doesn't exist in the SAP ASE, you need
# to specify <site>.ase_data_device_create_[x]_[y] and
# <site>.ase_log_device_create_[x]_[y] properties.
# See below.
#
# ASE HADR requires SAP ASE to have a database
# with cluster ID name (see "cluster_id" above).
# If you have not created this database, you can
# enter it here to have it created.
# cluster ID database
participating_database_1=AS1
materialize_participating_database_1=true
# user database
participating_database_2=pubs2
materialize_participating_database_2=true
# user database
# participating_database_3=userdb2
# materialize_participating_database_3=true
# Enable SSL - true or false
enable_ssl=true

```

```

# SSL common name - you are recommended to use the name of the SAP ASE server
ssl_common_name=YOUR_HADR_SERVERNAME

# Name of the SSL private key file
ssl_private_key_file=/tmp/hadr.key

# Name of the SSL public key file
ssl_public_key_file=/tmp/hadr.cert

# Name and location of the Root CA certificate. If you are using a self-signed
certificate, put your public key file here
ssl_ca_cert_file=/tmp/rootCA.pem

# SSL password to protect your private key
ssl_password=password
#####
# Site "site1" on host host1 with primary role
#####
# Host name where SAP ASE run
#
# Enter fully qualified domain name (FQDN)
# if your sites are on different subnet.
site1.ase_host_name=SFMACHINE.BIG.corp
# Site name
#
# Enter value that identifies this site,
# like a geographical location.
# Value must be unique.
site1.site_name=SFHADR1
# Site role
#
# Enter the role of this site.
# Valid values: primary, companion
site1.site_role=primary
# directory where SAP ASE installed
site1.ase_release_directory=/work/SAP1
# Directory that stored SAP ASE user data files
# (interfaces, RUN_<server>, error log, etc. files).
# Do not set value if your user data files are in
# SAP ASE installed directory (ase_release_directory).
site1.ase_user_data_directory=/work/SAP1
site1.ase_server_name=SFSAP1
site1.ase_server_port=5000
site1.backup_server_name=SFSAP1_BS
site1.backup_server_port=5001
# Directory to store database dumps
# in materialization
#
# Backup server must able to access this directory
site1.backup_server_dump_directory=/work/SAP1/data
# Data & log devices to create the databases specified
# in "participating_database_[x]" properties. You do
# not need to specify these properties if the database(s)
# already exist in the SAP ASE server.
#
# ase_data_device_create_[x]_[y] - property to create data device
# ase_log_device_create_[x]_[y] - property to create log device
#
#     where
#     x is number in "participating_database_[x]" property
#     y is number device to create
#
# Format: <logical_device_name>, <physical_device_path>, <size_in_MB>
#
# NOTE: Databases sizes on primary and companion
#       SAP ASE must be the same.
# Device for cluster ID database "LE1" (See "participating_database_1" property)
# Database size = 25MB
#       data device "le_data_dev" = 25MB

```

```

site1.ase_data_device_create_1_1=AS1_data_dev, /work/SAP1/data/AS1_dev1.dat, 25
# Devices for database "userdb1" (See "participating_database_2" property)
# Database Size = 100MB
#   data device 1 "db1_data_dev1" = 25MB
#   data device 2 "db1_data_dev2" = 25MB
#   data device 3 "db1_data_dev3" = 25MB
#   log device 1 "db1_log_dev1" = 25MB
# site1.ase_data_device_create_2_1=db1_data_dev1, /host1_eng/ase/data/
db1_dev1.dat, 25
# site1.ase_data_device_create_2_2=db1_data_dev2, /host1_eng/ase/data/
db1_dev2.dat, 25
# site1.ase_data_device_create_2_3=db1_data_dev3, /host1_eng/ase/data/
db1_dev3.dat, 25
# site1.ase_log_device_create_2_1=db1_log_dev1, /host1_eng/ase/data/
db1_dev1.log, 25
# Devices for database "userdb2" (See "participating_database_3" property)
# Database Size = 100MB
#   data device 1 "db2_data_dev1" = 25MB
#   data device 2 "db2_data_dev2" = 25MB
#   log device 1 "db2_log_dev1" = 25MB
#   log device 2 "db2_log_dev2" = 25MB
# site1.ase_data_device_create_3_1=db2_data_dev1, /host1_eng/ase/data/
db2_dev1.dat, 25
# site1.ase_data_device_create_3_2=db2_data_dev2, /host1_eng/ase/data/
db2_dev2.dat, 25
# site1.ase_log_device_create_3_1=db2_log_dev1, /host1_eng/ase/data/
db2_dev1.log, 25
# site1.ase_log_device_create_3_2=db2_log_dev2, /host1_eng/ase/data/
db2_dev2.log, 25
# Port numbers for Replication Server and Replication Management Agent on host1
#
# In remote topology, these are the companion Replication Server and
# Replication Management Agent.
#
# See "rsge.bootstrap.tds.port.number" properties in
# <SAP ASE installed directory>/DM/RMA-16_0/instances/AgentContainer/config/
bootstrap.prop
# for value
site1.rma_tds_port=7001
#
# Starting port number to use when setup Replication Server.
# Make sure next two ports (+1 and +2) are also available for use.
site1.srs_port=5005
# Device buffer for Replication Server on host1
# Recommend size = 128 * N
#   where N is the number of databases to replicate,
#   including the master and cluster ID databases.
#
site1.device_buffer_dir=/work/SAP1/data
site1.device_buffer_size=512
# Persistent queue directory for Replication Server running on host1
#
# For synchronous replication (synchronization_mode=sync),
# enter directory to an SSD (solid state drive) or other
# type of fast read/write storage device
site1.simple_persistent_queue_dir=/work/SAP1/data
site1.simple_persistent_queue_size=2000
# User defined dm data dir
site1.dm_database_file_directory=
site1.dm_translog_file_directory=
site1.dm_log_file_directory=
site1.dm_config_file_directory=
site1.dm_backup_file_directory_for_database=
#####
# Site "site2" on host host2 with companion role
#####
# Host name where SAP ASE run
#

```

```

# Enter fully qualified domain name (FQDN)
# if your sites are on different subnet.
site2.ase_host_name=SJMACHINE.BIG.corp
# Site name
#
# Enter value that identifies this site,
# like a geographical location.
# Value must be unique.
site2.site_name=SJHADR2
# Site role
#
# Enter the role of this site.
# Valid values: primary, companion
site2.site_role=companion
# directory where SAP ASE installed
site2.ase_release_directory=/work/SAP2
# Directory that stored SAP ASE user data files
# (interfaces, RUN_<server>, error log, etc. files).
# Do not set value if your user data files are in
# SAP ASE installed directory (ase_release_directory).
site2.ase_user_data_directory=/work/SAP2/
site2.ase_server_name=SJSAP2
site2.ase_server_port=5000
site2.backup_server_name=SJSAP2_BS
site2.backup_server_port=5001
# Directory to store database dumps
# in materialization
#
# Backup server must able to access this directory
site2.backup_server_dump_directory=/work/SAP2/data
# Data & log devices to create the databases specified
# in "participating_database_[x]" properties. You do
# not need to specify these properties if the database(s)
# already exist in the SAP ASE server.
#
# ase_data_device_create [x] [y] - property to create data device
# ase_log_device_create [x] [y] - property to create log device
#     where
#     x is number in "participating_database_[x]" property
#     y is number device to create
#
# Format: <logical_device_name>, <physical_device_path>, <size_in_MB>
#
# NOTE: Databases sizes on primary and companion
#     SAP ASE must be the same.
# Devices for database "LE1" (See "participating_database_1" property)
# Database size = 25MB
#     data device "le_data_dev" = 25MB
site2.ase_data_device_create_1_1=AS1_data_dev, /work/SAP2/data/le1_dev1.dat, 25
# Devices for database "userdb1" (See "participating_database_2" property)
# Database Size = 100MB
#     data device 1 "db1_data_dev1" = 25MB
#     data device 2 "db1_data_dev2" = 25MB
#     data device 3 "db1_data_dev3" = 25MB
#     log device 1 "db1_log_dev1" = 25MB
# site2.ase_data_device_create_2_1=db1_data_dev1, /host2_eng/ase/data/
db1_dev1.dat, 25
# site2.ase_data_device_create_2_2=db1_data_dev2, /host2_eng/ase/data/
db1_dev2.dat, 25
# site2.ase_data_device_create_2_3=db1_data_dev3, /host2_eng/ase/data/
db1_dev3.dat, 25
# site2.ase_log_device_create_2_1=db1_log_dev1, /host2_eng/ase/data/
db1_dev1.log, 25
# Devices for database "userdb2" (See "participating_database_3" property)
# Database Size = 100MB
#     data device 1 "db2_data_dev1" = 25MB
#     data device 2 "db2_data_dev2" = 25MB
#     log device 1 "db2_log_dev1" = 25MB

```

```

#         log device 2 "db2_log_dev2" = 25MB
# site2.ase_data_device_create_3_1=db2_data_dev1, /host2_eng/ase/data/
db2_dev1.dat, 25
# site2.ase_data_device_create_3_2=db2_data_dev2, /host2_eng/ase/data/
db2_dev2.dat, 25
# site2.ase_log_device_create_3_1=db2_log_dev1,    /host2_eng/ase/data/
db2_dev1.log, 25
# site2.ase_log_device_create_3_2=db2_log_dev2,    /host2_eng/ase/data/
db2_dev2.log, 25
# Port numbers for Replication Server and Replication Management Agent on host2
#
# In remote topology, these are the companion Replication Server and
# Replication Management Agent.
#
# See "rsge.bootstrap.tds.port.number" properties in
# <SAP ASE installed directory>/DM/RMA-16_0/instances/AgentContainer/config/
bootstrap.prop
# for value
site2.rma_tds_port=7001
#
# Starting port number to use when setup Replication Server.
# Make sure next two ports (+1 and +2) are also available for use.
site2.srs_port=5005
# Device buffer for Replication Server on host2
# Recommend size = 128 * N
#         where N is the number of databases to replicate,
#         including the master and cluster ID databases.
#
site2.device_buffer_dir=/work/SAP2/data
site2.device_buffer_size=512
# Persistent queue directory for Replication Server running on host2
#
# For synchronous replication (synchronization_mode=sync),
# enter directory to an SSD (solid state drive) or other
# type of fast read/write storage device
site2.simple_persistent_queue_dir=/work/SAP2/data
site2.simple_persistent_queue_size=2000
# User defined dm data dir
site2.dm_database_file_directory=
site2.dm_translog_file_directory=
site2.dm_log_file_directory=
site2.dm_config_file_directory=
site2.dm_backup_file_directory_for_database=

```

3.4 Configuring the SAP ASE Cockpit for the HADR System

When you first start the SAP ASE Cockpit, the *HADR* selection is activated if you installed the HADR system with the SAP installer. If you did not, register and authenticate the RMA to activate the HADR menu.

Prerequisites

The cockpit server is running.

Procedure

- Using a Web browser, connect to the SAP ASE Cockpit with this syntax:

```
https://<SAP_ASE_host>:<http_port>\cockpit\
```

Where:

- <SAP_ASE_host> is the name of the host on which you installed SAP ASE.
- <http_port> is the HTTPS cockpit port number you specified during installation (the default is 4283).

For example:

```
https://SFMACHINE1:4283\cockpit\
```

- Log in as the system administrator (the user with the sa_role).
- If you did not use the SAP installer or set SAP ASE for the monitoring options, the first time you log in, the SAP ASE Cockpit specifies a number of configuration parameters to reset.

The screenshot shows the SAP ASE Cockpit interface with a configuration dialog box open. The dialog box contains the following table:

Name	Current value	Required value
SQL batch capture	0	1
deadlock pipe active	0	1
deadlock pipe max messages	0	200
enable monitoring	0	1
enable spinlock monitoring	0	1
object lockwait timing	0	1
per object statistics active	0	1
process wait events	0	1
sql text pipe active	0	1
sql text pipe max messages	0	2000

The dialog box also includes an 'OK' button and a 'Cancel' button. The background interface shows various monitoring charts and system status information.

Reset these configuration parameters to the values indicated:

- SQL batch capture – 1
- deadlock pipe active – 1
- deadlock pipe max messages – 200

- enable monitoring-1
- enable spinlock monitoring-1
- max SQL text monitored-256
- object lockwait timing-1
- per object statistics active-1
- process wait events-1
- sql text pipe active-1
- sql text pipe max messages-2000
- statement pipe active-1
- statement pipe max messages-2000
- statement statistics active-1
- wait event timing-1

Click *OK*.

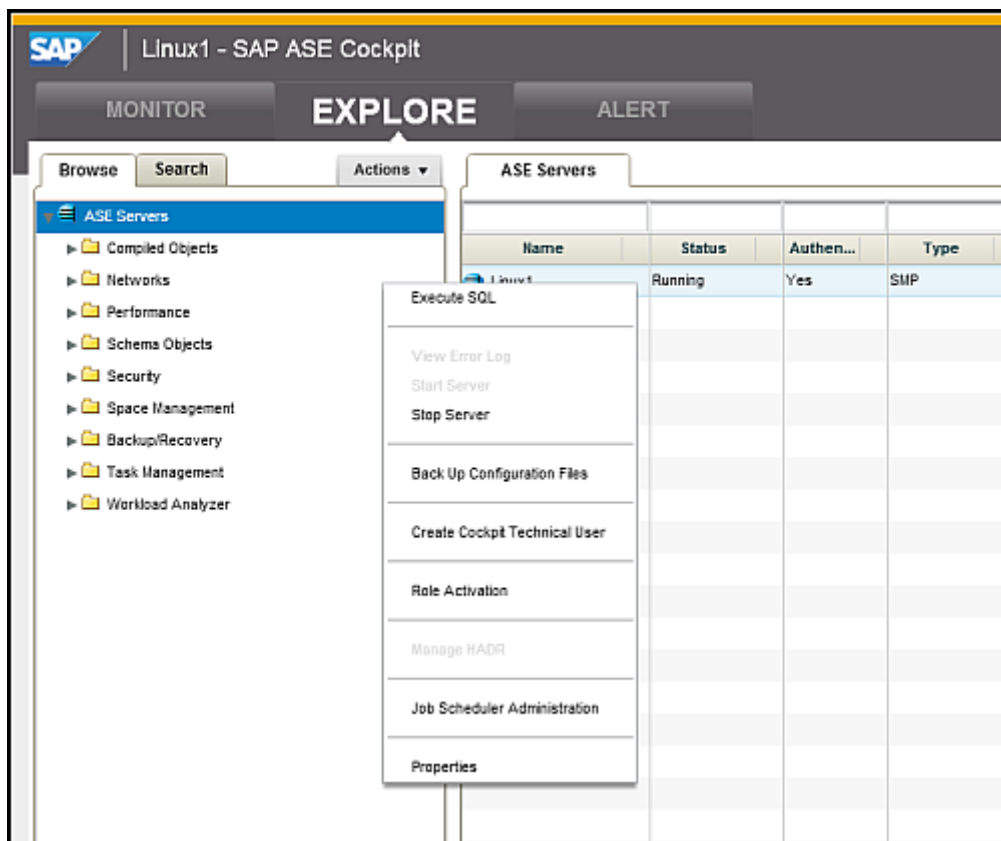
i Note

When you first start the SAP ASE Cockpit, there is a delay of up to 15 minutes before you see the statistics charts on the Monitoring tab (for example, the latency chart).


4. If necessary, restart the servers:
 - a. Select the *Explore* tab.
 - b. Select ► *ASE Servers* ► *<server_name>* ► *Start Server* ►.
 - c. Once the server is shut down, select ► *ASE Servers* ► *<server_name>* ► *Start Server* ►.
5. If necessary, create or configure the technical user. If you did not configure the Cockpit technical user in the SAP ASE Cockpit, configure it now.
 - a. Select the *Explore* tab.
 - b. Select ► *ASE Servers* ► *<server_name>* ► *Create Cockpit Technical User* ►.

i Note

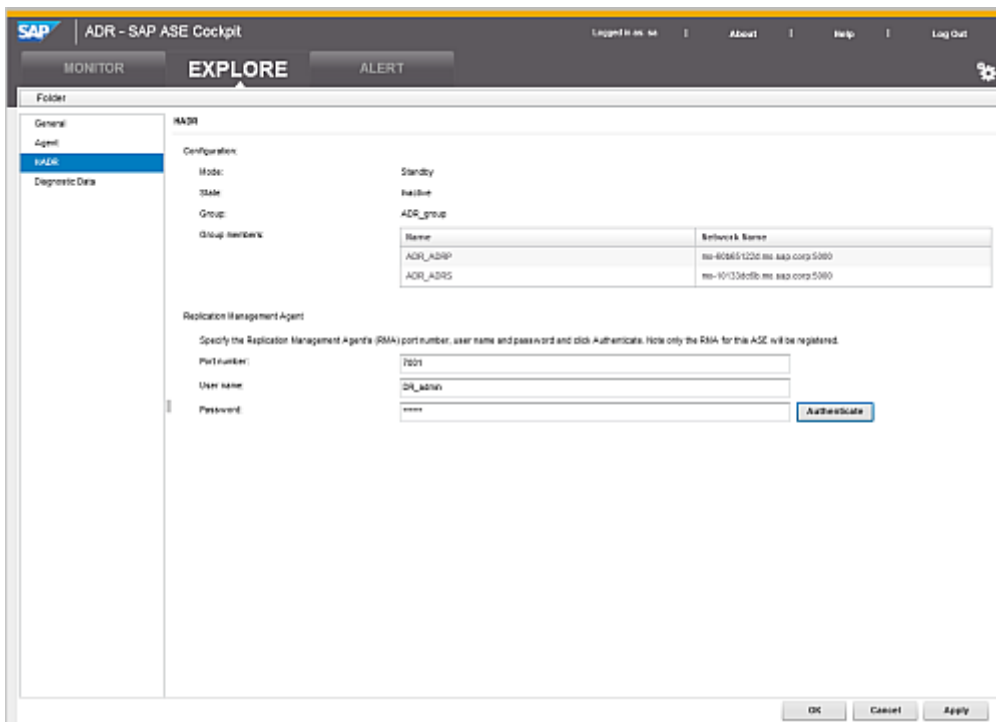
If the cockpit technical user already exists (for example, during an install using the SAP installer), the menu item reads *Update Cockpit Technical User*, which includes a wizard for updating a technical user in SAP ASE for cockpit service.



- c. Follow the prompts for the *Create Technical User* wizard.
6. (Optional, performed by SAP installer during typical install) Register the Replication Management Agent (RMA):
 - a. Select the *Explore* tab.
 - b. Select **ASE Servers** > <server_name> > *Properties* .
 - c. Select the *HADR* tab. The installer displays the Register Agent screen:

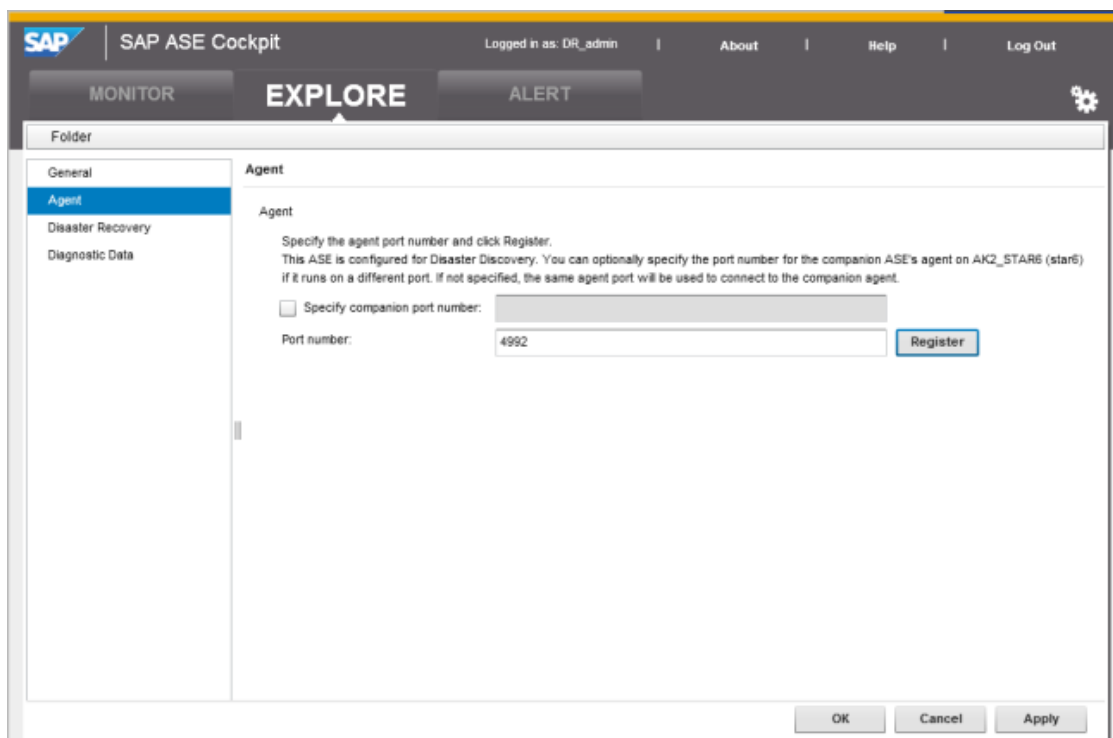


- d. Enter user name and password for the RMA administrator user (typically DR_admin)
- e. Select *Authenticate*.
1. Select the *Explore* tab.
2. Select **ASE Servers** > <server_name> > *Properties*.
3. Select the *HADR* tab. The installer displays the Register Agent screen:



4. Enter user name and password for the RMA administrator user (typically DR_admin).

5. Select *Authenticate*.
7. (Optional, performed by SAP Installer during typical install) Register the Agent.
 - a. Select the *Explore* tab.
 - b. Select **ASE Servers** > <server_name> > *Properties* >.
 - c. Select the *Agent* tab. The installer displays the Register Agent screen:



- d. Enter the agent port number (default is 4992) and click *Register*.
- e. Enter `uafadmin` as the user name and enter a password.
- f. Select *Authenticate*.

3.4.1 Configuration for Monitoring

When you first log into the SAP ASE Cockpit, you may see a *Configuration Parameters Validation* screen that displays the current values and the recommended values needed to enable monitoring. Change the current values to the recommended values, or to values that are consistent with the level of activity on the monitored SAP ASE for nonbinary values. You must have `sa_role` to change parameter values.

The `DR_admin` login does not have the `mon_role`. If a user logs in as `DR_admin`, the SAP ASE Cockpit displays an alert stating that the `mon_role` is missing, and general monitoring functionalities (such as Engine CPU Utilization) will be blank. However, HADR system monitoring is functional because this is performed by RMA.

You can assign the `mon_role` to `DR_admin`. See *Security > Manage Logins > Granting Roles to a Login* in the SAP ASE Cockpit online help.

3.5 Installing The Fault Manager

You can install the Fault Manager using the GUI installer or from the command line.

3.5.1 Installing and Configuring the Fault Manager

The Fault Manager is located on a machine that is separate from the machine hosting the primary and secondary companions, and uses a separate installer.

Prerequisites

i Note

The Fault Manager installer response file is automatically generated when you complete the HADR configuration on the second site. The response file is located in `$SYBASE/log/fault_manager_responses.txt`. Use this syntax to use this response file to install Fault Manager on third host:

```
<ASE_installer_directory>/FaultManager/setup.bin -f  
<fault_manager_responses.txt>
```

Where `<response_file_name>` is the absolute path to the generated responses file.

i Note

The steps below describe the installation process using the SAP installer. Often, it is much easier to use a response file to install the Fault Manager because many of the values are automatically filled out when you install and configure the servers (for example, primary and standby hosts, primary and standby SAP ASE directories and port numbers, primary and standby RMA hosts and port numbers, SAP ASE cockpit hosts and port numbers, and so on). See [Using setup.bin or setupConsole.exe with a Response File \[page 67\]](#) for information about installing and configuring the Fault Manager with a response file.

Keep in mind:

- The SAP ASE user provided during the Fault Manager installation must have the `sa_role` and `replication_role`. Grant these roles by logging into the primary server with `isql` and issuing (this example grants the roles to a user named `fmuser`):

```
create login fmuser with password "Sybase123"  
go  
grant role sa_role to fmuser  
go  
grant role replication_role to fmuser  
go  
sp_adduser fmuser  
go
```

- The Fault Manager log files are located in `<installation_directory>/log/FaultManager.log`.
- Configure SAP ASE server, Replication Management Agent, Replication Server, and SAP Host Agent on the primary and companion sites, and make sure they are running.
- The SAP ASE user you provide during the Fault Manager installation must have the `sa_role` and `replication_role`. Grant these roles by logging into the primary server with `isql` and issuing (this example grants the roles to a user named `fmuser`):

```
create login fmuser with password "Sybase123"
go
grant role sa_role to fmuser
go
grant role replication_role to fmuser
go
sp_adduser fmuser
go
```

- On the HP platform, the Fault Manager requires the C++ `libCsup11.so.1` library.
- On Linux, the Fault Manager requires GLIBC version 2.7 or later.
- The node running the Fault Manager must use the same platform as the HADR system nodes (however, it need not have the same operating system version).
- Set the number of file descriptors to 4096 (or higher) to start the Fault Manager. The default value for many systems is 1024. To determine the number of file descriptors to which your system is set, enter:
 - On the C-shell:

```
limit descriptors
```

- On the Bourne shell:

```
ulimit -n
```

To change the value for the number of file descriptors, enter:

- On the C-shell:

```
limit descriptors 4096
```

- On the Bourne shell:

```
ulimit -n 4096
```

- (AIX only) Set the data size limit to "unlimited":
 - On the Bourne shell – `ulimit -d unlimited`
 - On the C-shell – `limit datasize unlimited`
- (UNIX only) The `sapadm` operating system user requires a valid password. This user is created when the SAP Host Agent is installed by the SAP ASE installer but must be manually assigned a password.

Procedure

1. Run the Fault Manager installer from `<ASE_installer>/FaultManager/`, where `<fault_manager_responses.txt>` is from the second site's response file. `<fault_manager_responses.txt>` must use an absolute path to the generated responses file.

```
./setup.bin -f <fault_manager_responses.txt>
```

If you did not generate a response file, you can edit the sample response file. See [Using setup.bin or setupConsole.exe with a Response File \[page 67\]](#).

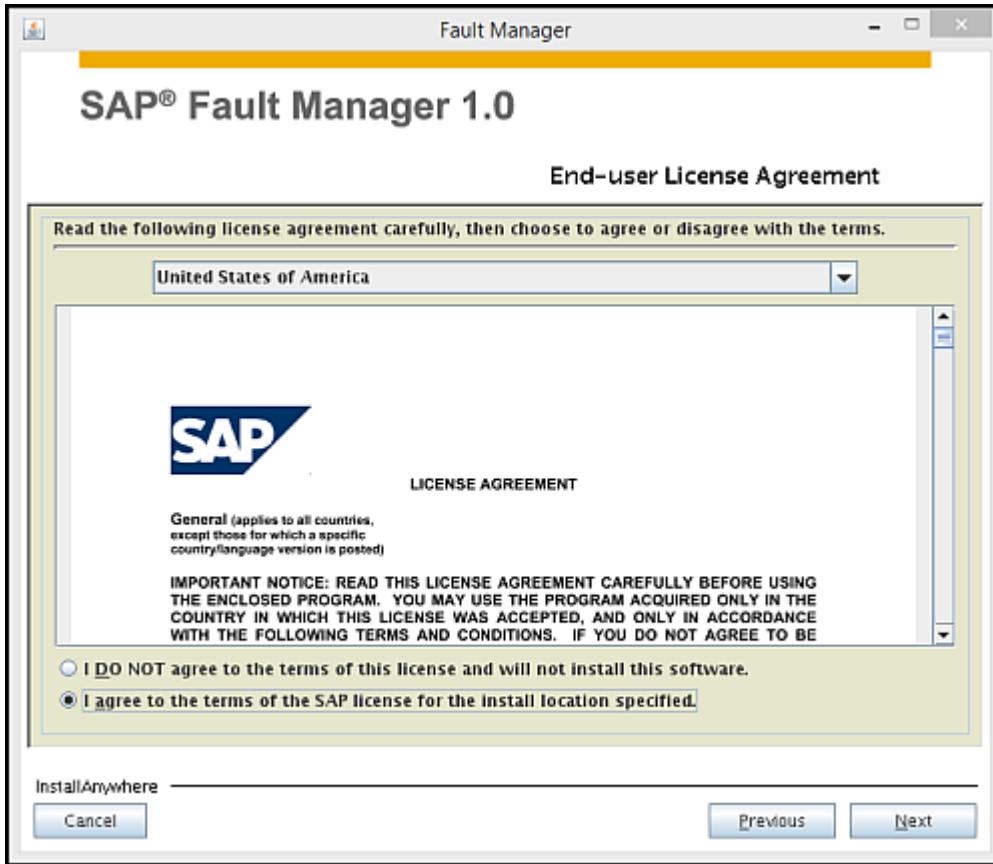
Alternatively, run the installer without a response file:

```
./setup.bin
```

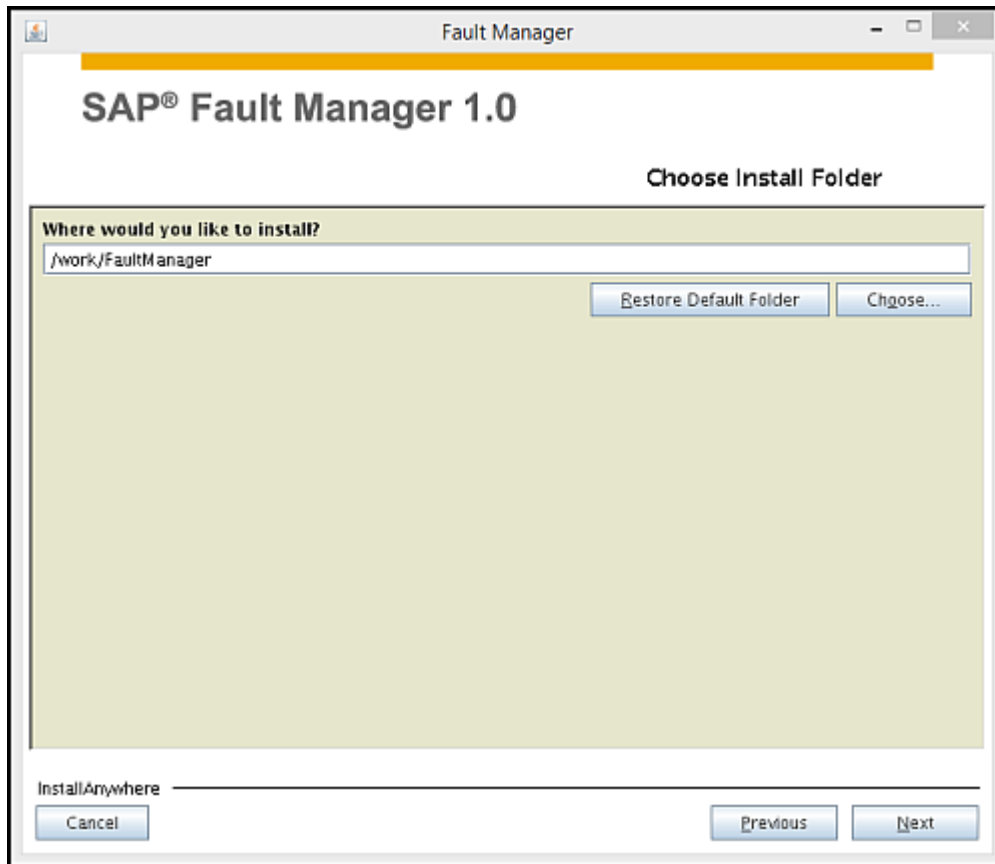
2. On the Introduction screen, click [Next](#).



3. In the End-user License Agreement screen, select the geographic location, and agree to the license terms. Click [Next](#).

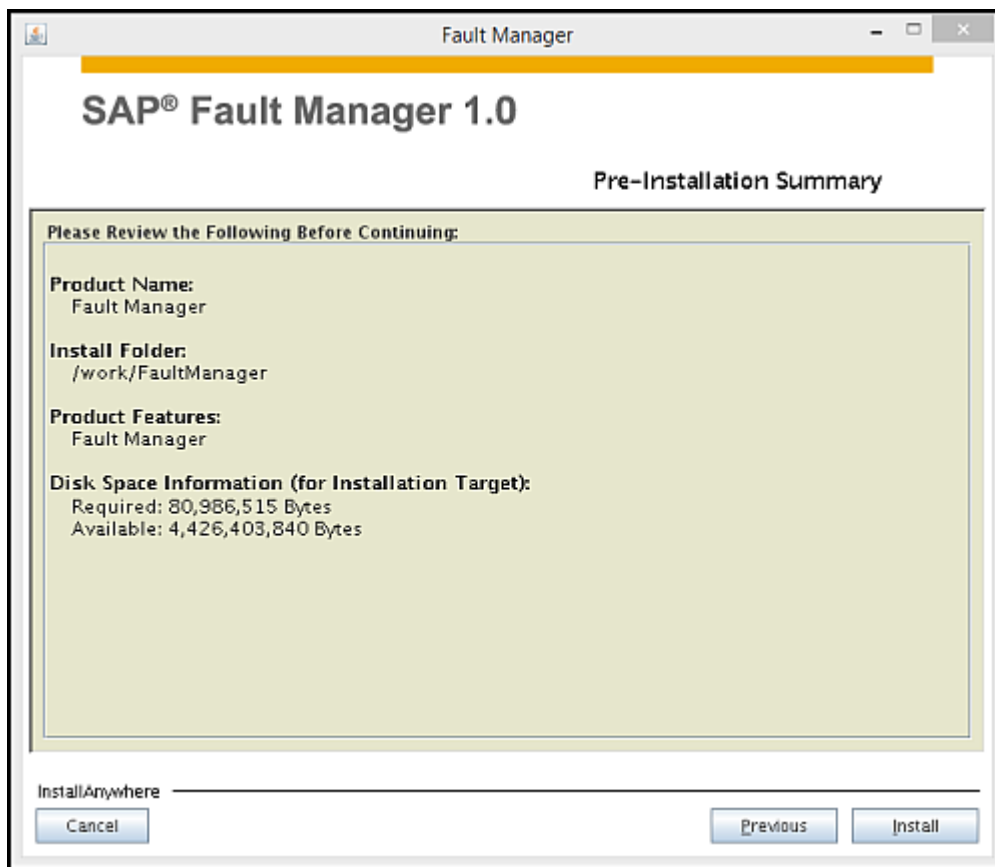


4. In the Choose Install Folder screen, accept the default directory or enter a new path to specify where to install Fault Manager.



Do not use the following when specifying the installation path; the installer does not recognize these, and displays an error:

- Double-byte characters
 - Single- or double-quote characters
5. On the Pre-Installation Summary screen, review the configuration summary. Click *Previous* to make changes. Click *Next* when you are ready to install the Fault Manager.



i Note

Installing the Fault Manager can take several minutes.

6. When the Fault Manager has installed, you see the Configure Fault Manager screen. Click *Yes* to configure the Fault Manager. If you click *No*, you can manually configure the Fault Manager later time using `sybdbfm` utility. See [Configuring the Fault Manager from the Command Line \[page 147\]](#).



Click *Next*.

7. In the Cluster ID screen, enter the ID of the cluster on which you are configuring the Fault Manager.



Click [Next](#).

8. In the Failover screen, choose whether to perform automatic failover if the primary SAP ASE is unreachable. The default is [Yes](#).



Click [Next](#).

9. In the SAP ASE on Primary screen, enter values for the primary site:

Option	Description
Site name	Name of the primary site.
SAP ASE host name	Host name of the primary SAP ASE.
SAP ASE name	Primary SAP ASE server name.
SAP ASE port	Default is 5000.
SAP ASE installed directory	Full path to the Adaptive Server release directory for the primary companion (\$SYBASE on UNIX, %SYBASE% on Windows).
SAP ASE installed user	ID of the user who installed the primary companion.

SAP® Fault Manager 1.0

SAP ASE on Primary

Enter the SAP ASE you have setup on primary site.

Site name

SAP ASE host name

SAP ASE Name

SAP ASE port

SAP ASE installed directory

SAP ASE installed user

InstallAnywhere

Click *Next*.

10. In the SAP ASE on Companion screen, enter values for the companion site:

Option	Description
Site name	Name of the companion site.
SAP ASE host name	Host name of the companion SAP ASE.
SAP ASE name	Companion SAP ASE server name.
SAP ASE port	Default is 5000.
SAP ASE installed directory	Full path to the SAP ASE release directory for the secondary companion (\$SYBASE on UNIX, %SYBASE% on Windows).
SAP ASE installed user	ID of the user who installed the secondary companion.

SAP® Fault Manager 1.0

SAP ASE on Companion

Enter the SAP ASE you have setup on companion site.

Site name

SAP ASE host name

SAP ASE Name

SAP ASE port

SAP ASE installed directory

SAP ASE installed user

InstallAnywhere _____

Click [Next](#).

11. Use virtual IPs (VIP) to redirect clients to the new primary server if you want to support older clients that cannot be linked with the new HADR or HA-aware SDK. You can create a VIP that is bound to the primary host and Fault Manager can move this VIP to new primary server as part of failover. The client application is aware only of the VIP, and when it attempts to reconnect, it will be connecting to the new primary server. Click [Yes](#) to enable the virtual IP.

Enter values for:

Option	Description
Virtual ASE host name	Name of the virtual host running SAP ASE.
Virtual ASE port	Default is 5000.
Network interface	Enter the network interface used.
Internet protocol version	Select the Internet protocol from drop-down list.

The screenshot shows the 'Virtual IP for SAP ASE' configuration window in SAP Fault Manager 1.0. The window has a title bar with 'SAP® Fault Manager 1.0' on the left and 'Virtual IP for SAP ASE' on the right. The main content area is titled 'Enable SAP ASE Virtual IP' and contains the following fields:

- Radio buttons for 'Yes' (selected) and 'No'.
- Text input for 'Virtual SAP ASE host name' with the value 'SFMACHINE'.
- Text input for 'Virtual SAP ASE port' with the value '5000'.
- Text input for 'Network interface' with the value 'eth0'.
- Dropdown menu for 'Internet protocol version' with 'ipv4' selected.

At the bottom of the window, there is a section labeled 'InstallAnywhere' with a horizontal line. Below this line are three buttons: 'Cancel' on the left, 'Previous' in the middle, and 'Next' on the right.

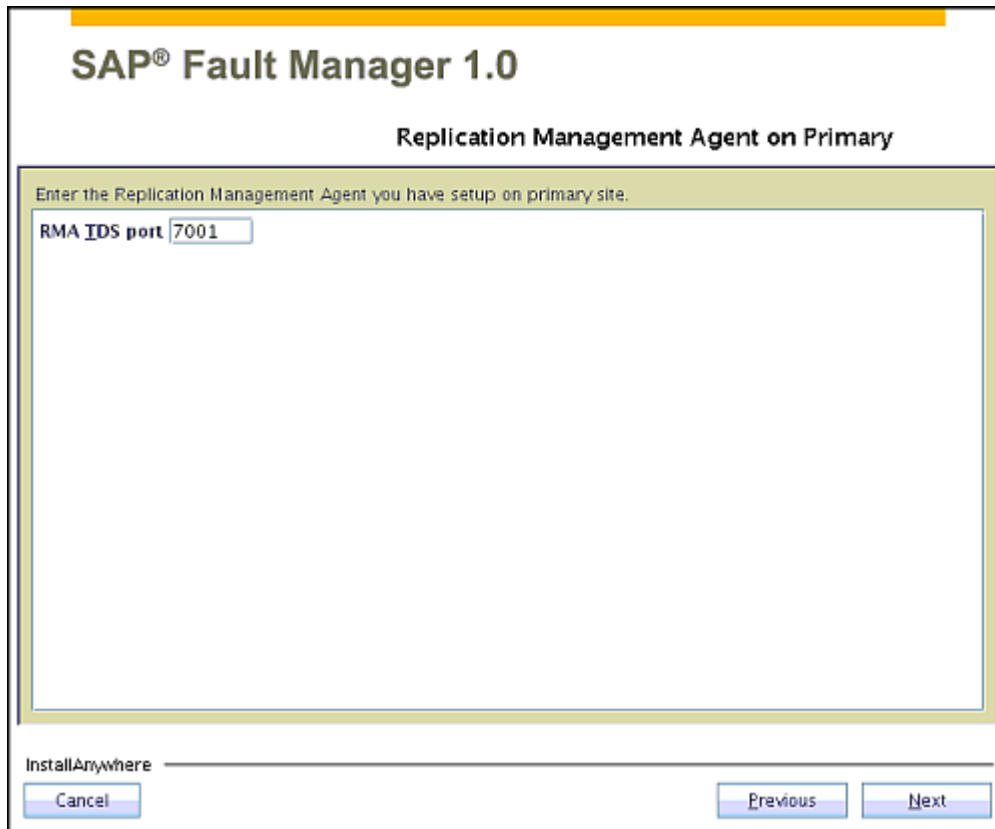
Click [Next](#).

12. On the ASE Cockpit screen, indicate whether to use SAP ASE Cockpit to manage the SAP ASE database HADR system. If you select [Yes](#), enter the primary and companion cockpit TDS ports.

The screenshot shows a dialog box titled "SAP® Fault Manager 1.0" with a sub-header "ASE Cockpit". The main text reads: "If you want to use ASE Cockpit to manage ASE HADR, enter the port numbers to the Cockpit you have setup on primary and companion sites." Below this is a question: "Do you want to use ASE Cockpit to manage ASE HADR?". There are two radio buttons: "Yes" (selected) and "No". Below the radio buttons are two input fields: "primary Cockpit TDS port" with the value "4998" and "companion Cockpit TDS port" with the value "4998". At the bottom left, there is a "Cancel" button. At the bottom right, there are "Previous" and "Next" buttons. The text "InstallAnywhere" is visible above the "Cancel" button.

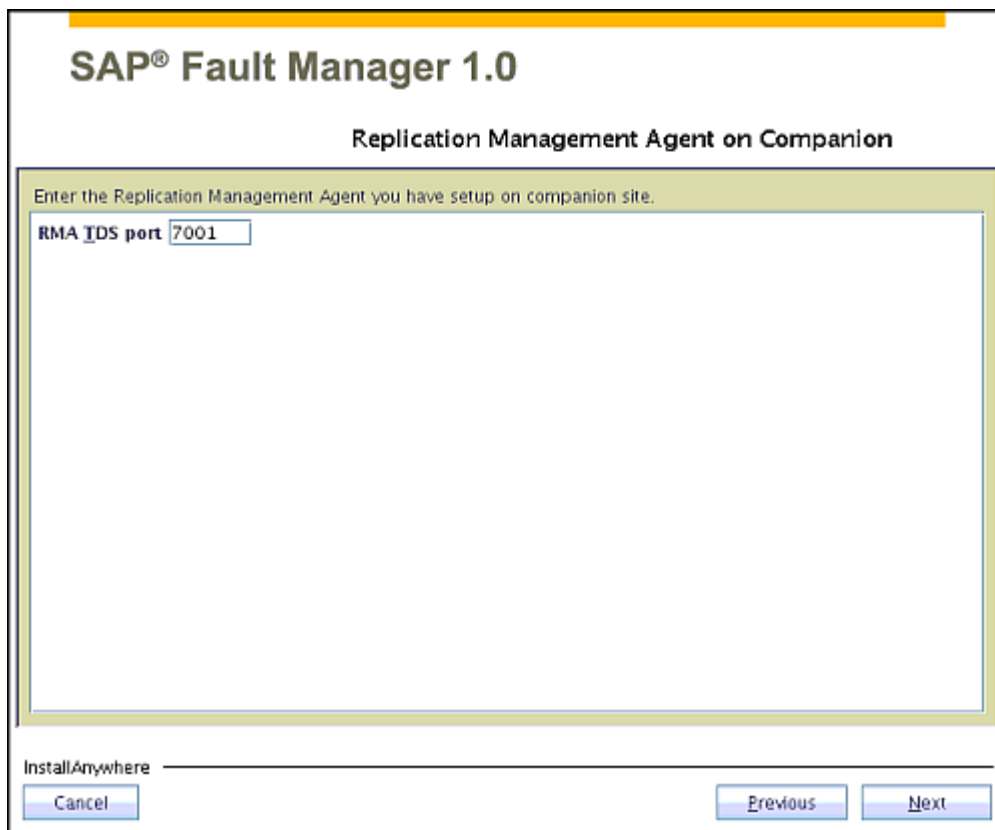
Click *Next*.

13. On the Replication Management Agent on Primary screen, enter the RMA TDS port number for connecting to the Replication Management Agent for the primary site. The default is 7001.



Click *Next*.

14. On the Replication Management Agent on Companion screen, enter the RMA TDS port number for connecting to the Replication Management Agent for the secondary companion site. The default is 7001.



Click *Next*.

15. In the Fault Manager Hosts and Ports screen, specify the Fault Manager host name, then its heartbeat port values for the Fault Manager (the primary Fault Manager heartbeat port number and the companion Fault Manager heartbeat port number cannot be the same):

Option	Description
Fault Manager heartbeat to heartbeat port	Default is 13797.
Primary Fault Manager heartbeat port	Default is 13777.
Secondary Fault Manager heartbeat port	Default is 13787.

SAP® Fault Manager 1.0

Fault Manager Host and Ports


Enter the Fault Manager host and ports.

Fault Manager host	<input type="text" value="FMHOST"/>
Fault Manager heartbeat to heartbeat port	<input type="text" value="13797"/>
Primary Fault Manager heartbeat port	<input type="text" value="13777"/>
Companion Fault Manager heartbeat port	<input type="text" value="13787"/>

InstallAnywhere

Click [Next](#).

16. In the Secure Store Directory screen, enter the full path to the directory for secure storage and click [Next](#).



17. In the Users for ASE HADR screen, enter the names and passwords for the users you used to set up the HADR system:
- *ASE system administrator user* – Name of the user authorized to administer SAP ASE. This user must have the sa_role and replication_role.
 - *ASE system administrator user password*
 - *Confirm system administrator user password*
 - *RMA administrator user* – Name of the user authorized to administer Replication Server.
 - *RMA administrator password*
 - *Confirm RMA administrator password*
 - *SAP Host Agent user* – Name of the user authorized to administer the host agent.
 - *SAP Host Agent user password*
 - *Confirm SAP Host Agent user password*
 - If you chose to use the SAP ASE Cockpit to manage the database HADR system, enter:
 - *Cockpit administrator user* – Name of the user authorized to administer SAP ASE Cockpit.
 - *Cockpit administrator password*
 - *Confirm Cockpit administrator user password*

Click [Next](#).

18. Review the configuration summary. Click [Previous](#) to make changes: Click [Next](#) to proceed with the configuration.

19. Click [Done](#).

20. (Optional) Check the Fault Manager status.

- a. Source the `$SYBASE/SYBASE.csh` (or `SYBASE.sh` for the Bourne shell) file to configure the environment variables.
- b. Move to the directory you specified for secure storage in the Secure Store Directory screen.
- c. Issue:

```
$SYBASE/FaultManager/bin/sybdbm status
```

3.5.2 Fault Manager Behavior During Automatic Failovers

The Fault Manager enters a bootstrap mode when you start it using the `hadm` command.

During the bootstrap mode, the Fault Manager requires that both SAP ASE nodes are running and configured for HADR, performs various operations, and checks on both nodes. Once the bootstrap is complete, Fault Manager continuously monitors the HADR system, and generates sanity reports until you stop the Fault Manager.

The Fault Manager operates in a revolving cycle of three phases:

- Diagnostic (DIAGNOSE) – Checks the health of all components of the HADR system.
- Action (ACTION) – If the Fault Manager finds that any of components are nonresponsive or have stopped during the Diagnostic phase, it takes the appropriate action to rectify the problem.
- Pause (PAUSING) - The Fault Manager sleeps for three seconds.

This scenario illustrates how the Fault Manager reacts during a failover. The scenario assumes you have two sites: London, which contains the primary server, and Paris, which contains the companion server. HADR is running normally, and replication is in synchronous mode.

If the London site becomes unavailable unexpectedly, and the Fault Manager `ha/syb/use_cockpit` profile parameter is set to 1, the Fault Manager sends this message to the SAP ASE Cockpit:

```
Primary host (London) cannot be contacted
```

If the Fault Manager `ha/syb/failover_if_insync` profile parameter is set to 1, Fault Manager automatically triggers an unplanned automatic failover to the companion server, Paris. When failover is initiated, the Fault Manager sends this message to the SAP ASE Cockpit:

```
Failover initiated from London to Paris
```

When the failover is complete, Fault Manager send this message to the SAP ASE Cockpit:

```
Failover succeeded from London to Paris
```

When the Paris companion server becomes the new primary companion, the Fault Manager sends this message to the SAP ASE Cockpit:

```
Primary ASE (Paris) contact is restored
```

The Fault Manager should be able to contact the London site when it becomes available. When contact is restored, the Fault Manager sends this message to the SAP ASE Cockpit:

```
Standby ASE (London) contact is restored
```

Although the Fault Manager should recognize the London site as a companion server, it is not yet available in the HADR system because replication is not yet restored. In this situation, the Fault Manager sends one of these messages to the SAP ASE cockpit:

```
Replication is IN DOUBT
Replication is DOWN
Replication is SUSPENDED
```

You then use the RMA to restore replication and make the London companion site available. If the Fault Manager `ha/syb/set_standby_available_after_failover` profile parameter is set to 1, the Fault Manager makes the host available using the `sap_host_available` RMA command. However, if the Fault Manager `ha/syb/set_standby_available_after_failover` profile parameter is set to 0, you manually issue the `sap_host_available` command from RMA.

Regardless of how it is issued, the `sap_host_available` command restores the replication in synchronous mode and completes the HADR restoration.

If a network issue occurs while the primary server is disconnected from the network, the heartbeat client running on the primary host deactivates the primary site, and the Fault Manager promotes the standby node to the primary server in Active mode.

3.5.3 Customizing the Fault Manager

Not all profile parameters appear in the profile file after you configure the Fault Manager. The Fault Manager uses default values for any parameters you do not specify.

Use the profile parameters in the subsequent sections to customize the Fault Manager, adding the appropriate setting to the profile file. Restart the Fault Manager anytime you add or change a profile parameter. See [Administering the Fault Manager \[page 150\]](#) for information about starting the Fault Manager.

3.5.4 Fault Manager Profile Parameters

The Fault Manager includes numerous profile parameters, which are set in the Fault Manager profile file, and are similar to configuration parameters for SAP ASE.

The profile file is initially populated with defaults, but users can modify these defaults to match their site's configuration.

Note

Restart the Fault Manager after adding or modifying any profile parameter. Use `$$SYBASE/FaultManager/bin/sybdbm restart` to restart the Fault Manager.

Table 1: Parameters for Fault Manager Setup

Parameter Name	Default Value	Possible Values / Range	Description
ha/syb/version	1	1	
ha/syb/trace	1	<ul style="list-style-type: none">1 – Basic verbose output2 – Medium verbose output3 – Maximum verbose output	Trace level determines the amount of diagnostic output written to the Fault Manager log (named <code>dev_sybdbfm</code>).
ha/syb/dbfmhost	The host on which you loaded the Fault Manager binary	IP address or hostname	The host on which the Fault Manager runs.
ha/syb/upload_hbeat	1	<ul style="list-style-type: none">0 – Fault Manager will not attempt to upload heartbeat client on the HADR system nodes1 – Fault Manager will upload heartbeat client on the HADR system nodes	Determines whether to upload heartbeat client to the HADR system nodes. If the HADR system nodes already have heartbeat clients running, use this option to disable the upload.

Parameter Name	Default Value	Possible Values / Range	Description
ha/syb/non_bs	1	Set to 1 for a Custom Application HADR system.	Indicates the type of HADR installation.
<div style="border: 1px solid #ccc; background-color: #f0f0f0; padding: 5px;"> <p>i Note</p> <p>Do not change this parameter.</p> </div>			
ha/syb/exedir	The current working directory in which you run the ./sybdbfm installation	Any valid directory path	Indicates the directory from which you run the Fault Manager.
ha/syb/h2hport	13797	Any valid unused port number	Describes the port the heartbeat client uses to hear communication.
SAPSYSTEMNAME	SYB	SYB	Internal use only. Do not change this parameter.
ha/syb/hbeat	1	<ul style="list-style-type: none"> 0 – disables heartbeat (not recommended) 1 – enables using heartbeat client 	Enables or disables heartbeat client usage.
ha/syb/use_cockpit	1	<ul style="list-style-type: none"> 0 – disables all communication with SAP ASE Cockpit 1 – enables all communication with SAP ASE Cockpit 	Enables or disables communication and error reporting of Fault manager with SAP ASE Cockpit.

3.5.4.1 Parameters for Fault Manager Setup

There are a number of parameters that affect the Fault Manager configuration

Table 2: Parameters for Fault Manager Setup

Parameter Name	Default Value	Possible Values / Range	Description
ha/syb/hb_shutdown_sys	0	<ul style="list-style-type: none"> 0 – (default) disabled 1 – enabled 	Shuts down host if heartbeat discovers a network breakdown and forces database to an inactive state and stop database fail.
ha/syb/hb_force_db_inactive	1	<ul style="list-style-type: none"> 0 – disabled 1 – (default) enabled 	Forces the database to an inactive state if the heartbeat discovers network breakdown.

Parameter Name	Default Value	Possible Values / Range	Description
ha/syb/hb_stop_db	1	<ul style="list-style-type: none"> 0 – disabled 1 – (default) enabled 	Stops the database if the heart-beat discovers a network breakdown and forces the database to inactive fails.
ha/syb/bootstrap_retries	0	Any positive numeric value	Continue bootstrapping the Fault Manager if there is an error. Continues to retry bootstrapping until the number of retries reaches the set value.
ha/syb/trace_file_size	20	Any positive numeric value	Truncate development traces after this value (in megabytes) is reached, and backup old trace file to <file_name>.old.

3.5.4.2 Parameters for the Fault Manager's Primary Site

There are a number of parameters that affect the primary site's configuration.

Parameter Name	Default Value	Possible Values / Range	Description
ha/syb/primary_dbhost	Host that runs Fault Manager.	Can take IP Address or hostname.	Host where primary SAP ASE will run.
ha/syb/primary_dbname	No default; blank.	Cluster ID database name.	Cluster ID database name on primary SAP ASE.
ha/syb/primary_dbport	4901	Any valid unused port number.	Port on which the primary SAP ASE runs.
ha/syb/primary_cockpit_host	Host that runs Fault Manager.	Can take IP address or hostname.	Host on which the primary SAP ASE is installed.
ha/syb/primary_cockpit_port	4998	Any valid unused port number set during SAP ASE installation.	Port where the Fault Manager connects to SAP ASE Cockpit to pass event messages.
ha/syb/primary_site	No default, blank.	Any valid logical site name that can be given in RMA to this HADR system node.	Logical site name given in RMA for this HADR system node.
ha/syb/primary_hbport	13777	Any valid unused port number. (The primary Fault Manager heartbeat port number and the standby Fault Manager heartbeat port number cannot be the same.)	Port from which the primary site heartbeat client makes contact.

Parameter Name	Default Value	Possible Values / Range	Description
ha/syb/primary_dr_host	Host that runs Fault Manager.	Can take IP address or hostname.	Host from which the primary site RMA runs.
ha/syb/primary_dr_port	4909	Any valid unused port number.	Port from which the primary node RMA runs.
<div style="border: 1px solid #ccc; padding: 10px; background-color: #f9f9f9;"> <p>i Note</p> <p>The HADR installation uses a default port number of 7001 for the Custom Application, but a default port number of 4909 for the Fault Manager.</p> <p>The Custom Application response file includes the correct port number for the Fault Manager.</p> </div>			
ha/syb/primary/ase_instance_name	No default; blank.	Any valid SAP ASE server name.	SAP ASE server name.
ha/syb/primary/ase_instance_path	/sybase/	Any valid directory path to the SAP ASE installation directory.	Location where SAP ASE has been installed.
ha/syb/primary/ase_instance_user	syb	Any valid user with appropriate permissions.	<p>(UNIX) Username who installed SAP ASE, or the login with the appropriate permissions to access the SAP ASE installation.</p> <p>(Windows) This option is not present because the SAP ASE database service (SAP SQLServer) is used.</p>

3.5.4.3 Parameters for the Fault Manager Companion Site

There are a number of parameters that affect the companion site's configuration.

Parameter Name	Default Value	Possible Values / Range	Description
ha/syb/standby_dbhost	Host that runs Fault Manager.	Can take IP address or hostname.	Host on which the standby SAP ASE runs.
ha/syb/standby_dbname	No default; blank.	Cluster ID database name.	Cluster ID database name on standby SAP ASE.
ha/syb/standby_dbport	4901	Any valid unused port number.	Port on which the standby SAP ASE runs.

Parameter Name	Default Value	Possible Values / Range	Description
ha/syb/ standby_cockpit_host	Host that runs Fault Manager.	Can take IP address or hostname.	Host where standby SAP ASE is installed. SAP ASE Cockpit must know the install directory to run.
ha/syb/ standby_cockpit_port	4998	Any valid unused port number set during SAP ASE installation.	Port on which the Fault Manager connects to SAP ASE Cockpit to pass event messages to it.
ha/syb/standby_site	No default; blank.	Any valid logical site name that can be given in RMA to this HADR system node.	Logical site name given in RMA for this HADR system node.
ha/syb/standby_hbport	13787	Any valid unused port number. (The standby Fault Manager heartbeat port number and the primary Fault Manager heartbeat port number cannot be the same)	Port from which the standby heartbeat client makes contact.
ha/syb/standby_dr_host	Host that runs Fault Manager.	Can take IP address or hostname.	Host from which the standby RMA runs.
ha/syb/standby_dr_port	4909	Any valid unused port number.	Port from which the RMA runs on the standby node.
ha/syb/standby/ ase_instance_name	No default; blank.	Any valid SAP ASE server name.	SAP ASE server name.
ha/syb/standby/ ase_instance_path	/sybase/	Any valid directory path to the SAP ASE installation directory.	Location where SAP ASE was installed.

Note

The HADR installation uses a default port number of 7001 for the Custom Application, but a default port number of 4909 for the Fault Manager.

The Custom Application response file includes the correct port number for the Fault Manager.

Parameter Name	Default Value	Possible Values / Range	Description
ha/syb/standby/ase_instance_user	syb	Any valid user with appropriate permissions.	(UNIX) Username who installed SAP ASE, or the login with the appropriate permissions to access the SAP ASE installation. (Windows) This option is not present because the SAP ASE database service (SAP SQLServer) is used.

3.5.4.4 Parameters for the Fault Manager Communication Frequency Checks

There are a number of parameters that affect the frequency of the Fault Manager's communication checks.

Parameter Name	Default Value	Possible Values / Range	Description
ha/syb/check_frequency	3 (seconds)	Any positive numeric value	The unit of frequency upon which the other units are based.
ha/syb/primary_ping_frequency	1 (1 unit of check_frequency; that is, 3 seconds)	Any positive numeric value	Frequency of the primary database shallow probe.
ha/syb/standby_ping_frequency	10 (10 units of check_frequency. That is, 30 seconds)	Any positive numeric value	Frequency of the standby database shallow probe.
ha/syb/primary_hostctrl_status_frequency	100 (100 units of check_frequency; that is, 300 seconds)	Any positive numeric value	Frequency of the primary database deep probe.
ha/syb/standby_hostctrl_status_frequency	100 (100 units of check_frequency; that is, 300 seconds)	Any positive numeric value	Frequency of the standby database deep probe.
ha/syb/report_status_frequency	100 (100 units of check_frequency; that is, 300 seconds)	Any positive numeric value	Frequency with which the status of the components is reported.
ha/syb/replication_status_check_frequency	100 (100 units of check_frequency; that is, 300 seconds)	Any positive numeric value	Frequency of deep probe to receive replication status.

3.5.4.5 Parameters For The Fault Manager Communication Timeout

There are a number of parameters that affect when the Fault Manager's communication times out.

Parameter Name	Default Value	Possible Values / Range	Description
ha/syb/ start_database_timeout	Max integer	Any positive numeric value	Timeout period for attempts to start the primary and companion databases.
ha/syb/ stop_database_timeout	60 seconds	Any positive numeric value	Timeout period for attempts to stop the primary and companion databases.
ha/syb/failover_timeout	10 minutes	Any positive numeric value	Timeout period for attempts for failover.
ha/syb/ replication_status_timeout	180 seconds	Any positive numeric value	Timeout period to receive replication status.
ha/syb/ standby_available_timeout	Max integer	Any positive numeric value	Timeout period for sap_host_available execution
ha/syb/dbctrl_timeout	30 seconds	Any positive numeric value	Timeout period for sapdbcontrol calls (the host agent uses sapdbcontrol to execute database-related commands, such as getting the status of the database).
ha/syb/ odbc_connect_timeout	3 seconds	Any positive numeric value	Timeout period for the shallow probe connect.
ha/syb/ odbc_command_timeout	3 seconds	Any positive numeric value	Timeout period for the shallow probe execution.
ha/syb/ upload_executable_timeout	60 seconds	Any positive numeric value	Timeout period for uploading the heartbeat client to the primary or companion host.
ha/syb/hb_fm_timeout	3 seconds	Any positive numeric value	Timeout period for the heartbeat client to determine if connection to Fault Manager is lost.
ha/syb/hb_hb_timeout	3 seconds	Any positive numeric value	Timeout period for any heartbeat client to determine if connection to the other heartbeat client is lost.
ha/syb/ hb_set_db_inactive_timeout	10 seconds	Any positive numeric value	Timeout period for the heartbeat client to set the SAP ASE database to "inactive."
ha/syb/hb_kill_db_timeout	10 seconds	Any positive numeric value	Timeout period for the heartbeat client to kill the SAP ASE database.

3.5.4.6 Parameters that Affect the Fault Manager's Actions

There are a number of parameters that affect the actions the Fault Manager performs.

Table 3: Parameters for the Fault Manager Action

Parameter Name	Default Value	Possible Values / Range	Description
ha/syb/allow_restart	0	<ul style="list-style-type: none"> 0 – disabled 1 – enabled 	Allow restart of primary SAP ASE in case of failure and ASYNC_REP.
ha/syb/allow_restart_companion	0	<ul style="list-style-type: none"> 0 – disabled 1 – enabled 	Allow restart of companion SAP ASE if it goes down.
ha/syb/set_standby_available_after_failover	0	<ul style="list-style-type: none"> 0 – disabled 1 – enabled 	After failover, make newly appointed companion server available in HADR environment by automatically executing sap_host_available.
ha/syb/chk_restart_repserver	0	<ul style="list-style-type: none"> 0 – disabled 1 – enabled 	Allow restart of Replication Server and RMA if they shut down.
ha/syb/failover_if_insync	1	<ul style="list-style-type: none"> 0 – monitors and reports only on the error case 1 – monitors and fails over if the case occurs 	Allow failover or just monitor and report.
ha/syb/failover_if_unresponsive	0	<ul style="list-style-type: none"> 0 – disabled 1 – enabled 	Allow failover if SAP ASE is unresponsive.
ha/syb/run_odbc_threaded	1	<ul style="list-style-type: none"> 0 – disabled 1 – enabled 	Run the shallow probes in a special thread.
ha/syb/waittime_if_unresponsive	60	Any positive numeric value	Set wait time to determine unresponsiveness of SAP ASE.

3.5.4.7 Parameters for the Fault Manager Virtual IP / Floating IP

There are a number of parameters that affect the Fault Manager's virtual and floating IP address.

Parameter Name	Default Value	Possible Values / Range	Description
ha/syb/support_floating_ip	0	<ul style="list-style-type: none">0 – floating IP not supported1 – move floating IP/hostname	Option to toggle usage of floating IP.
ha/syb/vdbhost	No default; blank.	Can take IP address or hostname	Floating IP/hostname.
ha/syb/vdbport	4901	Any valid unused port number	Port number for the floating IP.
ha/syb/vdb_interface	Network interface set by database	Network interfaces available	Option to set network interface.
ha/syb/vdb_ipv	ipv4	<ul style="list-style-type: none">ipv4 – IP address version 4ipv6 – IP address version6	Option to set IP version.

3.5.5 Upgrading the Fault Manager

Upgrade the Fault Manager by performing a binary overlay of the existing Fault Manager.

Procedure

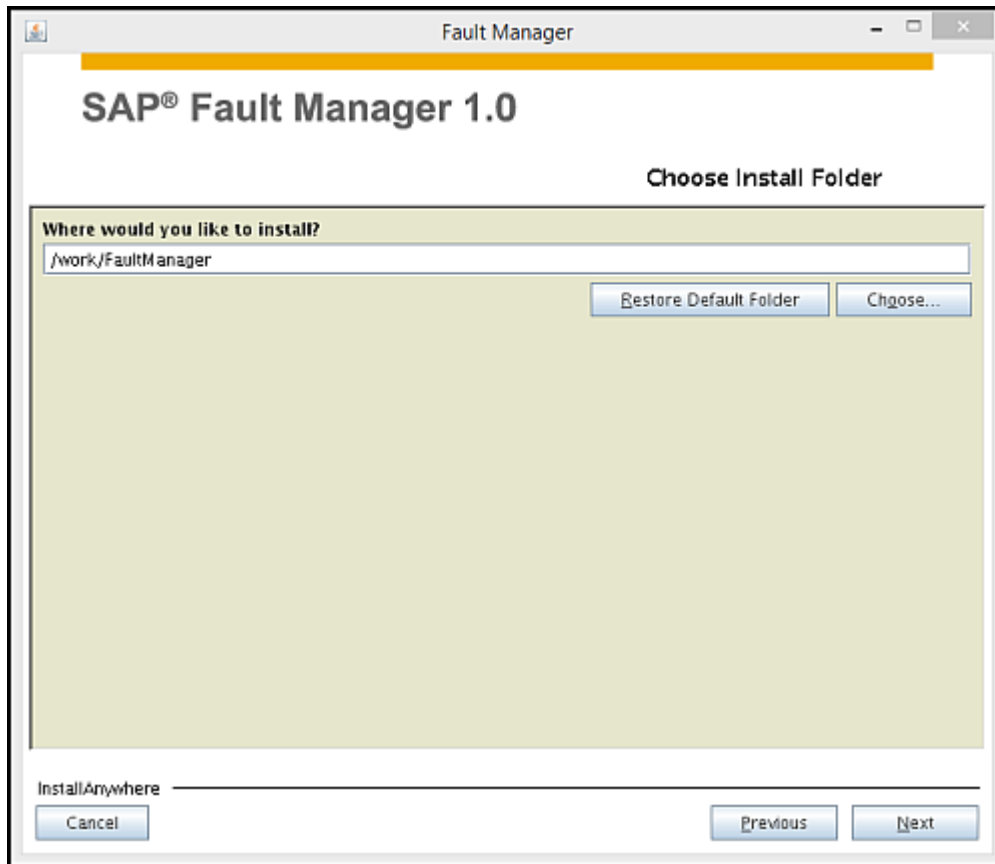
1. Stop Fault Manager:

```
<Fault_Manager_install_dir>/FaultManager/bin/sybdbfm stop
```

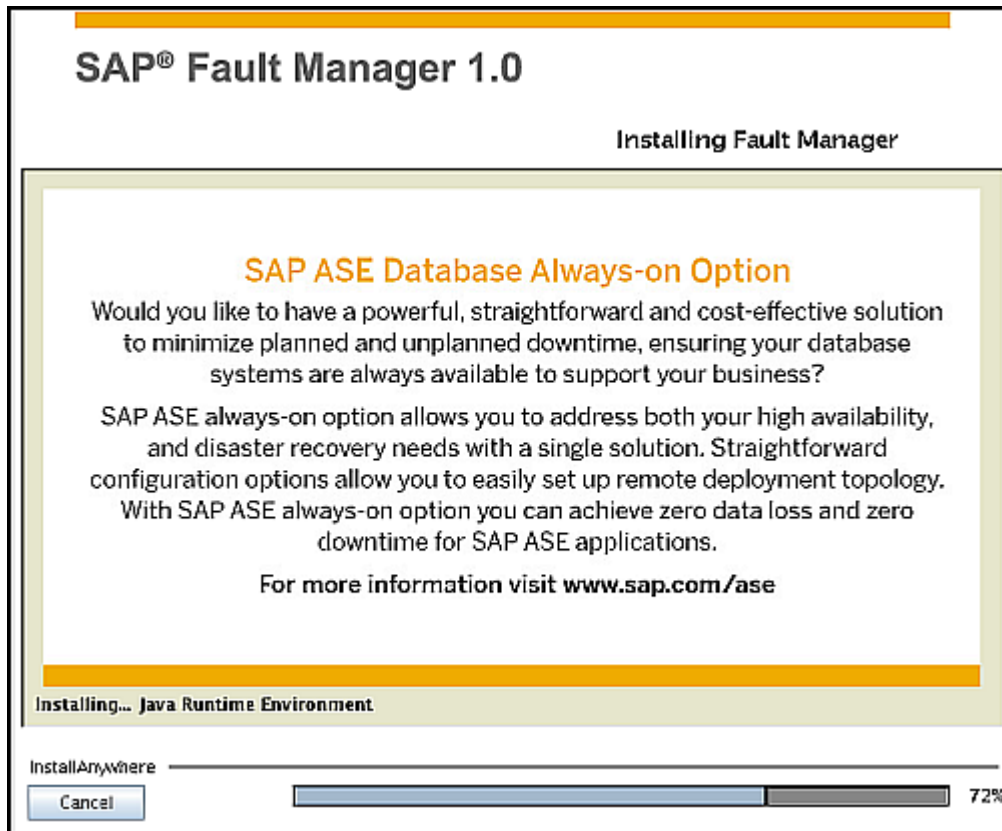
2. Run the Fault Manager installer from <ASE_installer>/FaultManager/:

```
<ASE_installer_directory>/FaultManager/setup.bin
```

3. Select the geographic location, and agree to the license terms. Click *Next*.
4. In the Choose Install Folder screen, enter the same installation path as the previous version of the Fault Manager.



5. Review the installation summary. Click *Previous* to make changes: Click *Install* for the installer to unload the files to the disk. The installer unloads the files.



6. In the Configure Fault Manager screen, select *No* to not configure the Fault Manager, but only overlay the binaries.



- Restart the Fault Manager:

```
<Fault_Manager_install_dir>/FaultManager/sybdbfm_<CID>
```

3.6 Using the Fault Manager with Custom Application

The Fault Manager monitors the health of the primary and standby servers, and triggers a failover if the primary server or host fails, and the HADR system is running in synchronous mode.

The Fault Manager is a standalone component that runs on a third node, preferably where the application server is running, and on the same platform as the HADR system nodes.

The Fault Manager functions in two modes: the Fault Manager mode and as the heartbeat client mode. The Fault Manager runs on a third host. In Fault Manager mode, it monitors SAP ASE, Replication Server, performs functions like initiating failover, and restarting the server, and acts as the server for the heartbeats that it receives from the heartbeat clients.

The Fault Manager heartbeat client mode runs on primary and standby hosts. In heartbeat client mode, the Fault Manager sends a heartbeat to the Fault Manager, checks for heartbeats from fellow heartbeat clients, and sends its own heartbeat to them (primarily to avoid a split-brain situation). If the heartbeat client on the primary host loses a connection with the Fault Manager and the fellow heartbeat client, the Fault Manager triggers a deactivation of the primary server. If the deactivation fails, the Fault Manager kills the SAP ASE process.

The Fault Manager checks the database state with the `saphostctrl` SAP host agent, which is a daemon process started on all participating nodes. The Fault Manager also uses `saphostctrl` to connect to the Replication Management Agent. See the chapter titled *The SAP Host Agent*.

The Fault Manager:

- Triggers a failover using `saphostctrl` if the primary server is down or if the primary node is down or unreachable, and the standby server is healthy and synchronously replicated.
- Restarts the primary server if it is down and replication is asynchronous.

Note

Stop or hibernate the Fault Manager when you perform any maintenance activity on SAP ASE or other components in the HADR system. Once hibernated, the Fault Manager process continues to run but will not monitor the database, and no failover occurs. The heartbeat processes are stopped during hibernation.

This table describes actions performed by the Fault Manager:

Scenario	Action Performed
Primary server is down.	Failover to companion, notify cockpit, if configured.
Primary host is down.	Failover to companion, notify cockpit, if configured.
Primary server is unreachable (network glitch or SAP ASE unresponsive).	<p>Failover to the companion if:</p> <ul style="list-style-type: none"> • SAP ASE is unreachable because it has become unresponsive and the <code>failover_if_unresponsive</code> parameter is set in the Fault Manager configuration file. • There is a network glitch. <p>Notify cockpit in both cases, if this is configured .</p>
Primary server reports an error condition.	<p>If client login and data access are unaffected, no action is taken. Fault Manager does not scan the SAP ASE log for errors.</p> <p>If the error results in login failures or data access errors, failover to companion.</p> <p>Login failures are identified by the Fault Manager or the SAP hostagent connecting to the SAP ASE servers.</p>
SAP ASE, Replication Server, or RMA on the companion host are down.	<p>Restart these components if the corresponding parameters in FM profile are set (for example, <code>chk_restart_repserver</code>).</p> <p>The Fault Manager uses the SAP hostagent to restart any of these services if they are down.</p>
Fault Manager is down.	For the Custom App version of HADR, manually restart the Fault Manager (see Administering the Fault Manager [page 150]). See the Business Suite documentation for instructions on restarting the Fault Manager.

Scenario	Action Performed
Failover fails.	Fault manager continues attempting failover until it succeeds. Manual intervention may be required to rectify the condition which is causing the failover to fail.
Fault Manager is unreachable from the 2 sites (primary and companion).	Automatic failover is disabled since the Fault Manager cannot reach the primary and companion sites. If the primary SAP ASE goes down in this situation, manual intervention is required to initiate failover. If the network between the primary and companion is satisfactory, then continue as is. However, if there is a network problem between primary and companion, deactivate the primary to avoid a split brain situation, and notify cockpit. In this situation, replication is turned off. SAP ASE cannot access to the database because there is no primary companion, and you must perform a manual intervention to resolve the issue.
Heartbeat from primary is missed for a preconfigured timeout.	Fault manager keeps trying to restart heartbeat on primary site. Fault manager status shows <code>DB host status as UNUSABLE</code> for the primary host.
Heartbeat from companion is missed for a preconfigured timeout.	Fault manager keeps trying to restart heartbeat on companion site. Fault manager status shows <code>DB host status as UNUSABLE</code> for the companion host.
Companion SAP ASE is down.	Notify cockpit, attempt to restart companion SAP ASE if <code>ha/syb/allow_restart_companion=1</code> is included in the <code>profile</code> file.
Companion host is down.	Replication is turned off. Notify cockpit. May require manual intervention to restart the host and other components.
Companion SAP ASE is unreachable.	Notify cockpit and attempt to restart companion SAP ASE if the corresponding Fault Manager parameter is set.
Replication Server or the RMA on the primary host are restored.	Notify cockpit.

3.6.1 Configuring the Fault Manager from the Command Line

Although you should use the SAP installer to install and configure it, you can use the `sybdbfm` to configure and run the Fault Manager. `sybdbfm` is located in `$SYBASE/FaultManager/bin`.

Syntax

```
sybdbfm [<options>]
hadm pf=<SYBHA.PFL> : start ha process.
hahb pf=<SYBHA.PFL> : start heartbeat.
install [pf=<SYBHA_INST.PFL>] : install.
uninstall [pf=<SYBHA.PFL>] : uninstall.
check pf=<SYBHA_INST.PFL> : check.
```

```
help : display help text.
hibernate : hibernate fault manager.
resume : resume hibernating fault manager.
restart : read in changed profile and restart.
stop : stop fault manager.
status : display status of fault manager.
version : display version.
```

Parameters

Where <options> is one of:

- `hadm pf=<full_path_to/SYBHA.pfl>` – Starts the Fault Manager process. Requires a profile file.
- `hahb pf=<full_path_to/SYBHA.pfl>` – Starts the heartbeat process. Requires a profile file. The Fault Manager (running in monitoring mode) internally starts the process in heartbeat mode on the hosts running SAP ASE.
- `install [pf=<full_path_to/SYBHA.pfl>]` – Configures the Fault Manager. You can use this parameter:
 - In an interactive mode, during which you provide information to command line prompts to install the Fault Manager.
 - With a profile file, which is provided by the `pf=<full_path_to/SYBHA.pfl>` parameter. The profile contains the details you would provide at interactive command prompt.
- `uninstall [pf=<full_path_to/SYBHA.pfl>]` – Uninstalls the Fault Manager. You can use this parameter:
 - In an interactive mode, during which you provide information to command line prompts to install the Fault Manager.
 - With a profile file, which is provided by the `pf=<full_path_to/SYBHA.pfl>` parameter. The profile contains the details you would provide at interactive command prompt.
- `check pf=<full_path_to/SYBHA.pfl>` – Performs the basic bootstrapping of the Fault Manager to confirm if the details provided in the profile file and the user credentials in the SecureStore file are correct and that the Fault Manager can run.
- `hibernate` – Hibernates the Fault Manager. Useful for pausing the Fault Manager for planned maintenance activities on the HADR components (for example, planned failover, upgrades, and so on).
- `resume` – Resumes running the Fault Manager from the hibernating state.
- `restart` – Restarts the Fault Manager using an altered profile file. Used when any profile parameter is modified or added.
- `stop` – Stops the Fault Manager and all heartbeat processes spawned on the HADR system nodes. Execute `stop` from the same directory from which you started the Fault Manager.
- `status` – Displays the status of the Fault Manager.
- `version` – Displays version information along with other build and operating system version information supported by the Fault Manager.

3.6.1.1 SecureStore

The Fault Manager requires username and password combinations to connect to SAP ASE, RMA, SAP ASE Cockpit, and SAP Host Agent. These usernames and passwords are stored in an encrypted format in the SecureStore.

During configuration, the Fault Manager adds usernames and passwords for the following users in the SecureStore:

- SADB_USER – SAP ASE user with the sa_role and replication_role roles.
- DR_USER – RMA user, used for connecting to RMA.
- SAPADM_USER – Operating system user, mostly used for sapadm for SAP HostAgent.
- SCC_USER – Administration user, mostly used for sccadmin.

Use the `rsecssfx` utility to perform this administration duty for SecureStore. Update any changed usernames and passwords in SecureStore. To do so, stop the Fault Manager, update the SecureStore using the `srecssfs` utility, and restart the Fault Manager. Stop the Fault Manager while the password is changed in the cluster components.

i Note

Source the `$SYBASE/SYBASE.csh` (`SYBASE.sh` for the Korn shell) file to configure the environment variables.

- Use the `put` parameter to add or update entries in the SecureStore. The syntax is:

```
$SYBASE/FaultManager/bin/rsecssfx put DB_CONNECT/SYB/<parameter_name> [-plain]
```

Where `<parameter_name>` is one of:

- SADB_USER – SAP ASE user with the sa_role and replication_role roles.
 - SADB_PASSWORD – Password for SADB_USER.
 - DR_USER – RMA user, used for connecting to RMA.
 - DR_PASSWORD – Password for DR_USER.
 - SAPADM_USER – Operating system user, mostly used for sapadm for SAPHostAgent.
 - SAPADM_PASSWORD – Password for SAPADM_USER.
 - SCC_USER – Administration user, mostly used for sccadmin.
 - SCC_PASSWORD – Password for SCC_USER.
- `-plain` – indicates you are storing the details in plain (human-readable) text. If you do not include `-plain`, `rsecssfx` encrypts and stores the details.

For example, this adds the SADB_USER user in plain text and the SADB_PASSWORD password in encrypted form:

```
<installation_directory>/FaultManager/bin/rsecssfx put DB_CONNECT/SYB/  
SADB_USER sa -plain  
<installation_directory>/FaultManager/bin/rsecssfx put DB_CONNECT/SYB/  
SADB_PASSWORD Sybase123
```

- Use the `list` parameter to list entries in the SecureStore. For example:

```
./FaultManager/bin/rsecssfx list  
-----  
| Record Key | Status | Time Stamp of Last Update |  
-----
```

```

| DB_CONNECT/SYB/DR_PASSWORD | Encrypted | 2015-06-09 08:48:29 UTC |
|-----|-----|-----|
| DB_CONNECT/SYB/DR_USER | Plaintext | 2015-06-09 08:48:25 UTC |
|-----|-----|-----|
| DB_CONNECT/SYB/SADB_PASSWORD | Encrypted | 2015-06-09 08:48:22 UTC |
|-----|-----|-----|
| DB_CONNECT/SYB/SADB_USER | Plaintext | 2015-06-09 08:48:18 UTC |
|-----|-----|-----|
| DB_CONNECT/SYB/SAPADM_PASSWORD | Encrypted | 2015-06-09 08:48:35 UTC |
|-----|-----|-----|
| DB_CONNECT/SYB/SAPADM_USER | Plaintext | 2015-06-09 08:48:31 UTC |
|-----|-----|-----|
| DB_CONNECT/SYB/SCC_PASSWORD | Encrypted | 2015-06-09 08:48:42 UTC |
|-----|-----|-----|
| DB_CONNECT/SYB/SCC_USER | Plaintext | 2015-06-09 08:48:38 UTC |
|-----|-----|-----|
Summary ----- Active Records : 8 (Encrypted: 4, Plain: 4, Wrong Key: 0,
Error: 0)
Defunct Records : 0 (180+ days: 0; Show: "list -withHistory", Remove:
"compact")

```

i Note

On Windows, include `<installation_directory>/FaultManager/bin` in the `<PATH>` environment variable.

3.6.2 Administering the Fault Manager

Start the Fault Manager from the installation directory, `installation_directory/FaultManager/sybdbfm_cluster_ID`.

Use the `sybdbfm` utility to view the status of the Fault Manager. For example:

```

$ sybdbfm status
fault manager running, pid = 17763, fault manager overall status = OK, currently
executing in mode PAUSING
*** sanity check report (1)***.
node 1: server star1, site hasite0.
db host status: OK.
db status OK hadr status PRIMARY.
node 2: server star2, site hasite1.
db host status: OK.
db status OK hadr status STANDBY.
replication status: SYNC_OK.

```

Edit the Fault Manager profile file to change any parameter. The profile file is named `SYBHA.PFL`, and is located in the install directory of the Fault Manager on all platforms. Restart the Fault Manager for the profile parameter changes to take effect.

You should continuously monitor the Fault Manager log (named `dev_sybdbfm`, and located in `<Fault_Manager_install_directory>/FaultManager`).

i Note

If a problem related to Fault Manager or the heartbeat requires you to consult SAP, back up the following data when the problem occurs:

- Fault Manager data available on the host running Fault Manager (`$SYBASE` below is the directory where Fault Manager is installed):
 - Fault Manager profile file – located in `$SYBASE/FaultManager/SYBHA.pfl`
 - Fault Manger log file – located in `$SYBASE/FaultManager/dev_sybdbfm`
- Host agent log from the primary and standby sites (requires sudo access):
 - Host Agent log file – located in `/usr/sap/hostctrl/work/dev_sapdbctrl`
 - Heart beat log file – located in `/usr/sap/hostctrl/work/dev_sybdbfm`

3.6.3 Uninstalling the Fault Manager

How you uninstall the Fault Manager depends on whether you installed it using the SAP installer or the `sybdbfm` utility.

1. If you used the `sybdbfm` utility to install the Fault Manager:
 1. Stop the Fault Manager. Move to the directory that contains `SYBHA.PFL` (the default is `<installed_directory>/FaultManager`), and issue:

```
sybdbfm stop
```

2. Remove SecureStore-related files by issuing this from the directory that contains `SYBHA.PFL`:

```
sybdbfm uninstall pf=<path_to_SYBHA.PFL>
```

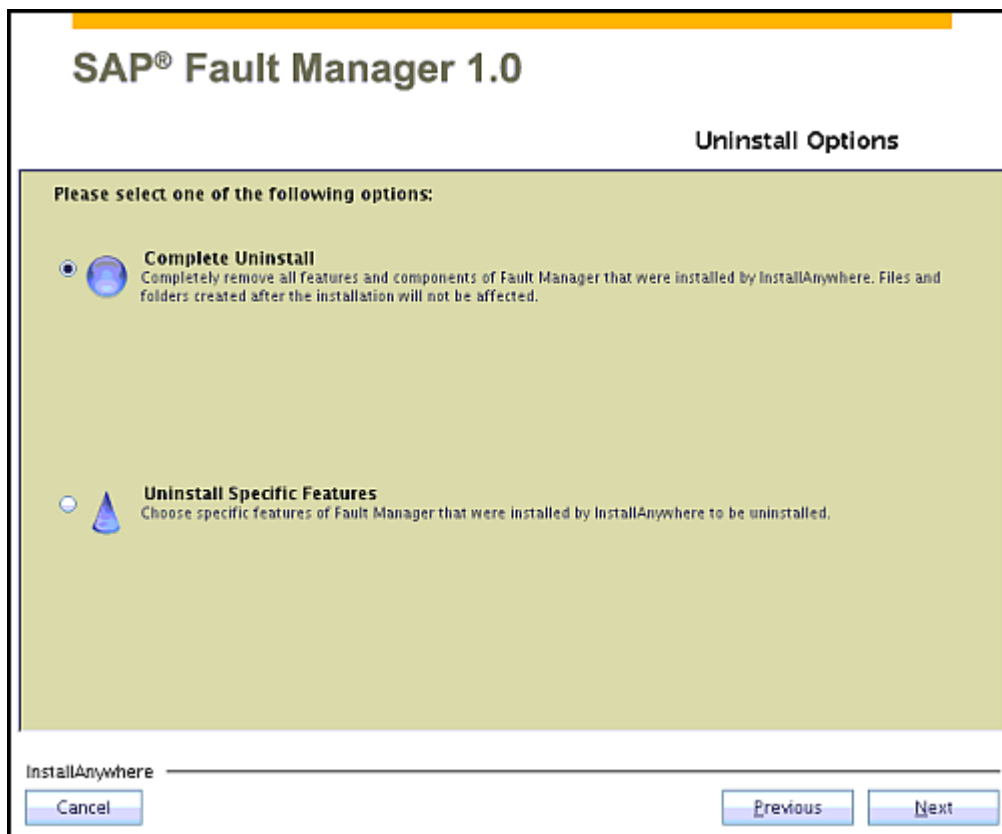
2. Move to `$SYBASE/sybuninstall/FaultManager` and issue:

```
./uninstall
```

When the `uninstall` utility displays the Uninstall Fault Manager introductory screen, click [Next](#):




3. In the Uninstall Options screen, select the appropriate option.

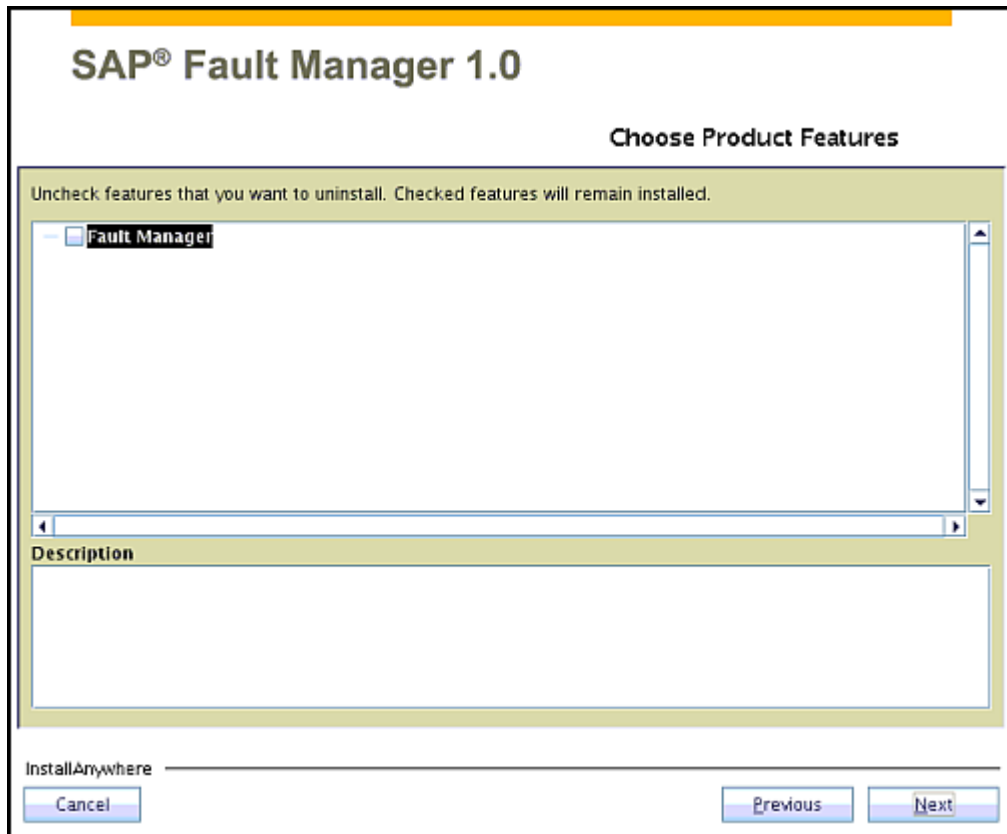


If you select:


- *Complete Uninstall* – the `uninstall` utility displays a summary of all features and components it will remove.



- *Uninstall Specific Features* – select which available features you want to uninstall from the list uninstall utility displays.
 1. Uncheck the features you want to remove and click *Next*.



2. The `uninstall` utility displays the *Summary* screen.



Click [Previous](#) to return to the previous screen. When you are ready to proceed, click [Next](#).

4. Delete the following files and directories manually (requires sudo permission) as part of uninstalling Fault Manager from the primary and companion hosts:
 1. Delete these directories:
 - `/usr/sap/hostctrl/exe/<SAP_ASE_server_name>`
 - `/usr/sap/hostctrl/work/<SAP_ASE_server_name>`
 - (If present) `/usr/sap/hostctrl/work/<SAP_ASE_server_name>_REP`
 2. Delete these `sybdbfm` files:
 - `/usr/sap/hostctrl/exe/sybdbfm`
 - `/usr/sap/hostctrl/exe/sybdbfm_<SAP_ASE_server_name>`
 - `/usr/sap/hostctrl/exe/SYBHA_<SAP_ASE_server_name>.PFL`
 - `/usr/sap/hostctrl/work/dev_h2h`
 - (If present) `/usr/sap/hostctrl/work/dev_h2h.old`
 - `/usr/sap/hostctrl/work/dev_hbeat`
 - (If present) `/usr/sap/hostctrl/work/dev_hbeat.old`
 - `/usr/sap/hostctrl/work/dev_sybdbfm`

3.7 Performing a Rolling Upgrade

The HADR feature allows SAP ASE applications to operate with zero down time while you are updating the SAP ASE software.

Generally, the steps for performing a rolling upgrade are:

- Upgrade the standby components
- Perform a failover
- Upgrade the components that are on standby after the failover

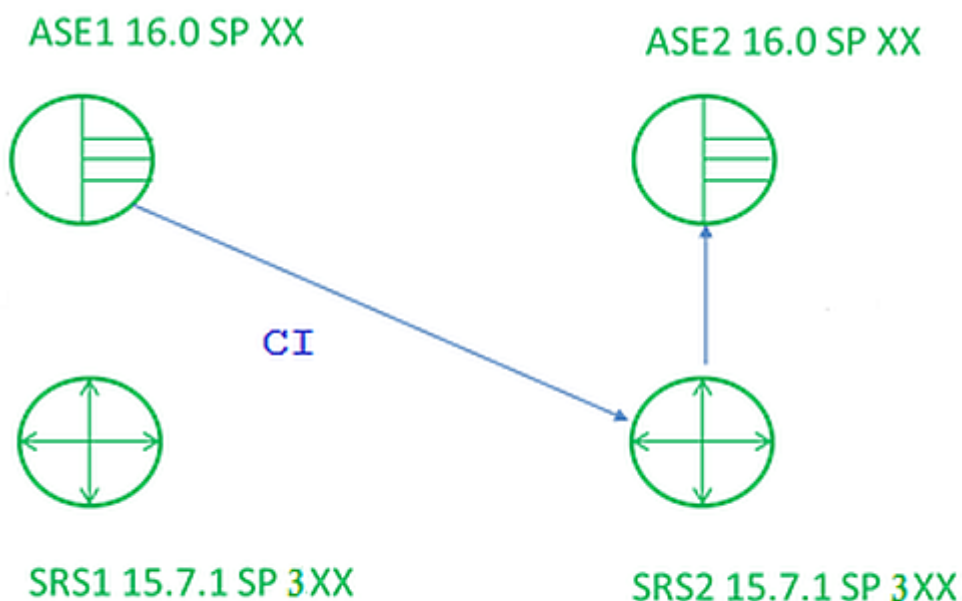
Complete the upgrade steps in a single sequence: partial upgrade is not supported (for example, you cannot upgrade some components now and then upgrade the other components at another time). Replication is suspended during some steps of a rolling upgrade, and if you perform a partial upgrade, logs continue to grow, which can result in logs or the SPQ running out of space. During a rolling upgrade, the versions between SAP ASE and Replication Server need not match.

The `RUN_rs instance name.sh` Replication Server runserver file is regenerated during an upgrade, and any user changes to this file are lost. If your site requires these changes, edit the runserver file after the upgrade is complete then restart Replication Server to make the environment settings take effect.

i Note

Before upgrading HADR with SAP Business Suite on SAP ASE, you may have to follow instructions from your application vendors over the general guidelines in this chapter. See SAP note [2808173](#) for more details.

The steps below describe performing a rolling upgrade on an HADR system with this remote topology:




In this topology, the primary server (ASE1) is installed on the same host as the inactive Replication Server (SRS1). The active Replication Server (SRS2) is installed on a remote host, along with the standby server (ASE2). Data changes that occur in ASE1 are sent by the Replication Agent thread to the active SRS2 running on the remote host. The active SRS2 then routes these changes to ASE2, which is running on the same host as the active Replication Server, SRS2. In this setup, the inactive Replication Server, SRS1, is not involved in data movement until failover occurs. The communication among ASE1, SRS1, and ASE2 is through a client interface (stream replication, indicated in this topology as "CI").

Run this command to determine which SAP ASE you are connected to in the HADR system:

```
select asehostname ()
```

The initial topology configuration looks like:



In this configuration, all components are running, and the standby server is almost in sync with the primary server. Prior to upgrade, site1 is the primary server and site2 is the companion server (in high-availability – HA – mode, the companion server is referred to as the standby server) with remote replication topology. The Replication Server versions prior the upgrade are compatible with the Replication Server versions after the upgrade. If you upgrade from a "1-Off" release, you can upgrade only the SAP ASE or Replication Servers.


To perform a rolling upgrade:

i Note

Stop the Fault Manager before you perform a rolling upgrade (even if you are performing planned activities like a planned failover). You can start the Fault Manager after the upgrade is complete. To stop the Fault Manager, issue this from the `<installation_directory>/FaultManager` directory:

```
<Fault_Manager_install_dir>/FaultManager/bin/sybdbfm stop
```

To perform a rolling upgrade, you first upgrade SRS1 on site1 to a higher version:



1. Verify that Replication Server is not running any `isql` processes during this Replication Server installation step. If there are `isql` processes running, Replication Server issues an error message stating "isql text file busy".
2. Log into RMA on site1 as the DR_admin user and issue `sap_upgrade_server` to prepare the Replication Server on site1 for upgrade:

```
sap_upgrade_server SRS, start, <site1_site_name>
```

3. Shut down RMA:

```
shutdown
```

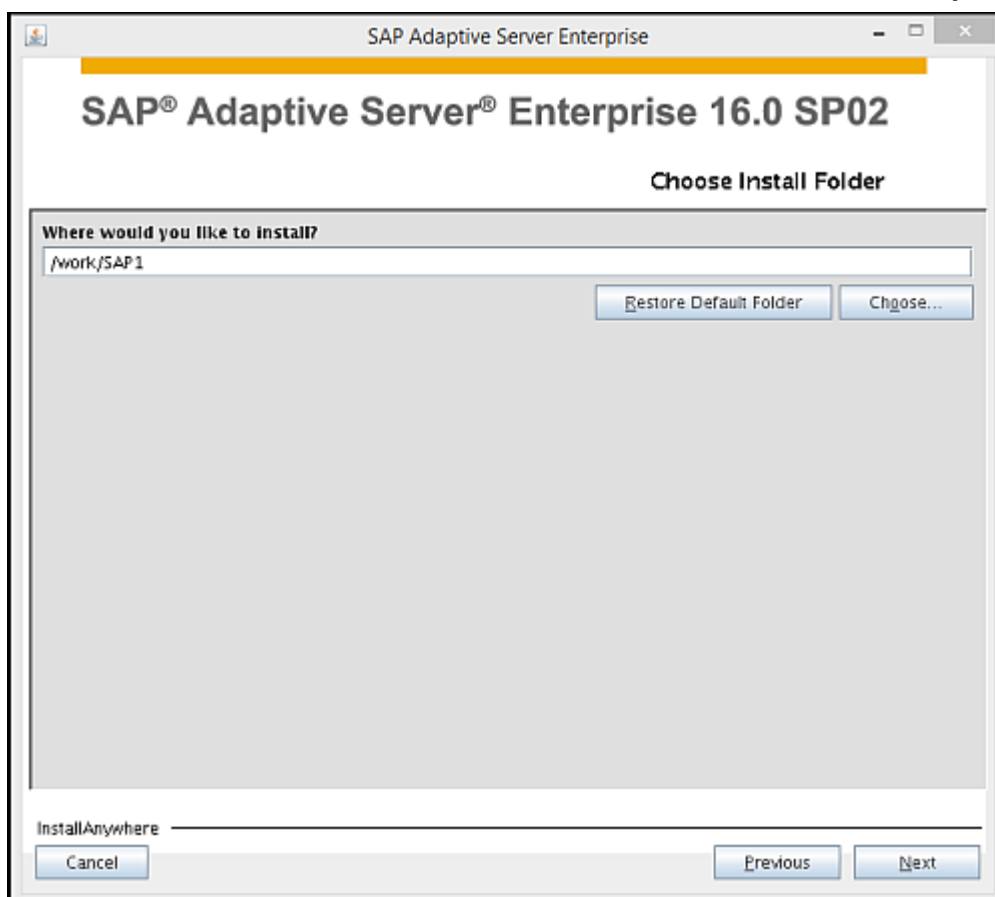
4. Remove the RMA service: On Windows, execute the following command from either the `%SYBASE%\RMA-16_0\compatibility\WinService\Win32\Release` directory, or the `%SYBASE%\RMA-16_0\compatibility\WinService\x64\Release` directory, to remove the RMA service –

```
drservice -remove <Cluster_ID>
```

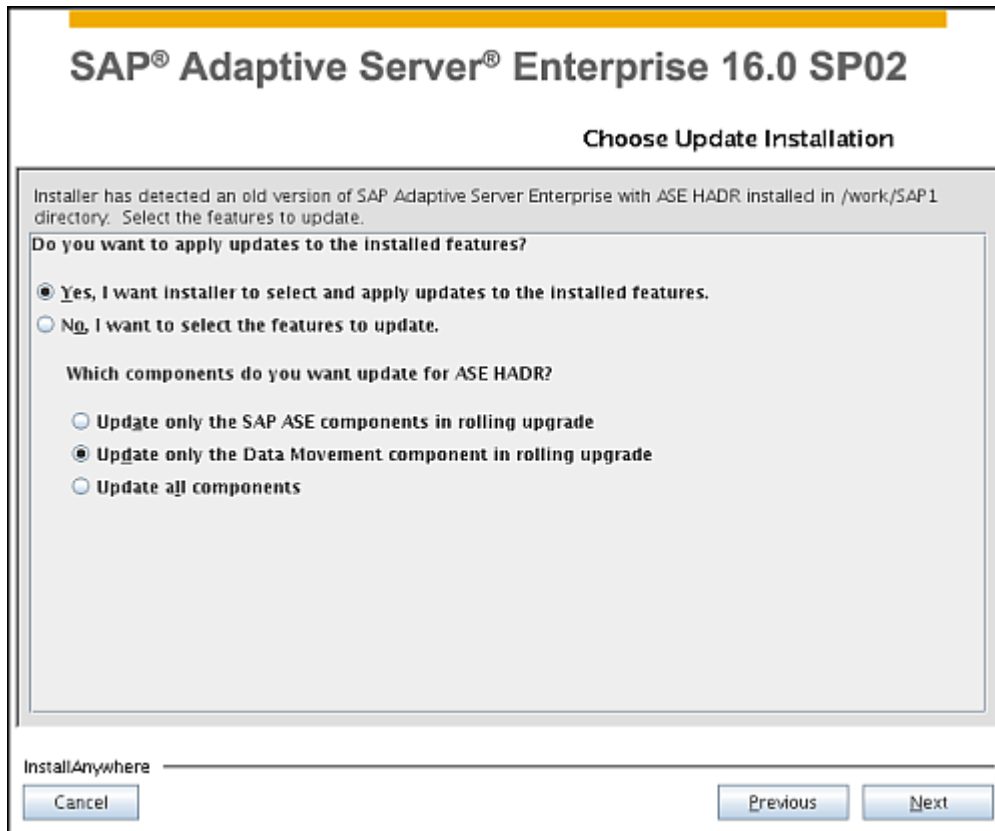
5. Start the SAP installer from the installation directory:

```
install_directory/setup.bin
```

6. In the Choose Install Folder screen, enter the current SAP ASE SAP installation directory, then click *Next*:



7. In the Choose Update Installation screen, determine if you want the installer to select and apply updates, then select *Update only the Data Movement component in rolling upgrade*.



i Note

The SAP installer must complete the software update before you continue to the next step.

8. Install a new RMA service on Windows: To install, and then start the new RMA service on Windows, execute the following command from either the %SYBASE%\RMA-16_0\compatibility\WinService\Win32\Release directory, or the %SYBASE%\RMA-16_0\compatibility\WinService\x64\Release directory –

```
drservice -install <Cluster_ID>
```

9. After the installer finishes the upgrade, start RMA:
 - (UNIX) – issue \$SYBASE/\$SYBASE_ASE/bin/rma
 - (Windows) – start the RMA Windows service by either of the following:
 - Starting *Sybase DR Agent - <cluster_ID>* from the Services panel
 - Issuing this command, where <cluster_ID> is the ID of the cluster:

```
net start SybaseDRAgent_<cluster_ID>
```

10. Log in to RMA on site1 as the DR_admin user and issue sap_upgrade_server to finish the upgrade for Replication Server on site1:

```
sap_upgrade_server SRS, finish,<site1_site_name>
```


i Note

At this point of the upgrade process, the HADR system is working normally with ASE1, SRS2, ASE2 at the older versions, and SRS1 at newer version.

11. Log into RMA on site1 as the DR_admin user and issue:

```
sap_failover <site1_site_name>, <site2_site_name>, 30
```

This command allows a 30-second grace period for any running transactions to complete before the deactivation starts. Failover will not succeed if there are still active transactions after 30 seconds. If this occurs, retry the command when the system is not busy, use a longer grace period, or use the `force` option to terminate the client connection (if it is safe) with:

```
sap_failover <site1_site_name>, <site2_site_name>, 30, force
```

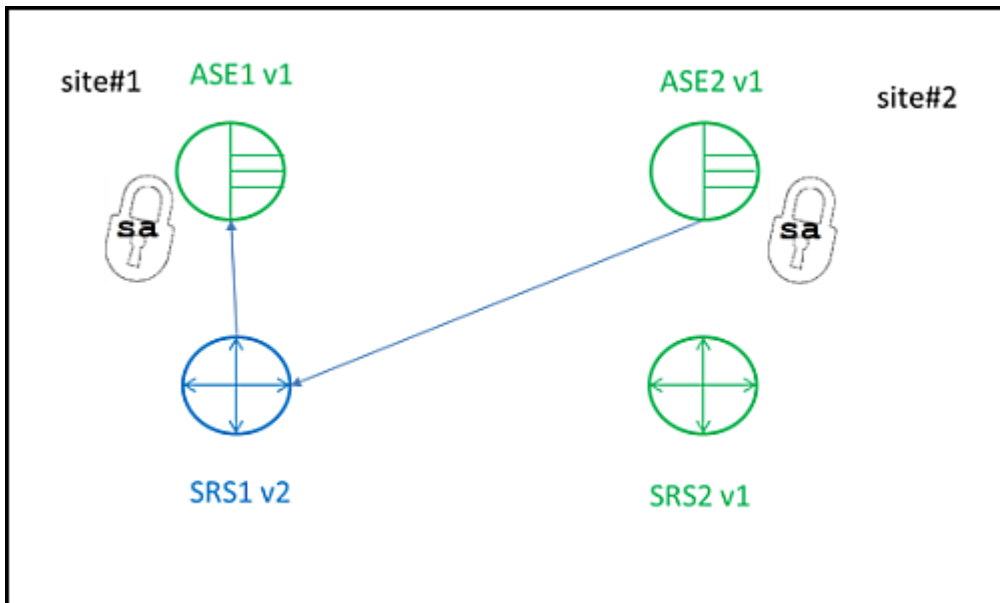
12. The `sap_failover` command may take a long time to finish. To check the status of the `sap_failover` command, issue this from the RMA:

```
sap_status task
```

13. Once the `sap_status` command returns `Completed`, resume replication by issuing this from the RMA:

```
sap_host_available <site1_site_name>
```

At this point, the topology looks like:



14. Verify that Replication Server is not running any `isql` processes during the Replication Server installation step below. If there are `isql` processes running, Replication Server issues an error message stating "`isql text file busy`".
15. Login to RMA on site2 as the DR_admin user and issue `sap_upgrade_server` to start the upgrade for Replication Server on site2:

```
sap_upgrade_server SRS, start,<site2_site_name>
```

16. Shut down RMA:

```
shutdown
```

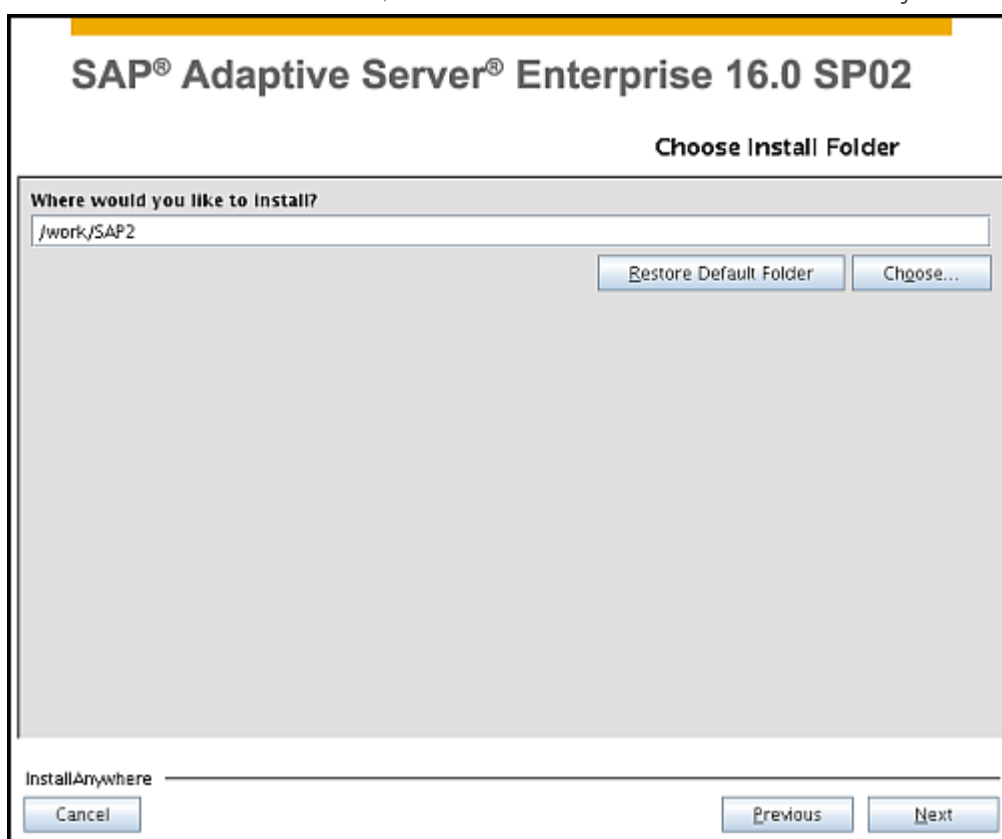
17. Remove the RMA service: On Windows, execute the following command from either the %SYBASE%\RMA-16_0\compatibility\WinService\Win32\Release directory, or the %SYBASE%\RMA-16_0\compatibility\WinService\x64\Release directory, to remove the RMA service –

```
drservice -remove <Cluster_ID>
```


18. Start the SAP installer from the installation directory:

```
<install_directory>/setup.bin
```

19. In the Choose Install Folder screen, enter the current ASE SAP installation directory.



20. In the Choose Update Installation screen, determine if you want the installer to select and apply updates, then select *Update only the Data Movement component in rolling upgrade*.



i Note

The SAP installer must complete the software update before you continue to the next step.

21. Install a new RMA service on Windows: To install, and then start the new RMA service on Windows, execute the following command from either the %SYBASE%\RMA-16_0\compatibility\WinService\Win32\Release directory, or the %SYBASE%\RMA-16_0\compatibility\WinService\x64\Release directory –

```
drservice -install <Cluster_ID>
```

22. After the SAP installer has finished the upgrade, start RMA:
 - (UNIX) – \$SYBASE/\$SYBASE_ASE/bin/rma
 - (Windows) – start the RMA Windows service by either of the following:
 - Starting *Sybase DR Agent* - <cluster_ID> from the Services panel
 - Issuing this command, where <cluster_ID> is the ID of the cluster:

```
net start SybaseDRAgent_<cluster_ID>
```

23. Log into RMA on site2 as the DR_admin user and issue sap_upgrade_server to finish the upgrade for Replication Server on site2:

```
sap_upgrade_server SRS, finish, <site2_site_name>
```


Note

At this point of the upgrade process, the HADR system is working normally with ASE1 and ASE2 at the older versions, and SRS1 and SRS2 at the newer version.

- (Skip this step if you do not lock the sa user) If the sa user is locked, temporarily unlock this user on ASE1 during the upgrade process by logging in as the user with SSO permission on ASE1 and issuing:

```
sp_locklogin sa, "unlock"
```


At this point, the topology looks like:



- Log into RMA on site1 as the DR_admin user and issue `sap_upgrade_server` to start the upgrade for SAP ASE on site1:

```
sap_upgrade_server ASE, start, <site1_site_name>
```

The topology now looks like:



26. Shut down SAP ASE and Backup Server on site1. Login to site1 SAP ASE server and issue:

```
shutdown SYB_BACKUP
go
shutdown
go
```

27. Shut down SAP ASE Cockpit. If the SAP ASE Cockpit is running:

- o In the foreground – At the cockpit> prompt, execute:

```
shutdown
```

Enter a reason for the shutdown.

i Note
Do not enter shutdown at a UNIX prompt: Doing so shuts down the operating system.

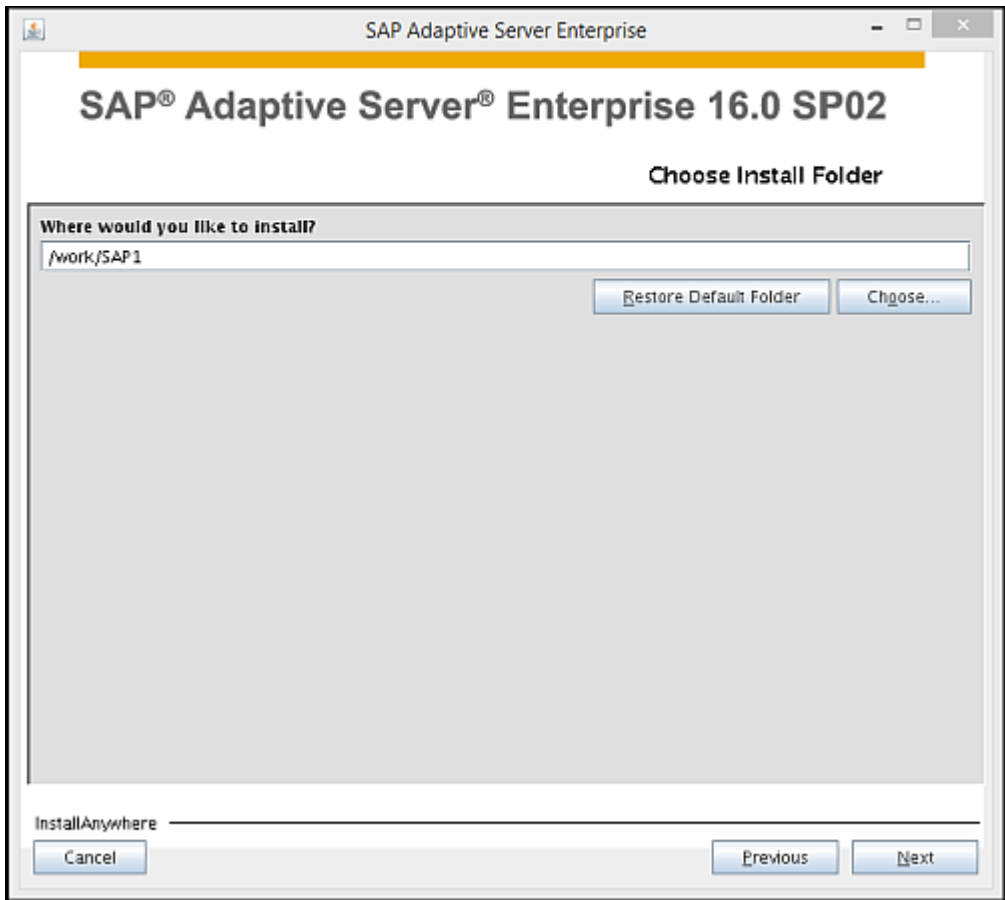
- o In the background – At the UNIX command line, execute:

```
$SYBASE/COCKPIT-4/bin/cockpit.sh --stop
```


28. Start the SAP installer from the installation directory:

```
<install_directory>/setup.bin
```

29. In the Choose Install Folder screen, enter the current ASE SAP installation directory.



30. In the Choose Install Folder screen, determine if you want the installer to select and apply updates, then select *Update only the SAP ASE component in rolling upgrade*.



i Note

The SAP installer must complete the software update before you continue to the next step.

31. After the SAP installer has completed the upgrade, use the `updatease` utility to upgrade SAP ASE, which runs `installmaster` and performs other tasks to bring SAP ASE up to date. `updatease` is available as a GUI display or a command line tool.

i Note

You need not perform this step if you update the SAP ASE server instance in the SAP installer.

1. Start `updatease`. From the `$$SYBASE/$SYBASE_ASE/bin/` directory, issue:

```
./updatease
```

2. Enter the server name and password.

The command line version of `updatease` displays messages similar to:

```
./updatease
Server: SFSAP1
ASE Password:
Updating SAP Adaptive Server Enterprise 'SFSAP1'...
Running installmaster script...
installmaster: 10% complete.
installmaster: 20% complete.
installmaster: 30% complete.
```

```

installmaster: 40% complete.
. . .
instmsgs.ebf: 80% complete.
instmsgs.ebf: 90% complete.
instmsgs.ebf: 100% complete.
The instmsgs.ebf script has been successfully installed.
Restarting SAP Adaptive Server Enterprise...
SAP Adaptive Server Enterprise has been restarted.
Updating SAP Adaptive Server Enterprise complete successfully.

```

32. Start Backup Server from the command line by issuing:

```
$SYBASE/ASE-16_0/install/RUN_<backupserver_name>
```

33. Start the SAP ASE Cockpit.

- In the foreground – from the command line, execute:

```
$SYBASE/COCKPIT-4/bin/cockpit.sh
```

- In the background:


From the Bourne shell (sh) or Bash, issue:

```
nohup ./cockpit.sh 2>&1 > cockpit-console.out &
```

From the C shell, issue:

```
nohup ./cockpit.sh >& cockpit-console.out &
```

The topology now looks like:



34. Log into RMA on site1 as the DR_admin user and issue `sap_upgrade_server` to complete the upgrade for SAP ASE on site1:


```
sap_upgrade_server ASE, finish,<site1_site_name>
```

35. (Skip this step if you do not lock the sa login) Log in to ASE1 as a user with `sapsso` permission and issue this once the upgrade is complete:

```
sp_locklogin sa, "lock"
```

i Note

At this point of the upgrade process, the HADR system is working normally, with ASE2 at the older versions and SRS1, SRS2, and ASE1 at the newer versions. The topology looks like:



36. Log into RMA on site2 as the DR_admin user and issue:

```
sap_failover <site2_site_name>, <site1_site_name>, 30
```

This command allows a 30-second grace period for any running transactions to complete before the deactivation starts. Failover will not succeed if there are still active transactions after 30 seconds. If this occurs, retry the command when the system is not busy, use a longer grace period, or use the `force` option to terminate the client connection (if it is safe) with:

```
sap_failover <site2_site_name>, <site1_site_name> 30, force
```


37. The `sap_failover` command may take a long time to finish. To check the status of the `sap_failover` command, issue this from the RMA:

```
sap_status task
```

38. Once the `sap_status` command returns `Completed`, resume replication from ASE1 to SRS2, but not to ASE2, by issuing this from the RMA (`suspend` ensures that no ticket is sent to verify the replication status during the upgrade process):

```
sap_host_available <site2_site_name>, suspend
```


The topology looks like:




The replication path from ASE1 to SRS2 is restored and all commits on ASE1 are synchronously committed on SRS2, meaning there is zero data loss if site1 is lost during this time. Because replication to ASE2 is not restored, the log records in SRS2 are not applied to ASE2. Make sure you have sufficient storage configured in the SPQ for the SRS2 for the short period of time until the upgrade process completes.

39. (Skip this step if you do not lock the sa login) Temporarily unlock the sa user on ASE2 for upgrade. Login as the user with sapsso permission on ASE2 and issue:

```
sp_locklogin sa, "unlock"
```

The topology looks like:



40. Log into RMA on site2 as the DR_admin user and issue `sap_upgrade_server` to start the upgrade for SAP ASE on site2 (`suspend` ensures that no ticket is sent to verify the replication status during the upgrade process):

```
sap_upgrade_server ASE, start,<site2_site_name>, suspend
```

41. Shut down SAP ASE and Backup Server on site2. Login to site2 SAP ASE server and issue:

```
shutdown SYB_BACKUP
go
shutdown
go
```

42. Shutdown SAP ASE Cockpit. If the SAP ASE Cockpit is running:

- o In the foreground – At the `cockpit>` prompt, execute:

```
shutdown
```

Enter a text reason for the shutdown.

i Note

Do not enter `shutdown` at a UNIX prompt; it shuts down the operating system.

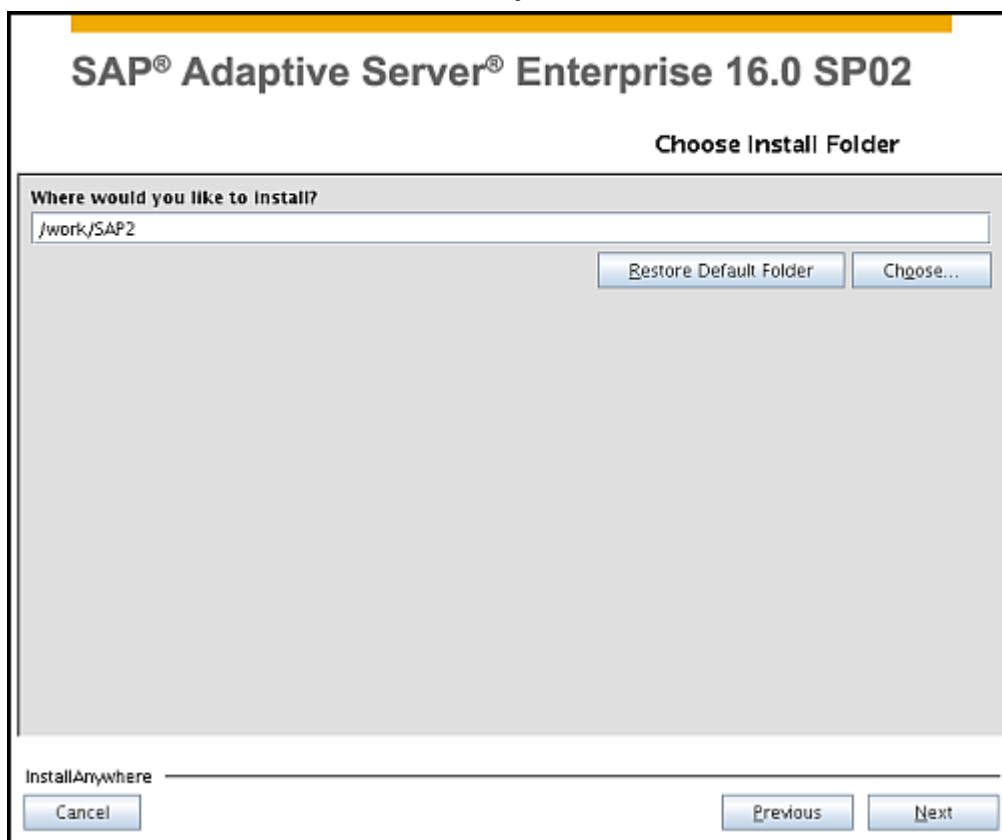
- o Running in the background – At the UNIX command line, execute:

```
$$SYBASE/COCKPIT-4/bin/cockpit.sh --stop
```


43. Start the SAP installer from the installation directory:

```
install_directory/setup.bin
```

44. Enter the current ASE SAP installation directory in the Choose Install Folder screen.



45. In the Choose Update Installation screen, determine if you want the installer to select and apply updates, then select *Update only the SAP ASE component in rolling upgrade*:



Note

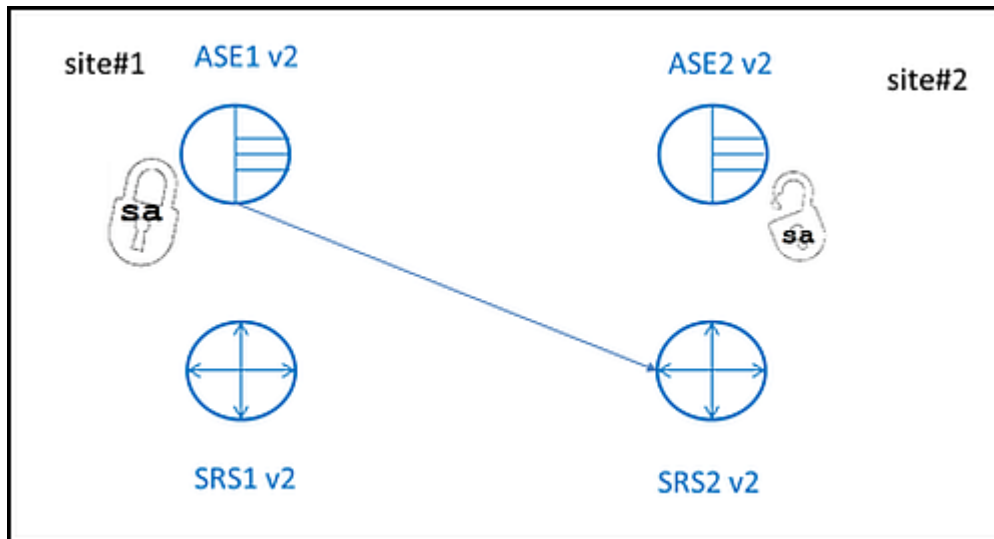
The SAP installer must complete the software update before you continue to the next step.

46. After the SAP installer has completed the upgrade, use the `updatease` utility to upgrade SAP ASE, which runs `installmaster` and performs other tasks to bring SAP ASE up to date. See the `updatease` instructions in the earlier step.

47. Start Backup Server from the command line by issuing:

```
$SYBASE/ASE-16_0/install/RUN_<backupserver_name>
```

The topology looks like:



48. Start the SAP ASE Cockpit.

- In the foreground – from the command line, execute:

```
$SYBASE/COCKPIT-4/bin/cockpit.sh
```

- In the background:

From the Bourne shell (sh) or Bash, issue:

```
nohup ./cockpit.sh 2>&1 > cockpit-console.out &
```

From the C shell, issue:

```
nohup ./cockpit.sh >& cockpit-console.out &
```

49. Log into RMA on site2 as the DR_admin user and issue `sap_upgrade_server` to complete the upgrade for SAP ASE on site2:


```
sap_upgrade_server ASE, finish,<site2_site_name>, suspend
```

50. (Skip this step if you do not lock the sa login) Log in to ASE2 as the user with `sapsso` permissions and issue this once the upgrade is complete:

```
sp_locklogin sa, "lock"
```

iNote

At this point of the upgrade process, the HADR system is working normally, with SRS1, SRS2, ASE1, and ASE2 at the newer versions. The topology looks like:



51. Run Fault Manager installer:

```
<installation_directory>/FaultManager/setup.bin
```

52. Choose the existing ASE-installed directory on the *Choose Install Folder* screen. Do not choose to configure Fault Manager in the installer.

53. Set the environment variables, and source SYBASE.sh for the Bourne shell:

```
source <installation_directory>/SYBASE.csh
```

54. Start the Fault Manager:

```
<Fault_Manager_install_dir>/FaultManager/sybdbfm_<CID>
```

3.8 Upgrading SAP ASE Version 15.7 DR to Version 16.0 HADR

These sections describe how to upgrade an SAP ASE version 15.7 disaster recovery (DR) solution to ASE 16.0 HADR solution in a Business Suite or a Custom Application environment.

Before you begin, it is important to consider:

- These steps require an application downtime, the length of which depends on the size of the databases you are upgrading.
- If you do not follow the upgrade steps precisely as documented and in the proper sequence, you will lose data.

3.8.1 Upgrade the SAP Host Agent

The SAPHOSTAGENT<SP-version>.SAR archive contains all of the required elements for centrally monitoring any host. It is available for all operating system platforms supported by SAP

Context

The SAP Host Agent is automatically installed during the installation of SAP systems or instances with SAP kernel 7.20 or higher.

Procedure

1. Go to the *Software Downloads* of the *SAP Support Portal* at <http://support.sap.com/swdc>.
2. Log on with your *SAP Support Portal ID*.
3. Choose ► *Installation and Upgrades* ► *By Alphabetical Index (A-Z)* ► *H* ► *SAP Host Agent* ► *SAP Host Agent 7.21* ► Select highest available version.
4. Select the appropriate SAPHOSTAGENT<SP-version>.SAR archive from the *Download* tab.

→ Recommendation

Always select the highest SP version of the SAPHOSTAGENT<SP-version>.SAR archive, even if you want to monitor a component of SAP NetWeaver with a lower release.

5. Make sure that the SAPCAR tool is available on the host where you want to install SAP Host Agent.

Decompressing the SAPHOSTAGENT<SP-version>.SAR archive requires the SAPCAR tool. To download the SAPCAR tool:

- a. Go to the *Software Downloads* of the *SAP Support Portal* at <http://support.sap.com/swdc>.
- b. Choose ► *Installation and Upgrades* ►
- c. Select the highest available version.

6. Log on as a user with the required authorization:

Windows	As a member of the local <code>Administrators</code> group
----------------	--

UNIX	As a user with <code>root</code> authorization
-------------	--

7. Copy the downloaded SAPHOSTAGENT<SP-version>.SAR archive to a temporary directory, for example:

Windows	<code>c:\temp\hostagent</code>
----------------	--------------------------------

UNIX	<code>/tmp/hostagent</code>
-------------	-----------------------------

8. Extract the SAPHOSTAGENT<SP-version>.SAR archive using SAPCAR.

Take the information provided by SAP Note [212876](#) into account when doing so. Use the following command for extraction, and execute them in the directory of the archive:

Windows `<path to SAPCAR>\sapcar.exe -xvf <path to temporary directory>\SAPHOSTAGENT<SP-version>.SAR -manifest SIGNATURE.SMF`

UNIX `/<path to SAPCAR>/sapcar -xvf <path to temporary directory>/SAPHOSTAGENT<SP-version>.SAR -manifest SIGNATURE.SMF`

Among others, the archive contains the `saphostexec` program.

9. Perform the upgrade by running the following command:

Windows `<path to temporary directory>\saphostexec.exe -upgrade`

→ Recommendation

You can use the additional parameter `-verify` to verify the content of the installation package against the SAP digital signature

UNIX `/<path to temporary directory>/saphostexec -upgrade`

→ Recommendation

You can use the additional parameter `-verify` to verify the content of the installation package against the SAP digital signature

The progress of the upgrade is displayed on the command line.

10. After the upgrade has finished successfully, you can check the version of the upgraded host agent by executing the following command from the directory of the SAP Host Agent executables:

Windows `"%ProgramFiles%\SAP\hostctrl\exe\saphostexec.exe" -version`

UNIX

- If you are logged on as a user with `root` authorization, the command is as follows: `/usr/sap/hostctrl/exe/saphostexec -version`
- If you are logged on as a member of the `sapsys` group, for example `<sapsid>adm`, the command is as follows: `/usr/sap/hostctrl/exe/hostexecstart -version`

3.8.2 Upgrading to an HADR System

Procedure

1. Log into SAP Primary Application Server using user `<sid>adm`:
2. Enter the password and respond to the prompts.

3. Stop the Netweaver system. At the prompt, enter:

```
stopsap r3
```

4. Check the ticket table on the standby site. Log in with `isql` and issue:

```
use master
go
select count(*) from rs_ticket_history
go
-----
          3
```

Verify this is the same ticket history in the replicated database (in this example, TIA):

```
use TIA
go
select count(*) from rs_ticket_history
go
-----
          3
```

5. Log in to the RMA, and check the SAP ASE transaction backlog and the Replication Server queue backlog using `sap_send_trace` RMA command (this example runs the command against the `PRI` database):

```
sap_send_trace PRI
TASKNAME          TYPE          VALUE
-----
Send Trace        Start Time    Mon Nov 16 04:42:02 EST 2015
Send Trace        Elapsed Time  00:00:00
Send Trace        Task Name     Send Trace
Send Trace        Task State    Completed
Send Trace        Short Description  Send a trace through the
Replication system using rs_ticket
Send Trace        Long Description  Successfully sent
Send Trace        Host Name     Big_Host
(7 rows affected)
1> sap_status resource
2> go
Name              Type          Value
-----
15:37:55.327     Start Time    2016-08-24
                  Elapsed Time  00:00:00
                  Estimated Failover Time  0
PRI               Replication device size (MB)  15360
PRI               Replication device usage      112
COM               Replication device size (MB)  15360
COM               Replication device usage      128
PRI.master        ASE transaction log (MB)      300
PRI.master        ASE transaction log backlog (MB)  0
PRI.master        Replication queue backlog (MB)  0
COM.master        Replication queue backlog (MB)  0
PRI.TIA           ASE transaction log (MB)      10240
PRI.TIA           ASE transaction log backlog (MB)  0
PRI.TIA           Replication queue backlog (MB)  0
COM.TIA           Replication queue backlog (MB)  0
(15 rows affected)
```


6. Recheck the ticket table on the standby site, making sure the tickets were received. Log in with `isql` and issue:

```
use master
go
select count(*) from rs_ticket_history
go
-----
4
```

Verify this is the same ticket history in the replicated database (in this example, TIA):

```
use TIA
go
select count(*) from rs_ticket_history
go
-----
4
```

7. Uninstall the primary or the standby Replication Server (RMA internally tears down the entire HADR system, including the primary and standby Replication Servers and the users and roles, once you issue `sap_teardown`). Log into either RMA and issue:

```
sap_teardown
```

8. Upgrade SAP ASE to version 16.0. You can upgrade the primary and the standby servers at the same time. Issue this from the command line:

```
saphostctrl -user sapadm - -function LiveDatabaseUpdate -dbname
<database_name> -dbtype syb -dbuser <user_name> -dbpass <password> -
updatemethod Execute -updateoption TASK=UPDATE_ASE -updateoption
DROP_LOCATION="<path_to_package>"
```

For example:

```
/usr/sap/hostctrl/exe/saphostctrl -user sapadm - -function LiveDatabaseUpdate
-database_name TIA -dbtype syb -dbuser sapsso -dbpass sybase123 -updatemethod
Execute -updateoption TASK=UPDATE_ASE -updateoption DROP_LOCATION="/hadr/
packages/SYBASE_LINUX_X86_64"
```

9. Install the Data Movement option on the primary and standby servers using the silent install method.
 - a. Prepare a response file according to the instructions in [Installing the HADR System with Response Files, Console, and Silent Mode \[page 67\]](#), indicating you are installing only the "SAP ASE Data Movement for HADR."
 - b. Log on to the host as user `syb<sid>`.
 - c. Execute the response file according to these instructions [Installing the HADR System in Silent Mode \[page 79\]](#).
10. Unlock the sa user on the primary and standby SAP ASE servers. Perform these steps on both companions:
 - a. Log in to the primary SAP ASE database as user `sapsso`.
 - b. Issue:

```
sp_locklogin sa, 'unlock'
```

11. Unlock user sa on the primary and standby SAP ASE servers. Perform these steps on both companions:
 - a. Log in to the primary SAP ASE database as user `sapsso`.

b. Issue:

```
sp_locklogin sa, 'unlock'
```

12. Configure SAP ASE for the HADR environment. Follow the instructions in [Installing HADR with an Existing System \[page 86\]](#) for configuring the primary and standby servers.

13. Lock user sa on the primary and standby SAP ASE servers. Perform these steps on both companions:

a. Log in to the primary SAP ASE database as user sapsso.

b. Issue:

```
sp_locklogin sa, 'lock'
```

14. Use the `sap_status` command to check the Replication Server status after the upgrade. Issue this at the RMA `isql` prompt:

```
sap_status path
```

You should see this line for the Replication Server Status in the output:

```
Replication Server Status    Active    The status of Replication Server.
```

15. Configure the RMA:

a. Log in to the RMA as user DR_admin

b. Issue this command to configure the Replication Server parameters:

```
sap_tune_rs <PrimarySiteName>, <RAM_Allocated_for_SRS>, <CPU Cores>
```

For example:

```
sap_tune_rs Big_Host, 8, 2
```

c. Execute the same command for the standby site:

```
sap_tune_rs Other_Big_Host, 8, 2
```

16. Start the SAP Application server. Issue:

```
startsap r3
```

4 Installing HADR for Business Suite

When you install HADR in a Business Suite environment, you install SAP ASE, add the Data Movement component to SAP ASE, and run the `setuphadr` utility on both primary and standby hosts.

Before installing or upgrading HADR with SAP Business Suite on SAP ASE, you may have to follow instructions from your application vendors over the general guidelines in this chapter. See SAP note [2808173](#) for more details.

4.1 Requirements and Limitations

- SAP recommends the following sizes for the server resources in an HADR system:

Server Size	Primary transaction log generation rate (GBs per hour)	Replication Server memory, in GBs	Replication Server CPU	Supports this number of large transactions in parallel
Small	7	4	2	1
Medium	15	8	4	2
Large	25	16	8	4
Extra Large	25	24	16	8

- HADR for Business Suite supports the Fault Manager. See SAP Note 1959660.
- The installation environment requires two hosts: a primary and a standby host.
- The installation release directory requires at least 80 GB of free space.
- Refer to the SAP ASE and Replication Server installation guides at help.sap.com for hardware and software prerequisite check.
- An HADR system in the Business Suite environment includes the system ID (SID) database (this database is created as part of the Business Suite installation).

Additional Parameters for `sap_set` in a Business Suite Environment

The RMA `sap_set <global_level_property>` variable includes additional parameters in a Business Suite environment:

- `<sap_sid>` – consists of three alphanumeric characters, and denotes the SAP System ID.
- `<installation_mode>` – specifies the HADR system type. For Business Suite, the mode is `BS`.

See [sap_set \[page 496\]](#) for the full syntax.

4.2 Prepare the Image

Download the following to prepare the image installation.

Software Required	Version	
<code>sapinst</code>	Use the latest available version, or at a minimum, SWPM10SP22_0.SAR.	The SAP Installer used to install SAP Software, including SAP Business Suite
Kernel	The version depends on which supported application version you are installing. See SAP Note 1554717.	The the main component of all SAP Applications. The Kernel contains the executable files for stating various SAP processes.
Exports	There are a number of supported Business Suite applications, including Net-Weaver, ERP, CRM, and so on.	Contains the tables, code, and transactions required for SAP Applications. An export media identifies which SAP application is being installed on the system
SAP ASE for Business Suite for HADR	16.0 SP03 PL08.	Software for the SAP HADR system

To download the image:

1. Download `SAPCAR.exe` from <http://support.sap.com/swdc> (SAP Software Download Center).
2. Use `SAPCAR.exe` to extract the `sapinst` SAR file to your local disk:

```
SAPCAR.exe -xvf SWPM10SP<version>.SAR
```

3. Use a file copy tool (such as Filezilla) to copy SAP ASE, the kernel, and exports images to the primary and companion hosts.

4.3 Installing the HADR System on the Primary

Configuring the HADR on the primary includes installing the Data Management software and running the `setuphadr` utility.

4.3.1 Installing the Business Suite Application

Install the Business Suite application using the application installation process

Context

The installation process varies depending on which installation application you use. The steps to install the Business Suite application depends on which application you use. The following example describes the NetWeaver installation process.

Procedure

1. Move to the `sapinst` directory, which was created when the `SAPCAR.exe` utility extracted files.
2. Execute the `sapinst` utility to start the SAP installation GUI
3. Select **SAP NetWeaver 7.5 > SAP ASE > SAP Systems > SAP Systems > Application Server ABAP > Standard System > Standard System** and click *Next*.
4. Define the parameters. Select either *Typical* or *Custom* parameter mode button (*Typical* does not display all input parameters), and click *Next*.
5. Enter the SAP SID and the destination drive. The SAP SID comprises three alphanumeric characters and must be unique on your system. The SAP SID becomes the name of the destination directory into which the NetWeaver software is loaded.
6. Click *Next*.
7. Enter and confirm the master password and click, then *Next*.
8. Define the kernel path by choosing *Provide the path to installation media (DVD or Blu-ray disc)* and entering the path in the box provided, or by selecting *Browse* to explore the system. The path you include depends on the platform:
 - (Linux) – `<Kernel_PATH>/DATA_UNITS/K_745_U_LINUX_X86_64`
 - (AIX) – `<Kernel_PATH>/DATA_UNITS/K_745_U_AIX_PPC64`
 - (HPUX) – `<Kernel_PATH>/DATA_UNITS/K_745_U_HPUX_IA64`
 - (Solaris) – `<Kernel_PATH>/DATA_UNITS/K_745_U_SOLARIS_SPARC`
 - (Windows) – `<Kernel_PATH>\DATA_UNITS\K_745_U_WINDOWS_X86_64`

Click *Next*.

9. Enter the path to the SAP ASE software package, then click *Next*.
 10. Enter the path to the NetWeaver software package and click *Next*.
 11. Enter the configuration information for SAP ASE:
 - Physical Memory – the amount of memory that SAP ASE is using at any time.
 - Number of Cores – the number of processors in your system.
 - Number of Database Connections – the maximum number of connections to SAP ASE.
- Click *Next*.
12. Enable ABAP table declustering by selecting the button for *Enable declustering/depooling of all ABAP tables*. Click *Next*.
 13. Select the *No SLD destination* option under the *Register in System Landscape Directory* heading. Click *Next*.
 14. Determine the secure storage key by selecting *Default Key* under the *Secure Storage Individual Key Information* heading. Click *Next*.
 15. The installer provides a summary of the configuration. Change any values that are incorrect:
 - *Parameter Settings > Parameter Mode – Typical* for *Parameter Mode*
 - *General SAP System Parameters* – One of:
 - *SAP System ID (SAPSID)* – the value for the global SAP SID, which you supplied earlier in this task.
 - *SAP Mount Directory* – the path to directory
 - *Master Password* – the master password, to be used users to log in to the system.
 - *Software Package Browser* – One of:
 - *Media Location* – the directory on which the installation media is mounted
 - *Package Location* – the directory in which the software package is located
 16. Click *Next* to start the package installation.

4.3.2 Install the Data Movement Component

Procedure

1. Log on the host as user syb<SID>.
2. Create and save a response file for your site using the following sample as input:

```
#
# This responses file installs "SAP ASE Data Movement for HADR" feature for
# Business Suite
#
RUN_SILENT=true
AGREE_TO_SYBASE_LICENSE=true
AGREE_TO_SAP_LICENSE=true
PRODUCTION_INSTALL=TRUE
INSTALL_SETUP_HADR_SAMPLE=true
# Windows only
DO_NOT_CREATE_SHORTCUT=true
REGISTER_UNINSTALLER_WINDOWS=false
INSTALL_USER_PROFILE=USER
```

```

DO_NOT_CREATE_RMA_WINDOW_SERVICE=true
#chadr
INSTALL_SCC_SERVICE=false
USER_INSTALL_DIR=<ASE_installed_directory>
# Install HADR ("SAP ASE Data Movement for HADR" feature)
DO_UPDATE_INSTALL=false
CHOSEN_INSTALL_SET=Custom
CHOSEN_FEATURE_LIST=fase_hadr
CHOSEN_INSTALL_FEATURE_LIST=fase_hadr
INSTALL_SAP_HOST_AGENT=FALSE
# License
SYBASE_PRODUCT_LICENSE_TYPE=license
SYSAM_LICENSE_SOURCE=proceed_without_license
SYSAM_PRODUCT_EDITION=Enterprise Edition
SYSAM_LICENSE_TYPE=AC : OEM Application Deployment CPU License
SYSAM_NOTIFICATION_ENABLE=false
# Do not configure new servers
SY_CONFIG_ASE_SERVER=false
SY_CONFIG_HADR_SERVER=false
SY_CONFIG_BS_SERVER=false
SY_CONFIG_XP_SERVER=false
SY_CONFIG_JS_SERVER=false
SY_CONFIG_SM_SERVER=false
SY_CONFIG_SCC_SERVER=false

```

3. In the line defining the `USER_INSTALL_DIR`, edit the value of `<ASE_installed_directory>` to point to your SAP ASE installation directory. For example:

```
USER_INSTALL_DIR=/sybase/<SID>
```

i Note

On Windows, use the double back slash (\\) to split paths. For example, enter "E:\sybase\<SID>" as "E:\\sybase\\<SID>".

4. Run the SAP ASE installer in silent mode to install the Data Movement component, where `<response_file>` is the absolute path of the file name you just created:
 - o (UNIX) – execute `setup.bin`. Use this syntax:

```
setup.bin -f <response_file> -i silent
```

- o (Windows) – run `setupConsole` utility, using this syntax:

```
.\setupConsole -f <response_file> -i silent
```

5. Log on the host as user `syb<SID>`.
6. Connect to SAP ASE as user `sapsso`:

```
$$SYBASE/$SYBASE_OCS/bin/isql -Usapsso -P<password> -S<server_name>
```

7. Unlock the user `sa`:

```
sp_locklogin sa, unlock
```

4.3.3 Run setuphadr to Configure HADR

Use the setuphadr utility to configure the HADR system.

Procedure

1. Log on to the primary machine as syb<SID>.
2. If you do not already have the setup_hadr.rs response file for your site, make a copy of the sample template described in [Sample setup_hadr.rs Response File for Business Suite \[page 192\]](#), and change the values based on the requirements for your site as you keep the following in mind:
 - If you copy and paste the text from a PDF, remove the running headers and footers similar to this:

```
HADR Users Guide
Installing HADR for Business Suite © 2015 SAP SE or an SAP affiliate
company. AA
All rights reserved. 151
```

- (Windows) Use the double back slash (\\) to split paths. For example, enter "E:\sybase" as "E:\\sybase".
- You do not need the <SID> database as it is automatically created by SWPM (sapinst).
- Set these properties on the primary site:

```
setup_site=<primary_site>
is_secondary_site_setup=false
```

See [Sample setup_hadr.rs Response File for Business Suite \[page 192\]](#) for examples of the changes required.

3. Run setuphadr with the response file:
 - (UNIX) – \$SYBASE/\$SYBASE_ASE/bin/setuphadr <path_to_response_file>
 - (Windows) – %SYBASE%\%SYBASE_ASE%\bin\setuphadr.bat <path_to_response_file>

The output looks similar to:

```
E:\>E:\sybase\NW7\ASE-16_0\bin\setuphadr setup_SFHADR.rs
Setup user databases
  Set "NW7" database "trunc log on chkpt" option to "false"...
Setup user databases...Success
Setup ASE HADR maintenance user
  Create maintenance login "NW7_maint"...
  Grant "sa_role" role to "NW7_maint"...
  Grant "replication_role" role to "NW7_maint"...
  Grant "replication_maint_role_gp" role to "NW7_maint"...
  Create "sap_maint_user_role" role...
  Grant set session authorization to "sap_maint_user_role"...
  Grant "sap_maint_user_role" role to "NW7_maint"...
  Add auto activated roles "sap_maint_user_role" to user "NW7_maint"...
  Allow "NW7_maint" to be known as dbo in "master" database...
  Allow "NW7_maint" to be known as dbo in "NW7" database...
Setup ASE HADR maintenance user...Success
Setup administrator user
  Create administrator login "DR_admin"...
  Grant "sa_role" role to "DR_admin"...
  Grant "sso_role" role to "DR_admin"...
  Grant "replication_role" role to "DR_admin"...
```



```

Grant "hadr_admin_role_gp" role to "DR_admin"...
Grant "sybase_ts_role" role to "DR_admin"...
Setup administrator user...Success
Setup Backup server allow hosts
Backup server on "PRIM" site: Add host "Huge_Machine1.corp" to allow
dump and load...
Setup Backup server allow hosts...Success
Setup complete on "PRIM" site. Please run Setup HADR on "COMP" site to
complete the setup.

```

4.4 Installing the HADR System on the Companion

Configuring the HADR on the companion includes installing NetWeaver, the Data Management software, and running `setuphadr` utility.

4.4.1 Install the Business Suite Application

Install the Business Suite application using the application installation process.

Context

The installation process varies depending on which installation application you use. The steps to install the Business Suite application depends on which application you use. The following example describes the NetWeaver installation process.

Procedure

1. Move to the `sapinst` directory, which was created when the `SAPCAR.exe` utility extracted files.
2. Execute the `sapinst` utility to start the SAP installation GUI
3. Select **▶ SAP NetWeaver 7.5 ▶ SAP ASE ▶ Database Replication ▶ Setup of Replication Environment ▶** and click *Next*.
4. Specify the Replication Server parameters, then click *Next*:
 - *SAP System ID* – comprises three alphanumeric characters and is the same as the SAP SID you entered for the primary
 - *Master Password* – is the same as the master password you entered for the primary
 - *SAP Global Host Name* – is the host name of the machine on which you are installing the software
 - *Set up a secondary database instance* – select to confirm.
 - *Install the replication server software* – leave blank
 - *Configure the replication system* – leave blank

- *Materialize the secondary database* – leave blank
5. Specify the Replication Server parameters, then click *Next*:
 - *Host Name of ASE Server* – the name of the machine hosting the primary server
 - *Port Number of ASE Server* –the port number to connect to the primary server.
 - *Password for sapsa* – the same as Master Password
 - *Password for sapsso* – the same as Master Password
 - *Password for DR_admin* –the same as Master Password
 6. Specify the path to the software package by choosing the path to installation media (DVD or Blu-ray disc) and entering the path in the box provided, or selecting Browse to explore the system.

4.4.2 Install the Data Movement Component

Use a response file to install the Data Movement component.

Procedure

1. Log on the host as user syb<SID>.
2. Create and save a response file for your site using the sample following as input:

```
#
# This responses file installs "SAP ASE Data Movement for HADR" feature for
Business Suite
#
RUN_SILENT=true
AGREE_TO_SYBASE_LICENSE=true
AGREE_TO_SAP_LICENSE=true
PRODUCTION_INSTALL=TRUE
INSTALL_SETUP_HADR_SAMPLE=true
# Windows only
DO_NOT_CREATE_SHORTCUT=true
REGISTER_UNINSTALLER_WINDOWS=false
INSTALL_USER_PROFILE=USER
DO_NOT_CREATE_RMA_WINDOW_SERVICE=true
#chadr
INSTALL_SCC_SERVICE=false
USER_INSTALL_DIR=<ASE_installed_directory>
# Install HADR ("SAP ASE Data Movement for HADR" feature)
DO_UPDATE_INSTALL=false
CHOSEN_INSTALL_SET=Custom
CHOSEN_FEATURE_LIST=fase_hadr
CHOSEN_INSTALL_FEATURE_LIST=fase_hadr
INSTALL_SAP_HOST_AGENT=FALSE
# License
SYBASE_PRODUCT_LICENSE_TYPE=license
SYSAM_LICENSE_SOURCE=proceed without license
SYSAM_PRODUCT_EDITION=Enterprise Edition
SYSAM_LICENSE_TYPE=AC : OEM Application Deployment CPU License
SYSAM_NOTIFICATION_ENABLE=false
# Do not configure new servers
SY_CONFIG_ASE_SERVER=false
SY_CONFIG_HADR_SERVER=false
SY_CONFIG_BS_SERVER=false
SY_CONFIG_XP_SERVER=false
```

```
SY_CONFIG_JS_SERVER=false
SY_CONFIG_SM_SERVER=false
SY_CONFIG_SCC_SERVER=false
```

3. In the line defining the `USER_INSTALL_DIR`, edit the value of `<ASE_installed_directory>` to point to your SAP ASE installation directory. For example:

```
USER_INSTALL_DIR=/sybase/<SID>
```

i Note

On Windows, use the double back slash (\\) to split paths. For example, enter "E:\sybase\<SID>" as "E:\\sybase\\<SID>".

4. Run the installer in silent mode to install the Data Movement component, where `<response_file>` is the absolute path of the file name you just created:

- o (UNIX) – execute the `setup.bin` using this syntax:

```
setup.bin -f <response_file> -i silent
```

- o (Windows) – run the `setupConsole` utility using this syntax:

```
.\setupConsole -f <response_file> -i silent
```

5. Log on the host as user `syb<SID>`.
6. Connect to SAP ASE as user `sapsso`. For example:

```
$SYBASE/$SYBASE_OCS/bin/isql -Usapsso -P<password> -S<server_name>
```

7. Unlock the user `sa`:

```
sp_locklogin sa, unlock
```

4.4.3 Run `setuphadr` to Configure HADR on the Companion

Use the `setuphadr` utility to configure the HADR system.

Procedure

1. Copy the `setup_hadr.rs` file to the companion machine.
2. Log on to the companion machine as `syb<SID>`.
3. Change these properties in the companion version of the `setup_hadr.rs` response file:
 - o `setup_site=site1` – change to `site2` on standby:

```
setup_site=COMP
```

- `is_secondary_site_setup=false` – change to "true" on standby:

```
is_secondary_site_setup=true
```

- `# user database` – comment out user database information
- `# Data & log devices to create the databases specified` – comment out user data and log device information

See [Sample setup_hadr.rs Response File for Business Suite \[page 192\]](#) for an example of the necessary changes.

4. As `syb<SID>`, run `setuphadr` with the response file:

- (UNIX) – `$SYBASE/$SYBASE_ASE/bin/setuphadr <path_to_response_file>`
- (Windows) – `%SYBASE%\%SYBASE_ASE%\bin\setuphadr.bat <path_to_response_file>`

The output looks similar to:

```
./ASE-16_0/bin/setuphadr setup_SJHADR.rs
Setup user databases
  Set "NW7" database "trunc log on chkpt" option to "false"...
Setup user databases...Success
Setup Backup server allow hosts
  Backup server on "COMP" site: Add host "Huge_Machine1.corp" to allow
dump and load...
  Backup server on "PRIM" site: Add host "Huge_Machine2.corp" to allow
dump and load...
Setup Backup server allow hosts...Success
Setup RMA
  Set SAP ID to "NW7"...
  Set installation mode to "BS"...
  Set site name "SFHADR1" with SAP ASE host:port to "Huge_Machine1.corp:
4901" and Replication Server host:port to "Huge_Machine1.corp:4905"...
  Set site name "SJHADR2" with SAP ASE host:port to "Huge_Machine2.corp:
4901" and Replication Server host:port to "Huge_Machine2.corp:4905"...
  Set site name "SFHADR1" with Backup server port to "4902"...
  Set site name "SJHADR2" with Backup server port to "4902"...
  Set site name "SFHADR1" databases dump directory to "/sybase/NW7/
data"...
  Set site name "SJHADR2" databases dump directory to "/sybase/NW7/
data"...
  Set site name "SFHADR1" synchronization mode to "sync"...
  Set site name "SJHADR2" synchronization mode to "sync"...
  Set site name "SFHADR1" distribution mode to "remote"...
  Set site name "SJHADR2" distribution mode to "remote"...
  Set site name "SFHADR1" distribution target to site name "SJHADR2"...
  Set site name "SJHADR2" distribution target to site name "SFHADR1"...
  Set maintenance user to "NW7_maint"...
  Set site name "SFHADR1" device buffer directory to "/sybase/NW7/
data"...
  Set site name "SJHADR2" device buffer directory to "/sybase/NW7/
data"...
  Set site name "SFHADR1" device buffer size to "512"...
  Set site name "SJHADR2" device buffer size to "512"...
  Set site name "SFHADR1" simple persistent queue directory to "/"
sybase/NW7/data"...
  Set site name "SJHADR2" simple persistent queue directory to "/"
sybase/NW7/data"...
  Set site name "SFHADR1" simple persistent queue size to "2000"...
  Set site name "SJHADR2" simple persistent queue size to "2000"...
  Set master, NW7 databases to participate in replication...
Setup RMA...Success
Setup Replication
  Setup replication from "SFHADR1" to "SJHADR2"...
  Configuring remote replication server.....
```

```

        Configuring local replication
server.....
    Setting up replication on 'standby' host for local database
'master'.....
    Setting up replication on 'standby' host for local database
'NW7'.....
Setup Replication...Success
Materialize Databases
    Materialize database "master"...
        Starting materialization of the master database from source
'SFHADR1' to target 'SJHADR2'...
        Completed materialization of the master database from source
'SFHADR1' to target 'SJHADR2'..
        Waiting 10 seconds: Before checking if Replication Connection
'S1_SJHADR2.master' is suspended.....
    Materialize database "NW7"...
        Executing ASE dump and load task for database 'NW7'.....
        Successfully verified materialization on database 'NW7'..
        Stop the Replication Agent for database 'master' on host
'SFMACHINE1.BIG.corp:4901' and data server 'NW7_SFHADR1'..
        Start the Replication Agent for database 'master' on host
'SFMACHINE1.BIG.corp:4901' and data server 'NW7_SFHADR1'..
        Stop the Replication Agent for database 'NW7' on host
'SFMACHINE1.BIG.corp:4901' and data server 'NW7_SFHADR1'..
        Configuring Replication Server: set 'hide_maintuser_pwd' to
'o'...
        Waiting 10 seconds: Before checking if Replication Connection
'S1_SJHADR2.NW7' is suspended.....
        Completed automatic materialization of database 'NW7' from
source 'SFHADR1' to target 'SJHADR2'...
Materialize Databases...Success

```

4.5 Post-Installation Tasks for Primary and Companion Servers

There are a number of tasks you must perform on the primary and companion servers after installation.

4.5.1 Restart Replication Server as a Service on Windows

On a Windows system, after the installation is complete on the primary and companion hosts, Replication Server is running, but not as a service to RMA.

Context

Perform the following on both the primary and companion servers:

Procedure

1. Log into RMA:

```
isql -U<login_name> -P<password> -S<host_name>:<RMA_port_number>
```

2. Issue this to include Replication Server as a service:

```
sap_set_replication_service <logical_name>, <short_hostname>\syb<sid>,  
<password>  
sap_set_replication_service <logical_name>, restart
```

4.5.2 Add the DR_admin Entry to SecureStore

Use the `rsecssfx put` command to add the `DR_admin` entry to SecureStore.

Procedure

As the `<SID>adm` user, enter this at the command line:

```
rsecssfx put DB_CONNECT/SYB/DR_USER DR_admin -plain  
rsecssfx put DB_CONNECT/SYB/DR_PASSWORD <password>
```

4.5.3 Configure Replication Server with sap_tune_rs

Perform these steps on the primary and companion servers to configure Replication Server to specify the maximum number of the CPUs and the maximum size of the memory for the Replication Server instance.

Procedure

1. Log into RMA:

```
isql -U<login_name> -P<password> -S<host_name>:<RMA_port_number>
```

2. Issue sap_tune_rs to configure Replication Server for your site:

```
sap_tune_rs <site_name>, <amount_of_RAM>, <number_of_CPUs>
```

For example, this tunes Replication Server on logical host SFHADR1 with 4 GB memory and 2 CPUs:

```
sap_tune_rs SJHADR2,4,2
```

4.5.4 Add the dbs_syb_ha and dbs_syb_server Users

If you would like to have your SAP Application Server automatically failover to the standby SAP ASE when you perform a database level failover, add the `dbs_syb_ha` and the `dbs_syb_server` user.

- On Windows:
 1. Log into the SAP Primary Application Server (PAS) host as the `<SID>adm` user.
 2. From the *System Properties* window, click the *Advanced* tab and select *Environment Variables*.
 3. Select *New*.
 4. Enter the following, then click *OK*:
 - *Variable Name* – `dbs_syb_ha`
 - *Variable value* – `1`
 5. Select the `dbs_syb_server` user variable, and click *Edit* to enter the following values:
 - *Variable Name* – `dbs_syb_server`
 - *Variable value* – `<host_name><standby_server>`
 6. Click *OK*.
 7. Click *OK*.
 8. Restart NetWeaver.
 9. Log into the SAP Management Console (sapmmc).
 10. Right-click on **Console Root** > **SAP Systems** > **NW7**.
 11. Select *Restart*.
 12. Make the same change on all additional SAP Application Servers.
- On Linux

1. Log into the primary host as the <SID>adm user and add these entries to the .dbenv.csh file:

```
setenv dbs_syb_ha 1
setenv dbs_syb_server <primary_server_name>:<standby_server_name>
```

2. Restart NetWeaver by issuing these commands on the primary server as the <SID>adm user:

```
sapcontrol -nr 00 -function StopSystem ALL
sapcontrol -nr 00 -function StartSystem ALL
```

4.6 Sample setup_hadr.rs Response File for Business Suite

This is a sample setup_hadr.rs file. The text changed for the installation described in this guide is in **bold**.

```
#####
# Setup HADR sample responses file
#
# This sample responses file setup ASE HADR on
# hosts "host1" (primary) and "host2" (companion).
#
# Prerequisite :
# - New SAP ASE and Backup servers setup and started on "host1" and "host2".
# - See HADR User Guide for requirements on SAP ASE servers.
# - Replication Management Agent (RMA) started on "host1" and "host2".
#
# Usage :
# 1. On host1 (primary), run:
#     $SYBASE/$SYBASE_ASE/bin/setuphadr <this_responses_file>
#
# 2. Change this responses file properties:
#     setup_site=COMP
#     is_secondary_site_setup=true
#
# 3. On host2 (companion), run
#     $SYBASE/$SYBASE_ASE/bin/setuphadr <responses_file_from_step_2>
#
#####
# ID that identifies this cluster
#
# Value must be unique,
# begin with a letter and
# 3 characters in length.
# Note: Set value to your SID incase of HADR on SAP Business Suite Installations
cluster_id=NW7
# Which_site being configured
#
# Note:
# You need to set "<setup_site_value>.*"
# properties in this responses file.
setup_site=PRIM
# Set installation_mode
#
# Valid values: true, false
#
# If set to true, installation_mode will be set to "BS".
# If set to false, installation_mode will be set to "nonBS"
# Note: Set value to true for HADR on SAP Business Suite installations
setup_bs=true
# Note: Set enable_ssl to false for HADR on SAP Business Suite Installations
#
```



```

# true OR false
enable_ssl=false
# common name, take SYBASE for example
#ssl_common_name=SYBASE
# private key file
#ssl_private_key_file=/tmp\hadr.key
# public key file
#ssl_public_key_file=/tmp\hadr.crt
# root CA cert
# NOTE: if you're using self-signed cert, put your public key file here
#ssl_ca_cert_file=/tmp\rootCA.pem
# ssl password
#ssl_password=Sybase
# Has the secondary site prepared for ASE HADR
#
# Valid values: true, false
#
# If set to true, "<secondary_setup_site_value>.*"
# properties must set in this responses file.
is_secondary_site_setup=false
# How data is replicated
#
# Valid values: sync, async
synchronization_mode=sync
# SAP ASE system administrator user\password
#
# setuphadr will prompt from standard input if not specified
ase_sa_user=sa
ase_sa_password=<sa_password>
# ASE HADR maintenance user\password
#
# For a Business Suite installation, name the user <cluster_id>_maint.
# Password must have at least 6 characters
# setuphadr will prompt from standard input if not specified
hadr_maintenance_user=NW7_maint
hadr_maintenance_password=<maint_password>
# Replication Management Agent administrator user\password
#
# Password must have at least 6 characters
# setuphadr will prompt from standard input if not specified
rma_admin_user=DR_admin
rma_admin_password=<rma_admin_password>

#If need to disable the checks for reference constraints
#
#Valid values: true, false
disable_referential_constraints=false
# If we need to config and start Replication Management Agent
#
# Valid values: true, false
config_start_rma=true
# If we need to create Replication Management Agent windows service
# Only affects windows
#
# Valid values: true, false
# If set to true, rma_service_user and rma_service_password will be used
create_rma_windows_service=true
# Replication Management Agent Service user\password
# Only needed for windows installations.
# Note: Set value of rma_service_user to syb<sid> user incase of HADR on SAP
Business Suite Installations
rma_service_user=sybnw7
rma_service_password=<rma_service_password>
# Databases that will participate in replication
# and "auto" materialize.
#
# ASE HADR requires SAP ASE to have a database
# with cluster ID name (see "cluster_id" above).

```

```

# cluster ID database
participating_database_1=NW7
materialize_participating_database_1=true
#####
# Site "PRIM" on host host1 with primary role
#####
# Host name where SAP ASE run
#
# Enter fully qualified domain name (FQDN)
# if your sites are on different subnet.
PRIM.ase_host_name=Huge_Machine1.corp
# We don't support ASE and SRS on different hosts yet
# This is virtual host name for SRS\RMA
# Optional property
#
# Enter fully qualified domain name (FQDN)
# if your sites are on different subnet.
PRIM.rma_host_name=Huge_Machine1.corp
# Site name
#
# Enter value that identifies this site,
# like a geographical location.
# Value must be unique.
PRIM.site_name=Site1
# Site role
#
# Enter the role of this site.
# Valid values: primary, companion
PRIM.site_role=primary
# directory where SAP ASE installed
PRIM.ase_release_directory=E:\sybase\NW7
# Directory that stored SAP ASE user data files
# (interfaces, RUN_<server>, error log, etc. files).
# Do not set value if your user data files are in
# SAP ASE installed directory (ase_release_directory).
PRIM.ase_user_data_directory=
PRIM.ase_server_name=NW7
PRIM.ase_server_port=4901
PRIM.backup_server_name=NW7_BS
PRIM.backup_server_port=4902
# Directory to store database dumps
# in materialization
#
# Backup server must able to access this directory
PRIM.backup_server_dump_directory=E:\sybase\NW7\data
# Port numbers for Replication Server and Replication Management Agent on host1
#
# In remote topology, these are the companion Replication Server and
# Replication Management Agent.
#
# See "rsge.bootstrap.tds.port.number" properties in
# <SAP ASE installed directory>\DM\RMA-16_0\instances\AgentContainer\config
# \bootstrap.prop
# for value
PRIM.rma_tds_port=4909
PRIM.rma_rmi_port=7000
# RMA RMI occupies five consecutive ports, with the configured port occupying
# the highest number.
#
# Starting port number to use when setup Replication Server.
# Make sure next two ports (+1 and +2) are also available for use.
PRIM.srs_port=4905
# Device buffer for Replication Server on host1
# Recommend size = 128 * N
#     where N is the number of databases to replicate,
#     including the master and cluster ID databases.
#

```

```

PRIM.device_buffer_dir=E:\\sybase\\NW7\\data
PRIM.device_buffer_size=20000
# Persistent queue directory for Replication Server running on host1
#
# For synchronous replication (synchronization_mode=sync),
# enter directory to an SSD (solid state drive) or other
# type of fast read\\write storage device
PRIM.simple_persistent_queue_dir=E:\\sybase\\NW7\\data
PRIM.simple_persistent_queue_size=20000
#####
# Site "COMP" on host host2 with companion role
#####
# Host name where SAP ASE run
#
# Enter fully qualified domain name (FQDN)
# if your sites are on different subnet.
COMP.ase_host_name=Huge_Machine2.corp
# We don't support ASE and SRS on different hosts yet
# This is virtual host name for SRS\\RMA
# Optional property
#
# Enter fully qualified domain name (FQDN)
# if your sites are on different subnet.
COMP.rma_host_name=Huge_Machine2.corp
# Site name
#
# Enter value that identifies this site,
# like a geographical location.
# Value must be unique.
COMP.site_name=Site2
# Site role
#
# Enter the role of this site.
# Valid values: primary, companion
COMP.site_role=companion
# directory where SAP ASE installed
COMP.ase_release_directory=E:\\sybase\\NW7
# Directory that stored SAP ASE user data files
# (interfaces, RUN_<server>, error log, etc. files).
# Do not set value if your user data files are in
# SAP ASE installed directory (ase_release_directory).
COMP.ase_user_data_directory=
COMP.ase_server_name=NW7
COMP.ase_server_port=4901
COMP.backup_server_name=NW7_BS
COMP.backup_server_port=4902
# Directory to store database dumps
# in materialization
#
# Backup server must able to access this directory
COMP.backup_server_dump_directory=E:\\sybase\\NW7\\data
# Port numbers for Replication Server and Replication Management Agent on host2
#
# In remote topology, these are the companion Replication Server and
# Replication Management Agent.
#
# See "rsge.bootstrap.tds.port.number" properties in
# <SAP ASE installed directory>\\DM\\RMA-16_0\\instances\\AgentContainer\\config
# \\bootstrap.prop
# for value
COMP.rma_tds_port=4909
COMP.rma_rmi_port=7000
# RMA RMI occupies five consecutive ports, with the configured port occupying
the highest number.
# Starting port number to use when setup Replication Server.
# Make sure next two ports (+1 and +2) are also available for use.
COMP.srs_port=4905

```

```

# Device buffer for Replication Server on host2
# Recommend size = 128 * N
#     where N is the number of databases to replicate,
#     including the master and cluster ID databases.
#
# Note: For HADR on SAP Business Suite Installations use SID database logsize *
1.5
COMP.device_buffer_dir=E:\\sybase\\NW7\\data
COMP.device_buffer_size=20000
# Persistent queue directory for Replication Server running on host2
#
# For synchronous replication (synchronization_mode=sync),
# enter directory to an SSD (solid state drive) or other
# type of fast read\\write storage device
# Note: For HADR on SAP Business Suite Installations use SID database logsize *
1.5
COMP.simple_persistent_queue_dir=E:\\sybase\\NW7\\data
COMP.simple_persistent_queue_size=

```

4.7 Using the Fault Manager with Business Suite

The Fault Manager monitors the health of the primary and standby servers, and triggers a failover if the primary server or host fails, and the HADR system is running in synchronous mode.

The Fault Manager is a standalone component that runs on a third node, preferably where the application server is running, and on the same platform as the HADR system nodes.

The Fault Manager functions in two modes: the Fault Manager mode and as the heartbeat client mode. The Fault Manager runs on a third host. In Fault Manager mode, it monitors SAP ASE, Replication Server, performs functions like initiating failover, and restarting the server, and acts as the server for the heartbeats that it receives from the heartbeat clients.

The Fault Manager heartbeat client mode runs on primary and standby hosts. In heartbeat client mode, the Fault Manager sends a heartbeat to the Fault Manager, checks for heartbeats from fellow heartbeat clients, and sends its own heartbeat to them (primarily to avoid a split-brain situation). If the heartbeat client on the primary host loses a connection with the Fault Manager and the fellow heartbeat client, the Fault Manager triggers a deactivation of the primary server. If the deactivation fails, the Fault Manager kills the SAP ASE process.

The Fault Manager checks the database state with the `saphostctrl` SAP host agent, which is a daemon process started on all participating nodes. The Fault Manager also uses `saphostctrl` to connect to the Replication Management Agent. See the chapter titled *The SAP Host Agent*.

The Fault Manager:

- Triggers a failover using `saphostctrl` if the primary server is down or if the primary node is down or unreachable, and the standby server is healthy and synchronously replicated.
- Restarts the primary server if it is down and replication is asynchronous.

i Note

Stop or hibernate the Fault Manager when you perform any maintenance activity on SAP ASE or other components in the HADR system. Once hibernated, the Fault Manager process continues to run but will not monitor the database, and no failover occurs. The heartbeat processes are stopped during hibernation.

This table describes actions performed by the Fault Manager:

Scenario	Action Performed
Primary server is down.	Failover to companion.
Primary host is down.	Failover to companion.
Primary server is unreachable (network glitch or SAP ASE unresponsive).	Retry shallow probe, try deep probe, probe companion, failover to companion, notify cockpit. Attempt restarting primary SAP ASE if possible when HA is off.
Primary server reports an error condition.	If client login and data access are unaffected, no action is taken. Fault Manager does not scan SAP ASE log for errors. If the error results in login failures or data access errors, failover to companion if <code>ha/syb/failover_if_unresponsive=1</code> is included in the <code>profile</code> file.
HA services on the companion are down, resulting in replication from the primary server to the companion server.	Attempt to restart HA services (may need manual intervention).
HA Services on companion are restored.	After syncing up the backlog of transaction logs, automatically switch to sync mode replication, turn on HA.
Fault Manager components on primary are down.	Attempt to restart failed components if <code>ha/syb/chk_restart_repserver=1</code> is included in the <code>profile</code> file, notify cockpit of success or failure of restart and HA on/off status (may need manual intervention).
Fault Manager is down.	In this version of the software, manually restart the Fault Manager (see Administering the Fault Manager [page 150]).
Failover fails.	Attempt to failover again until failover succeeds, or the condition causing failover is rectified (may need manual intervention).
Fault Manager components on companion are down.	Attempt to restart failed components.
Fault Manager is unreachable from the 2 sites (primary and companion).	If the network between the primary and companion is OK, then continue as is. However, if there is a network problem between primary and companion, deactivate the primary to avoid split brain, notify cockpit. HA is off and application has no access to database (will need manual intervention).
Primary and companion are unreachable from Fault Manager.	No action performed and replication continues as normal.
Heartbeat from primary is missed for a preconfigured timeout.	If SAP ASE is not reachable (confirmed by local agent), then failover. If SAP ASE is reachable and HA is not working, restart HA.
Heartbeat from companion is missed for a preconfigured timeout.	If SAP ASE is unreachable (confirmed by local agent), then failover. If SAP ASE is reachable and HA is not working, restart HA.
Companion SAP ASE is down.	Attempt to restart companion SAP ASE if <code>ha/syb/allow_restart_companion=1</code> is included in the <code>profile</code> file.
Companion host is down.	No action performed.

Scenario	Action Performed
Companion SAP ASE is unreachable (network glitch or SAP ASE unresponsive).	Attempt to restart companion SAP ASE when possible.
HA services on primary down, no impact to HA until failover.	Attempt to restart companion SAP ASE when possible.
HA services on primary restored.	No action performed.

4.7.1 Configuring the Fault Manager on Business Suite

4.7.1.1 Secure Store on Business Suite

The Fault Manager requires username and password combinations to connect to SAP ASE, RMA, and SAP Host Agent. These usernames and passwords are stored in an encrypted format in the SecureStore.

During configuration, the Fault Manager adds usernames and passwords for the following users in the SecureStore:

- SADB_USER – SAP ASE user with the sa_role and replication_role roles.
- DR_USER – RMA user, used for connecting to RMA.
- SAPADM_USER – Operating system user, mostly used for sapadm for SAP HostAgent.

Use the `rsecssfx` utility to perform this administration duty for SecureStore. Update any changed usernames and passwords in SecureStore. To do so, stop the Fault Manager, update the SecureStore using the `rsecssfx` utility, and restart the Fault Manager. Stop the Fault Manager while the password is changed in the cluster components.

- Use the SAP Host Agent to add or update entries in the SecureStore. The syntax is:

```
/usr/sap/hostctrl/exe/saphostctrl -user sapadm <password> -function
LiveDatabaseUpdate -dbname <SID> -dbtype syb -dbuser sapsa -dbpass <password>
-updatemethod Execute -updateoption TASK=SET_USER_PASSWORD -updateoption
USER=<parameter_name>
```

Where `<parameter_name>` is one of:

- SAPSA – SAP ASE user with the sa_role and replication_role roles.
 - DR_ADMIN – RMA user, used for connecting to RMA.
 - SAPADM – Operating system user, mostly used for sapadm for SAPHostAgent.
 - SAUPD – SAP ASE user sa created during SAP Netweaver installation
 - SAPSSO – SAP ASE user sapsso created during SAP Netweaver installation
 - SYBSID – SAP ASE user syb<SID> created during SAP Netweaver installation
 - SAPADM – SAP ASE user sapadm created during SAP Netweaver installation
- Use the `list` parameter to list entries in the SecureStore. For example:

```
./FaultManager/bin/rsecssfx list
```

```

| Record Key | Status | Time Stamp of Last Update |
|-----|-----|-----|
| DB_CONNECT/SYB/DR_PASSWORD | Encrypted | 2015-06-09 08:48:29 UTC |
|-----|-----|-----|
| DB_CONNECT/SYB/DR_USER | Plaintext | 2015-06-09 08:48:25 UTC |
|-----|-----|-----|
| DB_CONNECT/SYB/SADB_PASSWORD | Encrypted | 2015-06-09 08:48:22 UTC |
|-----|-----|-----|
| DB_CONNECT/SYB/SADB_USER | Plaintext | 2015-06-09 08:48:18 UTC |
|-----|-----|-----|
| DB_CONNECT/SYB/SAPADM_PASSWORD | Encrypted | 2015-06-09 08:48:35 UTC |
|-----|-----|-----|
| DB_CONNECT/SYB/SAPADM_USER | Plaintext | 2015-06-09 08:48:31 UTC |
|-----|-----|-----|
Summary ----- Active Records : 8 (Encrypted: 4, Plain: 4, Wrong Key: 0,
Error: 0)
Defunct Records : 0 (180+ days: 0; Show: "list -withHistory", Remove:
"compact")

```

i Note

On Windows, include `<installation_directory>/FaultManager/bin` in the `<PATH>` environment variable.

4.7.2 Administering the Fault Manager on Business Suite

Start the Fault Manager from the installation directory, `/usr/sap/<SID>/ASCS00/work`.

Fault Manager is installed by default as part of Kernel utilities during SAP Netweaver Installation in the following location `/usr/sap/<SID>/ASCS00/work`. For details refer to the SAP Note [1959660](#).

Use the `sybdbfm` utility to view the status of the Fault Manager. For example:

```

$ sybdbfm status
fault manager running, pid = 17763, fault manager overall status = OK, currently
executing in mode PAUSING
*** sanity check report (1)***.
node 1: server star1, site hasite0.
db host status: OK.
db status OK hadr status PRIMARY.
node 2: server star2, site hasitel.
db host status: OK.
db status OK hadr status STANDBY.
replication status: SYNC_OK.

```

Edit the Fault Manager profile file to change any parameter. The profile file is named `SYBHA.PFL`, and is located in the install directory of the Fault Manager on all platforms. Restart the Fault Manager for the profile parameter changes to take effect.

You should continuously monitor the Fault Manager log (named `dev_sybdbfm`, and located in `/usr/sap/<SID>/ASCS00/work`).

i Note

If a problem related to Fault Manager or the heartbeat requires you to consult SAP, back up the following data when the problem occurs:

- Fault Manager data available on the host running Fault Manager (\$SYBASE below is the directory where Fault Manager is installed):
 - Fault Manager profile file – located in /sapmnt/<SID>/profile/SYBHA.PFL
 - Fault Manger log file – located in /usr/sap/<SID>/ASCS00/work/dev_sybdbfm
- Host agent log from the primary and standby sites (requires sudo access):
 - Host Agent log file – located in /usr/sap/hostctrl/work/dev_sapdbctrl
 - Heart beat log file – located in /usr/sap/hostctrl/work/dev_sybdbfm

4.7.3 Uninstalling the Fault Manager from Business Suite

How you uninstall the Fault Manager depends on whether you installed it using the SAP installer or the sybdbfm utility.

1. If you used the sybdbfm utility to install the Fault Manager:
 1. Stop the Fault Manager. Move to the directory that contains SYBHA.PFL (the default is <installed_directory>/FaultManager), and issue:

```
sybdbfm stop
```

2. Remove SecureStore-related files by issuing this from the directory that contains SYBHA.PFL:

```
sybdbfm uninstall pf=<path_to_SYBHA.PFL>
```

2. (If you installed Fault Manager on a completely separate host that is not the SAP application server or the database host) Move to \$SYBASE/sybuninstall/FaultManager and issue:

```
./uninstall
```


When the `uninstall` utility displays the Uninstall Fault Manager introductory screen, click [Next](#):




3. (If you installed Fault Manager on a completely separate host that is not the SAP application server or the database host) In the Uninstall Options screen, select the appropriate option.




If you select:

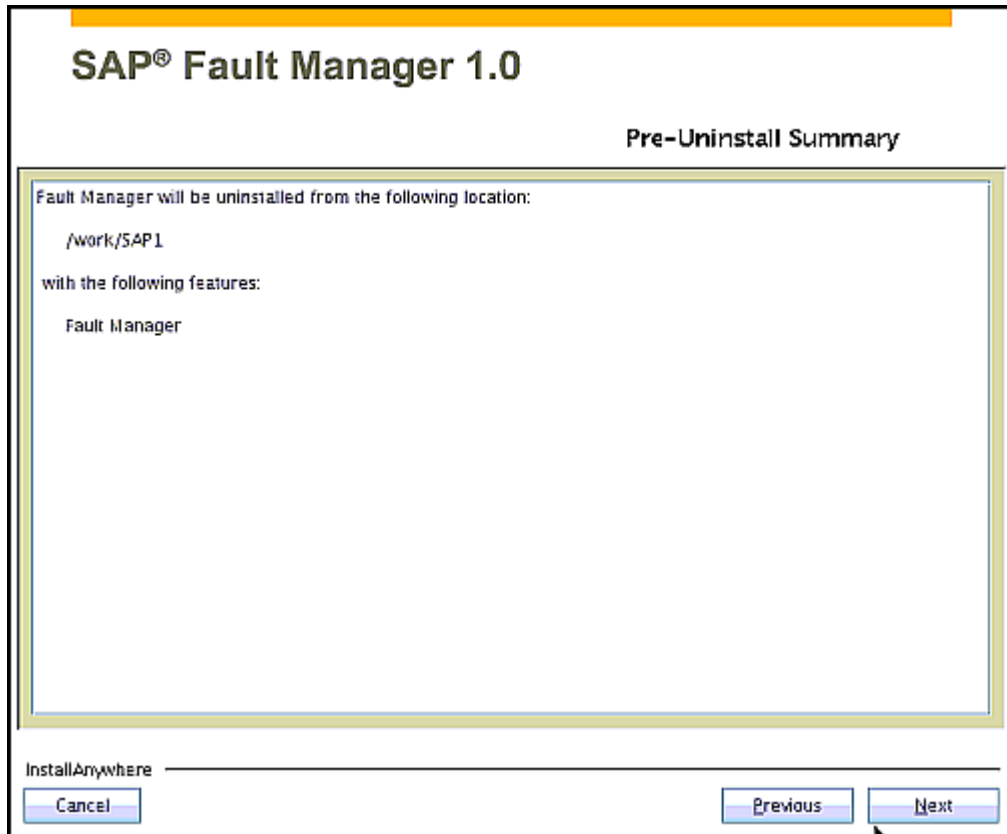
- *Complete Uninstall* – the `uninstall` utility displays a summary of all features and components it will remove.



- *Uninstall Specific Features* – select which available features you want to uninstall from the list uninstall utility displays.
 1. Uncheck the features you want to remove and click *Next*.



2. The `uninstall` utility displays the *Summary* screen.



Click [Previous](#) to return to the previous screen. When you are ready to proceed, click [Next](#).

4. Delete the following files and directories manually (requires sudo permission) as part of uninstalling Fault Manager from the primary and companion hosts:
 1. Delete these directories:
 - `/usr/sap/hostctrl/exe/<SAP_ASE_server_name>`
 - `/usr/sap/hostctrl/work/<SAP_ASE_server_name>`
 - (If present) `/usr/sap/hostctrl/work/<SAP_ASE_server_name>_REP`
 2. Delete these `sybdbfm` files:
 - `/usr/sap/hostctrl/exe/sybdbfm`
 - `/usr/sap/hostctrl/exe/sybdbfm_<SAP_ASE_server_name>`
 - `/usr/sap/hostctrl/exe/SYBHA_<SAP_ASE_server_name>.PFL`
 - `/usr/sap/hostctrl/work/dev_h2h`
 - (If present) `/usr/sap/hostctrl/work/dev_h2h.old`
 - `/usr/sap/hostctrl/work/dev_hbeat`
 - (If present) `/usr/sap/hostctrl/work/dev_hbeat.old`
 - `/usr/sap/hostctrl/work/dev_sybdbfm`

4.8 Performing a Rolling Upgrade

The HADR feature allows SAP ASE applications to operate with zero down time while you are updating the SAP ASE software.

Generally, the steps for performing a rolling upgrade are:

- Upgrade the standby components
- Perform a failover
- Upgrade the components that are on standby after the failover


Complete the upgrade steps in a single sequence: partial upgrade is not supported (for example, you cannot upgrade some components now and then upgrade the other components at another time). Replication is suspended during some steps of a rolling upgrade, and if you perform a partial upgrade, logs continue to grow, which can result in logs or the SPQ running out of space. During a rolling upgrade, the versions between SAP ASE and Replication Server need not match.

The `RUN_rs instance name.sh` Replication Server runserver file is regenerated during an upgrade, and any user changes to this file are lost. If your site requires these changes, edit the runserver file after the upgrade is complete then restart Replication Server to make the environment settings take effect.

i Note

Before upgrading HADR with SAP Business Suite on SAP ASE, you may have to follow instructions from your application vendors over the general guidelines in this chapter. See SAP note [2808173](#) for more details.

The steps below describe performing a rolling upgrade on an HADR system with this remote topology:




In this topology, the primary server (ASE1) is installed on the same host as the inactive Replication Server (SRS1). The active Replication Server (SRS2) is installed on a remote host, along with the standby server (ASE2). Data changes that occur in ASE1 are sent by the Replication Agent thread to the active SRS2 running on the remote host. The active SRS2 then routes these changes to ASE2, which is running on the same host as the active Replication Server, SRS2. In this setup, the inactive Replication Server, SRS1, is not involved in data movement until failover occurs. The communication among ASE1, SRS1, and ASE2 is through a client interface (stream replication, indicated in this topology as "CI").

Run this command to determine which SAP ASE you are connected to in the HADR system:

```
select asehostname ()
```

The initial topology configuration looks like:



In this configuration, all components are running, and the standby server is almost in sync with the primary server. Prior to upgrade, site1 is the primary server and site2 is the companion server (in high-availability – HA – mode, the companion server is referred to as the standby server) with remote replication topology. The Replication Server versions prior the upgrade are compatible with the Replication Server versions after the upgrade. If you upgrade from a "1-Off" release, you can upgrade only the SAP ASE or Replication Servers.


To perform a rolling upgrade:

i Note

Stop the Fault Manager before you perform a rolling upgrade (even if you are performing planned activities like a planned failover). You can start the Fault Manager after the upgrade is complete. To stop the Fault Manager, issue this from the `<installation_directory>/FaultManager` directory:

```
<Fault_Manager_install_dir>/FaultManager/bin/sybdbfm stop
```

To perform a rolling upgrade, you first upgrade SRS1 on site1 to a higher version:



1. Verify that Replication Server is not running any `isql` processes during this Replication Server installation step. If there are `isql` processes running, Replication Server issues an error message stating "isql text file busy".
2. Log into RMA on site1 as the DR_admin user and issue `sap_upgrade_server` to prepare the Replication Server on site1 for upgrade:

```
sap_upgrade_server SRS, start, <site1_site_name>
```

3. Shut down RMA:

```
shutdown
```


4. Remove the RMA service: On Windows, execute the following command from either the `%SYBASE%\RMA-16_0\compatibility\WinService\Win32\Release` directory, or the `%SYBASE%\RMA-16_0\compatibility\WinService\x64\Release` directory, to remove the RMA service –

```
drservice -remove <Cluster_ID>
```

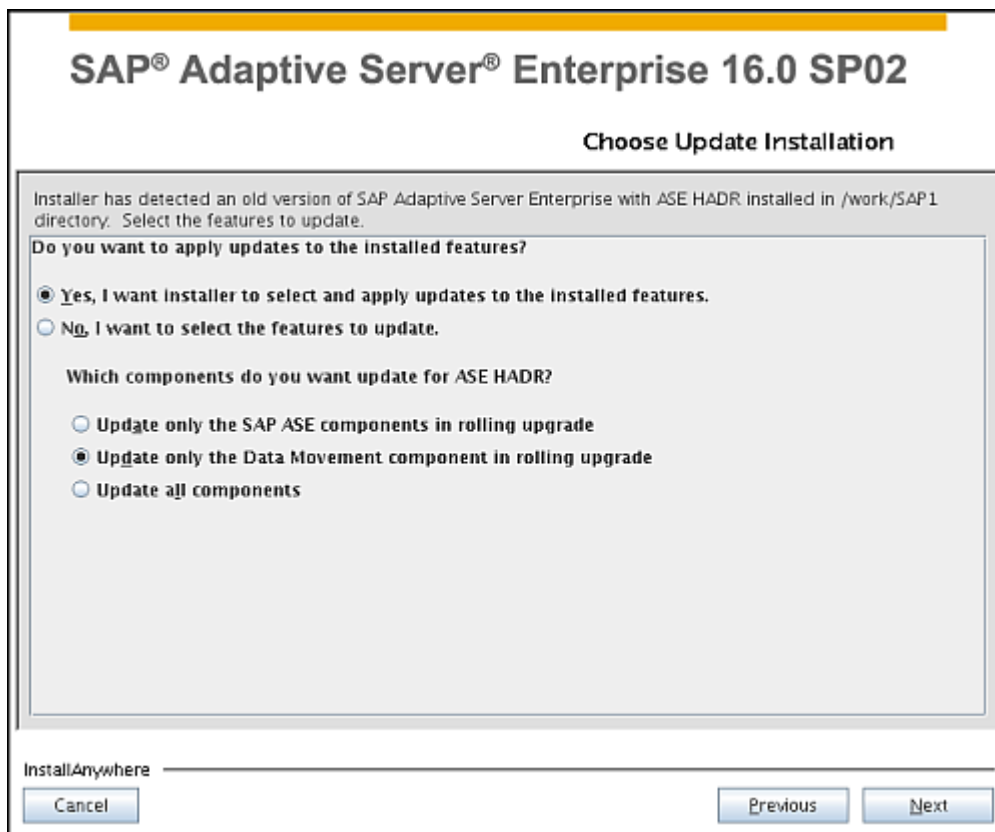
5. Start the SAP installer from the installation directory:

```
install_directory/setup.bin
```

6. In the Choose Install Folder screen, enter the current SAP ASE SAP installation directory, then click *Next*:



7. In the Choose Update Installation screen, determine if you want the installer to select and apply updates, then select *Update only the Data Movement component in rolling upgrade*.



i Note

The SAP installer must complete the software update before you continue to the next step.

8. Install a new RMA service on Windows: To install, and then start the new RMA service on Windows, execute the following command from either the %SYBASE%\RMA-16_0\compatibility\WinService\Win32\Release directory, or the %SYBASE%\RMA-16_0\compatibility\WinService\x64\Release directory –

```
drservice -install <Cluster_ID>
```

9. After the installer finishes the upgrade, start RMA:
 - (UNIX) – issue \$SYBASE/\$SYBASE_ASE/bin/rma
 - (Windows) – start the RMA Windows service by either of the following:
 - Starting *Sybase DR Agent* - <cluster_ID> from the Services panel
 - Issuing this command, where <cluster_ID> is the ID of the cluster:

```
net start SybaseDRAgent_<cluster_ID>
```

10. Log in to RMA on site1 as the DR_admin user and issue sap_upgrade_server to finish the upgrade for Replication Server on site1:

```
sap_upgrade_server SRS, finish,<site1_site_name>
```

Note

At this point of the upgrade process, the HADR system is working normally with ASE1, SRS2, ASE2 at the older versions, and SRS1 at newer version.

11. Log into RMA on site1 as the DR_admin user and issue:

```
sap_failover <site1_site_name>, <site2_site_name>, 30
```

This command allows a 30-second grace period for any running transactions to complete before the deactivation starts. Failover will not succeed if there are still active transactions after 30 seconds. If this occurs, retry the command when the system is not busy, use a longer grace period, or use the `force` option to terminate the client connection (if it is safe) with:

```
sap_failover <site1_site_name>, <site2_site_name>, 30, force
```


12. The `sap_failover` command may take a long time to finish. To check the status of the `sap_failover` command, issue this from the RMA:

```
sap_status task
```

13. Once the `sap_status` command returns `Completed`, resume replication by issuing this from the RMA:

```
sap_host_available <site1_site_name>
```

At this point, the topology looks like:



14. Verify that Replication Server is not running any `isql` processes during the Replication Server installation step below. If there are `isql` processes running, Replication Server issues an error message stating "`isql text file busy`".
15. Login to RMA on site2 as the DR_admin user and issue `sap_upgrade_server` to start the upgrade for Replication Server on site2:

```
sap_upgrade_server SRS, start,<site2_site_name>
```

16. Shut down RMA:

```
shutdown
```

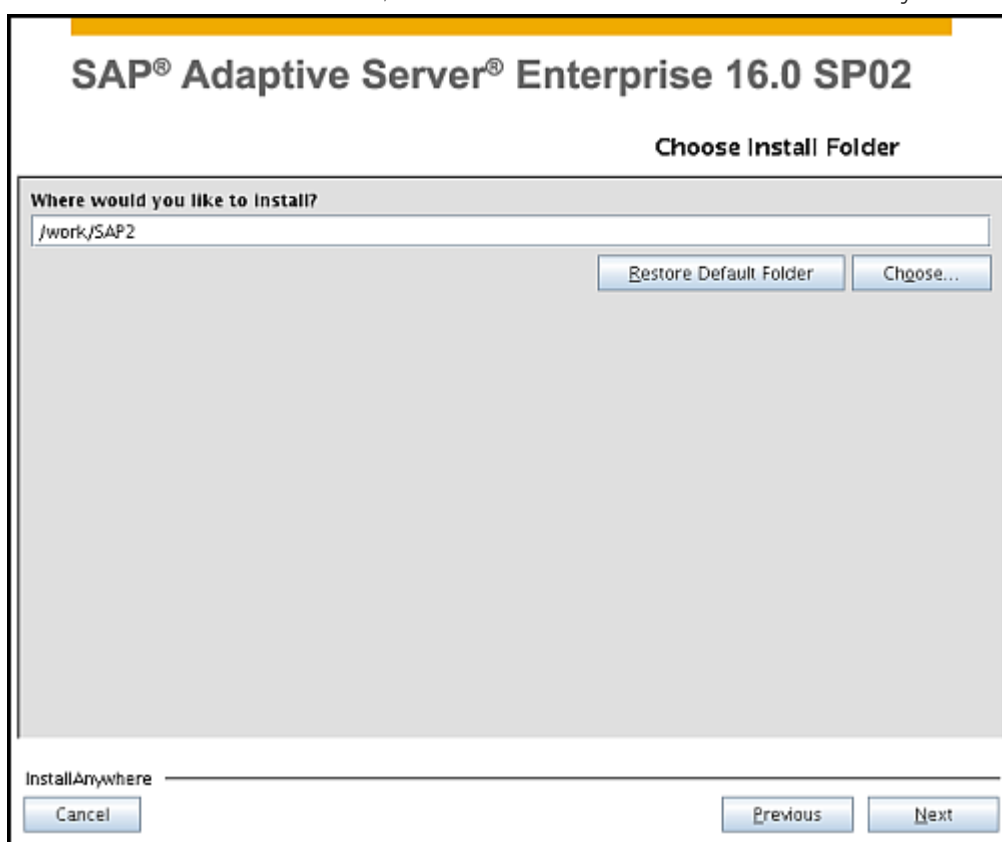
17. Remove the RMA service: On Windows, execute the following command from either the %SYBASE%\RMA-16_0\compatibility\WinService\Win32\Release directory, or the %SYBASE%\RMA-16_0\compatibility\WinService\x64\Release directory, to remove the RMA service –

```
drservice -remove <Cluster_ID>
```

18. Start the SAP installer from the installation directory:

```
<install_directory>/setup.bin
```

19. In the Choose Install Folder screen, enter the current ASE SAP installation directory.



20. In the Choose Update Installation screen, determine if you want the installer to select and apply updates, then select *Update only the Data Movement component in rolling upgrade*.



i Note

The SAP installer must complete the software update before you continue to the next step.

21. Install a new RMA service on Windows: To install, and then start the new RMA service on Windows, execute the following command from either the %SYBASE%\RMA-16_0\compatibility\WinService\Win32\Release directory, or the %SYBASE%\RMA-16_0\compatibility\WinService\x64\Release directory –

```
drservice -install <Cluster_ID>
```

22. After the SAP installer has finished the upgrade, start RMA:
 - (UNIX) – \$SYBASE/\$SYBASE_ASE/bin/rma
 - (Windows) – start the RMA Windows service by either of the following:
 - Starting *Sybase DR Agent* - <cluster_ID> from the Services panel
 - Issuing this command, where <cluster_ID> is the ID of the cluster:

```
net start SybaseDRAgent_<cluster_ID>
```

23. Log into RMA on site2 as the DR_admin user and issue sap_upgrade_server to finish the upgrade for Replication Server on site2:

```
sap_upgrade_server SRS, finish, <site2_site_name>
```


Note

At this point of the upgrade process, the HADR system is working normally with ASE1 and ASE2 at the older versions, and SRS1 and SRS2 at the newer version.

24. (Skip this step if you do not lock the sa user) If the sa user is locked, temporarily unlock this user on ASE1 during the upgrade process by logging in as the user with SSO permission on ASE1 and issuing:

```
sp_locklogin sa, "unlock"
```

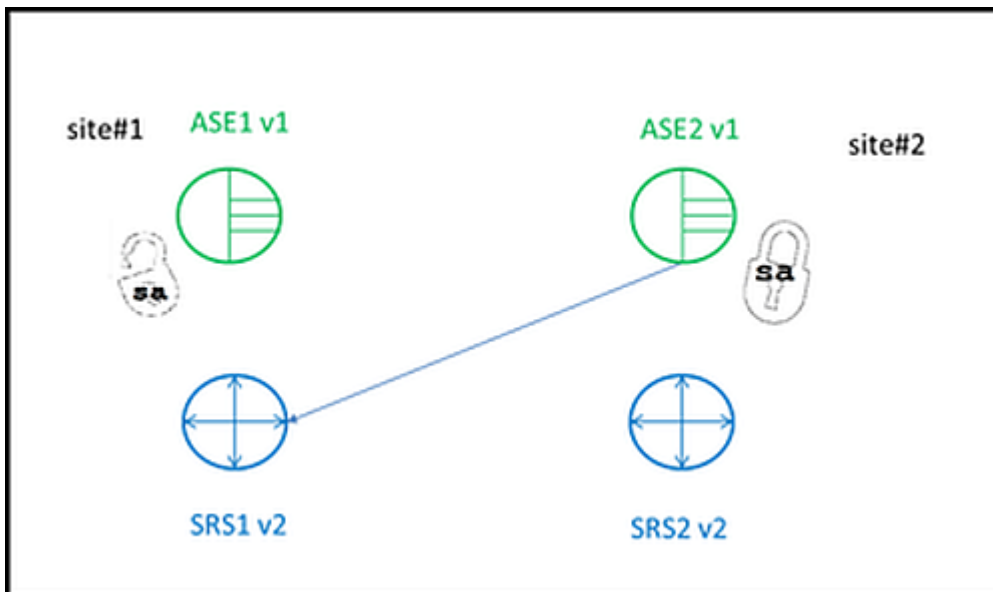
At this point, the topology looks like:



25. Log into RMA on site1 as the DR_admin user and issue `sap_upgrade_server` to start the upgrade for SAP ASE on site1:

```
sap_upgrade_server ASE, start, <site1_site_name>
```

The topology now looks like:



26. Shut down SAP ASE and Backup Server on site1. Login to site1 SAP ASE server and issue:

```
shutdown SYB_BACKUP
go
shutdown
go
```

27. Shut down SAP ASE Cockpit. If the SAP ASE Cockpit is running:

- o In the foreground – At the `cockpit>` prompt, execute:

```
shutdown
```

Enter a reason for the shutdown.

i Note

Do not enter `shutdown` at a UNIX prompt: Doing so shuts down the operating system.


- o In the background – At the UNIX command line, execute:

```
$SYBASE/COCKPIT-4/bin/cockpit.sh --stop
```

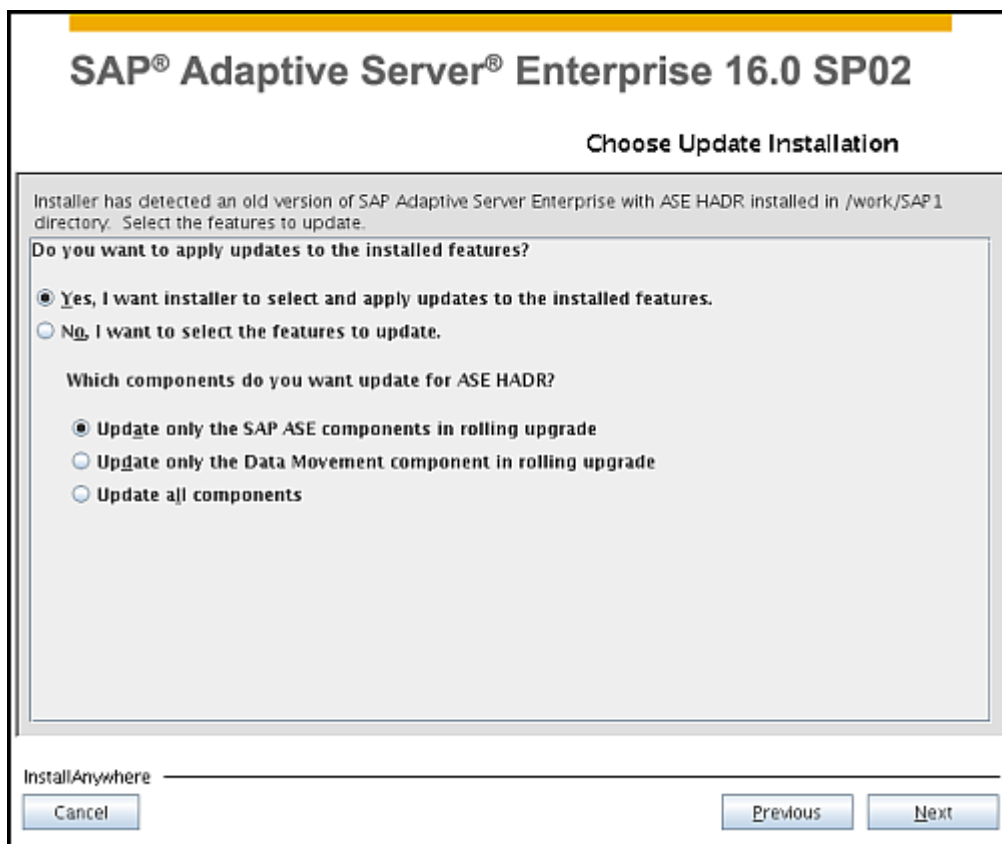
28. Start the SAP installer from the installation directory:

```
<install_directory>/setup.bin
```

29. In the Choose Install Folder screen, enter the current ASE SAP installation directory.



30. In the Choose Install Folder screen, determine if you want the installer to select and apply updates, then select *Update only the SAP ASE component in rolling upgrade*.



i Note

The SAP installer must complete the software update before you continue to the next step.

31. After the SAP installer has completed the upgrade, use the `updatease` utility to upgrade SAP ASE, which runs `installmaster` and performs other tasks to bring SAP ASE up to date. `updatease` is available as a GUI display or a command line tool.

i Note

You need not perform this step if you update the SAP ASE server instance in the SAP installer.

1. Start `updatease`. From the `$$SYBASE/$SYBASE_ASE/bin/` directory, issue:

```
./updatease
```

2. Enter the server name and password.

The command line version of `updatease` displays messages similar to:

```
./updatease
Server: SFSAP1
ASE Password:
Updating SAP Adaptive Server Enterprise 'SFSAP1'...
Running installmaster script...
installmaster: 10% complete.
installmaster: 20% complete.
installmaster: 30% complete.
```

```

installmaster: 40% complete.
. . .
instmsgs.ebf: 80% complete.
instmsgs.ebf: 90% complete.
instmsgs.ebf: 100% complete.
The instmsgs.ebf script has been successfully installed.
Restarting SAP Adaptive Server Enterprise...
SAP Adaptive Server Enterprise has been restarted.
Updating SAP Adaptive Server Enterprise complete successfully.

```

32. Start Backup Server from the command line by issuing:

```
$SYBASE/ASE-16_0/install/RUN_<backupserver_name>
```

33. Start the SAP ASE Cockpit.

- In the foreground – from the command line, execute:

```
$SYBASE/COCKPIT-4/bin/cockpit.sh
```

- In the background:

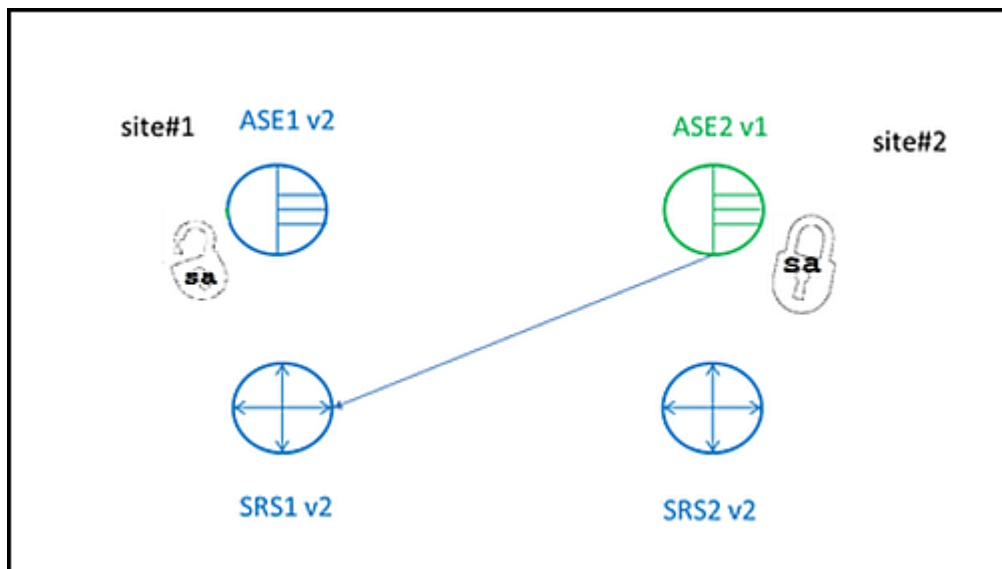
From the Bourne shell (sh) or Bash, issue:

```
nohup ./cockpit.sh 2>&1 > cockpit-console.out &
```

From the C shell, issue:

```
nohup ./cockpit.sh >& cockpit-console.out &
```

The topology now looks like:



34. Log into RMA on site1 as the DR_admin user and issue `sap_upgrade_server` to complete the upgrade for SAP ASE on site1:


```
sap_upgrade_server ASE, finish,<site1_site_name>
```

35. (Skip this step if you do not lock the sa login) Log in to ASE1 as a user with `sapsso` permission and issue this once the upgrade is complete:

```
sp_locklogin sa, "lock"
```


Note

At this point of the upgrade process, the HADR system is working normally, with ASE2 at the older versions and SRS1, SRS2, and ASE1 at the newer versions. The topology looks like:



36. Log into RMA on site2 as the DR_admin user and issue:

```
sap_failover <site2_site_name>, <site1_site_name>, 30
```

This command allows a 30-second grace period for any running transactions to complete before the deactivation starts. Failover will not succeed if there are still active transactions after 30 seconds. If this occurs, retry the command when the system is not busy, use a longer grace period, or use the `force` option to terminate the client connection (if it is safe) with:

```
sap_failover <site2_site_name>, <site1_site_name> 30, force
```


37. The `sap_failover` command may take a long time to finish. To check the status of the `sap_failover` command, issue this from the RMA:

```
sap_status task
```

38. Once the `sap_status` command returns `Completed`, resume replication from ASE1 to SRS2, but not to ASE2, by issuing this from the RMA (`suspend` ensures that no ticket is sent to verify the replication status during the upgrade process):

```
sap_host_available <site2_site_name>, suspend
```

The topology looks like:




The replication path from ASE1 to SRS2 is restored and all commits on ASE1 are synchronously committed on SRS2, meaning there is zero data loss if site1 is lost during this time. Because replication to ASE2 is not restored, the log records in SRS2 are not applied to ASE2. Make sure you have sufficient storage configured in the SPQ for the SRS2 for the short period of time until the upgrade process completes.

39. (Skip this step if you do not lock the sa login) Temporarily unlock the sa user on ASE2 for upgrade. Login as the user with sapsso permission on ASE2 and issue:

```
sp_locklogin sa, "unlock"
```

The topology looks like:



40. Log into RMA on site2 as the DR_admin user and issue `sap_upgrade_server` to start the upgrade for SAP ASE on site2 (`suspend` ensures that no ticket is sent to verify the replication status during the upgrade process):

```
sap_upgrade_server ASE, start, <site2_site_name>, suspend
```

41. Shut down SAP ASE and Backup Server on site2. Login to site2 SAP ASE server and issue:

```
shutdown SYB_BACKUP
go
shutdown
go
```

42. Shutdown SAP ASE Cockpit. If the SAP ASE Cockpit is running:

- o In the foreground – At the `cockpit>` prompt, execute:

```
shutdown
```

Enter a text reason for the shutdown.

i Note

Do not enter `shutdown` at a UNIX prompt; it shuts down the operating system.

- o Running in the background – At the UNIX command line, execute:

```
$$SYBASE/COCKPIT-4/bin/cockpit.sh --stop
```


43. Start the SAP installer from the installation directory:

```
install_directory/setup.bin
```

44. Enter the current ASE SAP installation directory in the Choose Install Folder screen.



45. In the Choose Update Installation screen, determine if you want the installer to select and apply updates, then select *Update only the SAP ASE component in rolling upgrade*:



Note


The SAP installer must complete the software update before you continue to the next step.

46. After the SAP installer has completed the upgrade, use the `updatease` utility to upgrade SAP ASE, which runs `installmaster` and performs other tasks to bring SAP ASE up to date. See the `updatease` instructions in the earlier step.

47. Start Backup Server from the command line by issuing:

```
$SYBASE/ASE-16_0/install/RUN_<backupserver_name>
```

The topology looks like:



48. Start the SAP ASE Cockpit.

- In the foreground – from the command line, execute:

```
$SYBASE/COCKPIT-4/bin/cockpit.sh
```

- In the background:

From the Bourne shell (sh) or Bash, issue:

```
nohup ./cockpit.sh 2>&1 > cockpit-console.out &
```

From the C shell, issue:

```
nohup ./cockpit.sh >& cockpit-console.out &
```

49. Log into RMA on site2 as the DR_admin user and issue `sap_upgrade_server` to complete the upgrade for SAP ASE on site2:


```
sap_upgrade_server ASE, finish,<site2_site_name>, suspend
```

50. (Skip this step if you do not lock the sa login) Log in to ASE2 as the user with `sapsso` permissions and issue this once the upgrade is complete:

```
sp_locklogin sa, "lock"
```

iNote

At this point of the upgrade process, the HADR system is working normally, with SRS1, SRS2, ASE1, and ASE2 at the newer versions. The topology looks like:



51. Run Fault Manager installer:

```
<installation_directory>/FaultManager/setup.bin
```

52. Choose the existing ASE-installed directory on the *Choose Install Folder* screen. Do not choose to configure Fault Manager in the installer.

53. Set the environment variables, and source SYBASE.sh for the Bourne shell:

```
source <installation_directory>/SYBASE.csh
```

54. Start the Fault Manager:

```
<Fault_Manager_install_dir>/FaultManager/sybdbfm_<CID>
```

4.9 Upgrading SAP ASE Version 15.7 DR to Version 16.0 HADR

These sections describe how to upgrade an SAP ASE version 15.7 disaster recovery (DR) solution to ASE 16.0 HADR solution in a Business Suite or a Custom Application environment.

Before you begin, it is important to consider:

- These steps require an application downtime, the length of which depends on the size of the databases you are upgrading.
- If you do not follow the upgrade steps precisely as documented and in the proper sequence, you will lose data.

4.9.1 Upgrade the SAP Host Agent

The `SAPHOSTAGENT<SP-version>.SAR` archive contains all of the required elements for centrally monitoring any host. It is available for all operating system platforms supported by SAP.

Context

The SAP Host Agent is automatically installed during the installation of SAP systems or instances with SAP kernel 7.20 or higher.

Procedure

1. Go to the *Software Downloads* of the *SAP Support Portal* at <http://support.sap.com/swdc>.
2. Log on with your *SAP Support Portal ID*.
3. Choose ► *Installation and Upgrades* ► *By Alphabetical Index (A-Z)* ► *H* ► *SAP Host Agent* ► *SAP Host Agent 7.21* ► Select highest available version.
4. Select the appropriate `SAPHOSTAGENT<SP-version>.SAR` archive from the *Download* tab.

→ Recommendation

Always select the highest SP version of the `SAPHOSTAGENT<SP-version>.SAR` archive, even if you want to monitor a component of SAP NetWeaver with a lower release.

5. Make sure that the `SAPCAR` tool is available on the host where you want to install SAP Host Agent.

Decompressing the `SAPHOSTAGENT<SP-version>.SAR` archive requires the `SAPCAR` tool. To download the `SAPCAR` tool:

- a. Go to the *Software Downloads* of the *SAP Support Portal* at <http://support.sap.com/swdc>.
- b. Choose ► *Installation and Upgrades* ►
- c. Select the highest available version.

6. Log on as a user with the required authorization:

Windows	As a member of the local <code>Administrators</code> group
----------------	--

UNIX	As a user with <code>root</code> authorization
-------------	--

7. Copy the downloaded `SAPHOSTAGENT<SP-version>.SAR` archive to a temporary directory, for example:

Windows	<code>c:\temp\hostagent</code>
----------------	--------------------------------

UNIX	<code>/tmp/hostagent</code>
-------------	-----------------------------

8. Extract the `SAPHOSTAGENT<SP-version>.SAR` archive using `SAPCAR`.

Take the information provided by SAP Note [212876](#) into account when doing so. Use the following command for extraction, and execute them in the directory of the archive:

Windows `<path to SAPCAR>\sapcar.exe -xvf <path to temporary directory>\SAPHOSTAGENT<SP-version>.SAR -manifest SIGNATURE.SMF`

UNIX `/<path to SAPCAR>/sapcar -xvf <path to temporary directory>/SAPHOSTAGENT<SP-version>.SAR -manifest SIGNATURE.SMF`

Among others, the archive contains the `saphostexec` program.

9. Perform the upgrade by running the following command:

Windows `<path to temporary directory>\saphostexec.exe -upgrade`

→ Recommendation

You can use the additional parameter `-verify` to verify the content of the installation package against the SAP digital signature

UNIX `/<path to temporary directory>/saphostexec -upgrade`

→ Recommendation

You can use the additional parameter `-verify` to verify the content of the installation package against the SAP digital signature

The progress of the upgrade is displayed on the command line.

10. After the upgrade has finished successfully, you can check the version of the upgraded host agent by executing the following command from the directory of the SAP Host Agent executables:

Windows `"%ProgramFiles%\SAP\hostctrl\exe\saphostexec.exe" -version`

UNIX

- If you are logged on as a user with `root` authorization, the command is as follows: `/usr/sap/hostctrl/exe/saphostexec -version`
- If you are logged on as a member of the `sapsys` group, for example `<sapsid>adm`, the command is as follows: `/usr/sap/hostctrl/exe/hostexecstart -version`

4.9.2 Upgrading to an HADR System

Procedure

1. Log into SAP Primary Application Server using user `<sid>adm`:
2. Enter the password and respond to the prompts.

- Stop the Netweaver system. At the prompt, enter:

```
stopsap r3
```

- Check the ticket table on the standby site. Log in with `isql` and issue:

```
use master
go
select count(*) from rs_ticket_history
go
-----
          3
```

Verify this is the same ticket history in the replicated database (in this example, TIA):

```
use TIA
go
select count(*) from rs_ticket_history
go
-----
          3
```

- Log in to the RMA, and check the SAP ASE transaction backlog and the Replication Server queue backlog using `sap_send_trace` RMA command (this example runs the command against the `PRI` database):

```
sap_send_trace PRI
TASKNAME          TYPE          VALUE
-----
Send Trace        Start Time    Mon Nov 16 04:42:02 EST 2015
Send Trace        Elapsed Time  00:00:00
Send Trace        Task Name     Send Trace
Send Trace        Task State    Completed
Send Trace        Short Description  Send a trace through the
Replication system using rs_ticket
Send Trace        Long Description  Successfully sent
Send Trace        Host Name     Big_Host
(7 rows affected)
1> sap_status resource
2> go
Name              Type          Value
-----
15:37:55.327     Start Time    2016-08-24
                  Elapsed Time  00:00:00
                  Estimated Failover Time  0
PRI               Replication device size (MB)  15360
PRI               Replication device usage      112
COM               Replication device size (MB)  15360
COM               Replication device usage      128
PRI.master        ASE transaction log (MB)      300
PRI.master        ASE transaction log backlog (MB)  0
PRI.master        Replication queue backlog (MB)  0
COM.master        Replication queue backlog (MB)  0
PRI.TIA           ASE transaction log (MB)      10240
PRI.TIA           ASE transaction log backlog (MB)  0
PRI.TIA           Replication queue backlog (MB)  0
COM.TIA           Replication queue backlog (MB)  0
(15 rows affected)
```

6. Recheck the ticket table on the standby site, making sure the tickets were received. Log in with `isql` and issue:

```
use master
go
select count(*) from rs_ticket_history
go
-----
          4
```

Verify this is the same ticket history in the replicated database (in this example, TIA):

```
use TIA
go
select count(*) from rs_ticket_history
go
-----
          4
```

7. Uninstall the primary or the standby Replication Server (RMA internally tears down the entire HADR system, including the primary and standby Replication Servers and the users and roles, once you issue `sap_teardown`). Log into either RMA and issue:

```
sap_teardown
```

8. Upgrade SAP ASE to version 16.0. You can upgrade the primary and the standby servers at the same time. Issue this from the command line:

```
saphostctrl -user sapadm - -function LiveDatabaseUpdate -dbname
<database_name> -dbtype syb -dbuser <user_name> -dbpass <password> -
updatemethod Execute -updateoption TASK=UPDATE_ASE -updateoption
DROP_LOCATION="<path_to_package>"
```

For example:

```
/usr/sap/hostctrl/exe/saphostctrl -user sapadm - -function LiveDatabaseUpdate
-database TIA -dbtype syb -dbuser sapsso -dbpass sybase123 -updatemethod
Execute -updateoption TASK=UPDATE_ASE -updateoption DROP_LOCATION="/hadr/
packages/SYBASE_LINUX_X86_64"
```

9. Install the Data Movement option on the primary and standby servers using the silent install method.
 - a. Prepare a response file according to the instructions in [Installing the HADR System with Response Files, Console, and Silent Mode \[page 67\]](#), indicating you are installing only the "SAP ASE Data Movement for HADR."
 - b. Log on to the host as user `syb<sid>`.
 - c. Execute the response file according to these instructions [Installing the HADR System in Silent Mode \[page 79\]](#).
10. Unlock the sa user on the primary and standby SAP ASE servers. Perform these steps on both companions:
 - a. Log in to the primary SAP ASE database as user `sapsso`.
 - b. Issue:

```
sp_locklogin sa, 'unlock'
```

11. Unlock user sa on the primary and standby SAP ASE servers. Perform these steps on both companions:
 - a. Log in to the primary SAP ASE database as user `sapsso`.

b. Issue:

```
sp_locklogin sa, 'unlock'
```

12. Configure SAP ASE for the HADR environment. Follow the instructions in [Installing HADR with an Existing System \[page 86\]](#) for configuring the primary and standby servers.

13. Lock user sa on the primary and standby SAP ASE servers. Perform these steps on both companions:

a. Log in to the primary SAP ASE database as user sapsso.

b. Issue:

```
sp_locklogin sa, 'lock'
```

14. Use the `sap_status` command to check the Replication Server status after the upgrade. Issue this at the RMA `isql` prompt:

```
sap_status path
```

You should see this line for the Replication Server Status in the output:

```
Replication Server Status    Active    The status of Replication Server.
```

15. Configure the RMA:

a. Log in to the RMA as user DR_admin

b. Issue this command to configure the Replication Server parameters:

```
sap_tune_rs <PrimarySiteName>, <RAM_Allocated_for_SRS>, <CPU Cores>
```

For example:

```
sap_tune_rs Big_Host, 8, 2
```

c. Execute the same command for the standby site:

```
sap_tune_rs Other_Big_Host, 8, 2
```

16. Start the SAP Application server. Issue:

```
startsap r3
```

5 HADR Security

An HADR system allows you to maintain confidentiality of data by encrypting client-server communications using Secure Sockets Layer (SSL) session-based security.

i Note

SSL is supported on Linux x64, Linux x86-64, and Windows x64.

SSL is the standard for securing the transmission of sensitive information over the Internet, including credit card numbers, stock trades, and banking transactions.

SSL uses certificates issued by certificate authorities (CAs) to establish and verify identities. A certificate is like an electronic passport; it contains all the information necessary to identify an entity, including the public key of the certified entity and the signature of the issuing CA.

You can enable SSL security for the following HADR scenarios:

- Inside an HADR system. See [Enabling SSL for the HADR System \[page 228\]](#).
- For an external replication system. See [Configuring SSL for External Replication \[page 229\]](#).
- For the Fault Manager. See [Configuring the Fault Manager in an SSL-Enabled HADR Environment \[page 231\]](#).

Implementing HADR SSL features requires a knowledgeable system security officer familiar with the security policies and needs of your site, and who has a general understanding of SSL and public-key cryptography.

5.1 Enabling SSL for the HADR System

To enable SSL functionality when you set up a new HADR system, use the `setuphadr` utility.

Configure the following parameters in the `setup_hadr.rs` file to enable SSL functionality. Then, run the `setuphadr` utility to configure the interface files, and enable SSL automatically.

```
# Enable SSL - true or false
enable_ssl=true

# SSL common name - This is the name of your SAP ASE server. It must be the
common name you use to generate your server
# certificates.
ssl_common_name=YOUR_HADR_SERVERNAME

# Checks whether SSL is enabled on your SAP ASE server.
# If SSL is not enabled on your server, set the value of this parameter to false
(which is also the default
# value for this parameter).
# If SSL is enabled on your ASE server, set the value to true.
# After you set the value to true, setuphadr extracts the necessary certificates
and private keys
# from your existing configuration.
# Note: When this parameter is set to true, the configurations
ssl_private_key_file, and ssl_public_key_file, are ignored.
```

```
ase_ssl_enabled=false

# Name of the SSL private key file
ssl_private_key_file=/tmp/hadr.key

# Name of the SSL public key file
ssl_public_key_file=/tmp/hadr.cert

# Name and location of the Root CA certificate. If you are using a self-signed
certificate, put your
# public key file here.
ssl_ca_cert_file=/tmp/rootCA.pem

# SSL password to protect your private key - this is the same password you used
while creating your certificates
# and private keys.
ssl_password=password
```

i Note

The `setuphadr` utility uses the following default values for the backup server credential if you don't specify your own values in the response file:

```
bs_admin_user=sa
bs_admin_password=
```

For details on key generation, see *Enabling SSL* in the *SAP ASE Security Administration Guide*.

i Note

The RMI ports used by RMA also support SSL encryption.

Related Information

[Sample setup_hadr.rs Response File \[page 104\]](#)

5.2 Configuring SSL for External Replication

The process required to enable SSL for an external replication system differs from the process required to enable SSL for an HADR system. In an external replication system, manual configurations are required to enable SSL.

5.2.1 Enable SSL for Replicating into an HADR System

You can enable SSL security to replicate data from an external SAP Replication Server to an HADR system.

Prerequisites

SSL is enabled in your HADR system. See [Enabling SSL for the HADR System \[page 228\]](#).

Procedure

1. Ensure that `$$SYBASE/config/trusted.txt` contains the CA certificates and public keys needed to access the SSL-enabled HADR system.
2. In the interfaces file, add an SSL entry to the external Replication Server.

Example:

```
SAMPLE_RS
master tcp ether localhost 11752 ssl="CN=SAMPLE_RS.sap.com"
query tcp ether localhost 11752 ssl="CN=SAMPLE_RS.sap.com"
```

In this example, `ssl="CN=SAMPLE_RS.sap.com"` is an SSL entry, and, `"CN=SAMPLE_RS.sap.com"` is the Common Name of the certificate used by the server.

3. If necessary, establish a connection between the external Replication Server and the HADR System:

```
create connection to <SAP_ASE_server.dbname>
set error class to rs_sqlserver_error_class
set function string class to rs_sqlserver_function_class
set username to <maintenance_user_name>
set password to <maintenance_user_password>
connect to hadr
```

If the connection is already established, restart the system to enable SSL.

5.2.2 Enable SSL for Replicating Out from an HADR System

You can enable SSL security to replicate data from an HADR system to an external SAP Replication Server.

Prerequisites

SSL is enabled in your HADR system. See [Enabling SSL for the HADR System \[page 228\]](#).

Procedure

1. Set up and enable SSL services on the SAP Replication Server. See *Setting Up SSL Security on Replication Server* and *Enable SSL Security on Replication Server* in the *SAP Replication Server Administration Guide: Volume 1*.
2. Establish a connection between the HADR system and the external SAP Replication Server:

```
create connection to <SAP_ASE_server.dbname>
set error class to rs_sqlserver_error_class
set function string class to rs_sqlserver_function_class
set username to <maintenance_user_name>
set password to <maintenance_user_password>
with primary only
connect to hadr
set spq_agent_username to <spq_agent_username>
set spq_agent_password to <password>
```

3. After the connection is established, perform the following tasks:
 - a. Configure the SAP Agent by adding an SSL entry into the SAP Replication Server server name:

```
configure spq_agent <dbname> set <rs_servername> to
<rs_servername_with_ssl_entry>
```

- b. Stop the SPQ Agent thread for external replication:

```
suspend spq_agent <dbname>
```

- c. Start the SPQ Agent thread for external replication:

```
resume spq_agent <dbname>
```

5.3 Configuring the Fault Manager in an SSL-Enabled HADR Environment

Configuring the Fault Manager in an SSL-Enabled HADR environment requires tasks that include creating and configuring personal security environments for the client and server.

i Note

- SSL is supported on platforms Linux x64 on x86-64, and Windows x64.
- Tasks in this section are applicable to HADR for Custom Application, and not to HADR for Business Suite.

The following table lists the shared library recommendations to extend SSL support in the SAP Host Agent and the heartbeat client. Shared libraries for all platforms are provided in the Fault Manager installer.

Table 4: Shared Library Recommendations

Platform	Required shared libraries
Windows	<ul style="list-style-type: none"> • sapcrypto.dll • slcryptokernel.dll • slcryptokernel.dll.sha256
Linux/Unix	<ul style="list-style-type: none"> • libslcryptokernel.so • libsapcrypto.so • libslcryptokernel.so.sha256

For more information, refer to the tasks that follow.

5.3.1 Create a Server PSE for the SAP Host Agent on Database Hosts

Perform this task to create a personal security environment (PSE), `SAPSSLS.pse`, for the SAP Host Agent on both database hosts.

Procedure

1. Create the security directory using the following command:

- (Linux)

```
sudo mkdir /usr/sap/hostctrl/exe/sec
```

- (Windows)

```
%PROGRAMFILES%\SAP\hostctrl\exe> mkdir sec
```

2. (Linux) Assign the ownership for the `sec` directory to the `sapadm:sapsys` user:

```
sudo chown sapadm:sapsys /usr/sap/hostctrl/exe/sec
```

3. Use the following commands to set up the shared library search path (`LD_LIBRARY_PATH`, `LIBPATH`, or `SHLIB_PATH`) and `SECUDIR` environment variables, and change to the `exe` directory of SAP Host Agent:

- (Linux)

```
export LD_LIBRARY_PATH=/usr/sap/hostctrl/exe/  
export SECUDIR=/usr/sap/hostctrl/exe/sec  
cd /usr/sap/hostctrl/exe
```

- (Windows) Use the following command to set the `SECUDIR` environment variable:

```
%PROGRAMFILES%\SAP\hostctrl\exe> set SECUDIR=%PROGRAMFILES%\SAP\hostctrl  
\exe\sec
```


i Note

Set up SECUDIR as an absolute path to avoid trouble with the `sapgenpse` tool.

4. Logged in as the `sapadm` user, execute the following using a fully qualified domain name as your host name (such as `myhost.wdf.sap.corp`). This command creates a server PSE file named `SAPSSLS.pse`, which authenticates `myhost.wdf.sap.corp` for incoming SSL connections. Access to this file will require a password; include the `-r` option to direct the certificate-signing request to a file; omit it if you plan to copy and paste the certificate signing request (CSR) into a web formula:

- o (Linux)

```
sudo -u sapadm LD_LIBRARY_PATH=/usr/sap/hostctrl/exe SECUDIR=/usr/sap/hostctrl/exe/sec /usr/sap/hostctrl/exe/sapgenpse gen_pse -p SAPSSLS.pse -x <password> -r /tmp/myhost-csr.p10 "CN=myhost.wdf.sap.corp, O=SAP AG, C=DE"
```

- o (Windows)

```
%PROGRAMFILES%\SAP\hostctrl\exe\sapgenpse.exe gen_pse -p SAPSSLS.pse -x <password> -r myhost-csr.p10 "CN=myhost.wdf.sap.corp, O=SAP AG, C=DE"
```

5. Grant the server PSE access to the `sapadm` operating system user:

- o (Linux)

```
sudo -u sapadm LD_LIBRARY_PATH=/usr/sap/hostctrl/exe SECUDIR=/usr/sap/hostctrl/exe/sec /usr/sap/hostctrl/exe/sapgenpse seclogin -p SAPSSLS.pse -x <password> -O sapadm
```

- o (Windows)

```
%PROGRAMFILES%\SAP\hostctrl\exe\sapgenpse.exe seclogin -p SAPSSLS.pse -x <password> -O sapadm
```

6. Get the certificate:
 - a. Send the certificate signing request to an appropriate certifying authority (CA).
 - a. If the CA replies to the request file with a CA response file (which contains the signed certificate in the PKCS#7 format), use the file as input for importing the signed certificate into the server PSE.

i Note

When the PKCS#7 format is used for signing the certificate, the default name of the certificate file is `myhost.p7b`.

7. Import the signed certificate into the server PSE:

- o (Linux)

```
sudo -u sapadm LD_LIBRARY_PATH=/usr/sap/hostctrl/exe SECUDIR=/usr/sap/hostctrl/exe/sec /usr/sap/hostctrl/exe/sapgenpse import_own_cert -p SAPSSLS.pse -x <password> -c /tmp/myhost.p7b
```

- o (Windows)

```
%PROGRAMFILES%\SAP\hostctrl\exe\sapgenpse.exe import_own_cert -p SAPSSLS.pse -x <password> -c myhost.p7b
```

8. Verify the certificate chain:

- o (Linux)

```
sudo -u sapadm LD_LIBRARY_PATH=/usr/sap/hostctrl/exe SECUDIR=/usr/sap/hostctrl/exe/sec /usr/sap/hostctrl/exe/sapgenpse get_my_name -p SAPSSLS.pse -x <password> -v
```

- o (Windows)

```
%PROGRAMFILES%\SAP\hostctrl\exe\sapgenpse.exe get_my_name -p SAPSSLS.pse -x <password> -v
```

9. Export the certificate, so it can be used by the Fault Manager:

- o (Linux)

```
sudo -u sapadm LD_LIBRARY_PATH=/usr/sap/hostctrl/exe SECUDIR=/usr/sap/hostctrl/exe/sec /usr/sap/hostctrl/exe/sapgenpse export_own_cert -p SAPSSLS.pse -x <password> -r -f x509 -o /tmp/<serverCA>.cer
```

- o (Windows)

```
%PROGRAMFILES%\SAP\hostctrl\exe\sapgenpse.exe export_own_cert -p SAPSSLS.pse -x <password> -r -f x509 -o <serverCA>.cer
```

Next Steps

i Note

For detailed commands and instructions to create a PSE for the SAP Host Agent on UNIX, Windows and IBMi environments, see *SSL Configuration for the SAP Host Agent* under *SAP NetWeaver AS for ABAP innovation package*.

5.3.2 Install the Fault Manager

Install, but do not configure the Fault Manager.

Use the SAP GUI installer (run the installer without the response file) to install the Fault Manager, and enter the required parameters manually.

See [Installing and Configuring the Fault Manager \[page 115\]](#).

5.3.3 Create a Client PSE on the Fault Manager Host for the SAP Host Agent

Perform this task to create a personal security environment (PSE), `SAPSSLC.pse`, for the client () on the Fault Manager host for the SAP Host Agent.

Procedure

1. Create the `FM_INSTALL_DIRECTORY/sec` directory.
2. Export `SECUDIR` to the `FM_INSTALL_DIRECTORY/sec` directory:

```
LD_LIBRARY_PATH=$SYBASE/Faultmanager/lib:$LD_LIBRARY_PATH
SECUDIR=$SYBASE/FaultManager/sec
cd $SYBASE/FaultManager/
```

Where `$SYBASE` is release area of the Fault Manager as generated by the installer.

3. Create the client PSE, `SAPSSLC.pse`, and the certificate signing request (CSR):

- o (Linux)

```
$SYBASE/FaultManager/bin/sapgenpse gen_pse -p SAPSSLC.pse -x <password>
-r /tmp/myhost-csr.p10 "CN=myhost.wdf.sap.corp, O=SAP AG, C=DE"
```

- o (Windows)

```
%SYBASE\FaultManager\bin\sapgenpse.exe gen_pse -p SAPSSLC.pse -x <password>
-r \tmp\myhost-csr.p10 "CN=myhost.wdf.sap.corp, O=SAP AG, C=DE"
```

i Note

Use the fully qualified domain name (FQDN) as your host name. For example:
`myhost.wdf.sap.corp`.

4. Configure the user with permission to start and access the Fault Manager without a password:

- o (Linux)

```
$SYBASE/FaultManager/bin/sapgenpse seclogin -p SAPSSLC.pse -x <password> -
O <FM_OS_USER>
```

- o (Windows)

```
%SYBASE\FaultManager\bin\sapgenpse.exe seclogin -p SAPSSLC.pse -x
<password> -O <FM_OS_USER>
```

5. Get the certificate:

- a. Send the certificate signing request to an appropriate certifying authority (CA).
- b. If the CA replies to the request file with a CA response file (which contains the signed certificate in the PKCS#7 format), use the file as input for importing the signed certificate into the server PSE.

i Note

When the PKCS#7 format is used for signing the certificate, the default name of the certificate file is `myhost.p7b`.

6. Import the signed certificate:

- o (Linux)

```
$SYBASE/FaultManager/bin/sapgenpse import_own_cert -p SAPSSLC.pse -x  
<password> -c /tmp/myhost.p7b
```

- o (Windows)

```
%SYBASE\FaultManager\bin\sapgenpse.exe import_own_cert -p SAPSSLC.pse -x  
<password> -c \tmp\myhost.p7b
```

7. Verify the certificate chain:

- o (Linux)

```
$SYBASE/FaultManager/bin/sapgenpse get_my_name -p SAPSSLC.pse -x  
<password> -v
```

- o (Windows)

```
%SYBASE\FaultManager\bin\sapgenpse.exe get_my_name -p SAPSSLC.pse -x  
<password> -v
```

8. Import the certificate you exported after you created the server PSE. See task [Create a Server PSE for the SAP Host Agent on Database Hosts \[page 232\]](#) for more information.

- o (Linux)

```
$SYBASE/FaultManager/bin/sapgenpse maintain_pk -p SAPSSLC.pse -x  
<password> -a /tmp/serverCA.cer
```

- o (Windows)

```
%SYBASE\FaultManager\bin\sapgenpse.exe maintain_pk -p SAPSSLC.pse -x  
<password> -a \tmp\serverCA.cer
```

5.3.4 Create a Server PSE on the Fault Manager Host for the Heartbeat Client

Perform this task to create a server personal security environment (PSE), `SAPSSLS.pse`, on the Fault Manager host for the Heartbeat client.

Procedure

1. Set up the shared library search path (`LD_LIBRARY_PATH`, `LIBPATH` or `SHLIB_PATH`) and `SECUDIR` environment variables, and change to the `exe` directory of SAP Host Agent as you did in when you created the server PSE for the SAP host agent on the database host.
2. Create the server PSE, `SAPSSLS.pse`, and the certificate signing request (CSR). Run the following command as the `FM_OS` user so that the created files are owned by this user:

- o (Linux)

```
$SYBASE/FaultManager/bin/sapgenpse gen_pse -p SAPSSLS.pse -x <password> -r /tmp/myhost-csr.p10 "CN=myhost.wdf.sap.corp, O=SAP AG, C=DE"
```

- o (Windows)

```
%SYBASE\FaultManager\bin\sapgenpse.exe gen_pse -p SAPSSLS.pse -x <password> -r \tmp\myhost-csr.p10 "CN=myhost.wdf.sap.corp, O=SAP AG, C=DE"
```

i Note

Use the fully qualified domain name (FQDN) as your host name. For example:
`myhost.wdf.sap.corp`.

3. Grant SAP Host Agent access to the server PSE:

- o (Linux)

```
$SYBASE/FaultManager/bin/sapgenpse seclogin -p SAPSSLS.pse -x <password> -O <FM_OS_USER>
```

- o (Windows)

```
%SYBASE\FaultManager\bin\sapgenpse.exe seclogin -p SAPSSLS.pse -x <password> -O <FM_OS_USER>
```

4. Get the certificate:
 - a. Send the certificate signing request to an appropriate certifying authority (CA).
 - b. If the CA replies to the request file with a CA response file (which contains the signed certificate in the PKCS#7 format), use the file as input for importing the signed certificate into the server PSE.

i Note

When the PKCS#7 format is used for signing the certificate, the default name of the certificate file is `myhost.p7b`.

5. Import the certificate:

- (Linux)

```
$SYBASE/FaultManager/bin/sapgenpse import_own_cert -p SAPSSL.S.pse -x  
<password> -c /tmp/myhost.p7b
```

- (Windows)

```
%SYBASE\FaultManager\bin\sapgenpse.exe import_own_cert -p SAPSSL.S.pse -x  
<password> -c \tmp\myhost.p7b
```

6. Verify the certificate chain:

- (Linux)

```
$SYBASE/FaultManager/bin/sapgenpse get_my_name -p SAPSSL.S.pse -x  
<password> -v
```

- (Windows)

```
%SYBASE\FaultManager\bin\sapgenpse.exe get_my_name -p SAPSSL.S.pse -x  
<password> -v
```

7. Export the certificate, so it can be used by the heartbeat host:

- (Linux)

```
$SYBASE/FaultManager/bin/sapgenpse export_own_cert -p SAPSSL.S.pse -x  
<password> -r -f x509 -o /tmp/<serverCA>.cer
```

- (Windows)

```
%SYBASE\FaultManager\bin\sapgenpse.exe export_own_cert -p SAPSSL.S.pse -x  
<password> -r -f x509 -o \tmp\<serverCA>.cer
```

5.3.5 Create a Client PSE for the Heartbeat Client on Database Hosts

Perform this task to create a client personal security environment (PSE), `SAPSSLC.pse`, for the heartbeat client on database hosts, by using the `sapgenpse` binary located on the Fault Manager host.

Context

- Execute all commands in this task as the root user (for Windows, this is the user that has Administrator privileges).
- The `sapgenpse` binary is located in the `$SYBASE/FaultManager/bin/sapgenpse` directory on the Fault Manager host. Do not use the binary located in `/usr/sap/hostctrl/exe/`.

Procedure

1. Create a directory to save the client PSE in. Perform one of:

- Issue this command:

```
mkdir /usr/sap/hostctrl/exe/sec/<CID>/hb/sec
```

- Issue this command, specifying a different path to save the Heartbeat client PSE in than the default path, specify the path in the `<hb_secudir>` parameter (`ha/syb/hb_secudir = <hb_secudir>`) in the Fault Manager profile file, then execute the following:

```
mkdir <hb_secudir>
```

Where `<hb_secudir>` is your custom path.

2. (Linux) Set up the shared library search path (`LD_LIBRARY_PATH`, `LIBPATH` or `SHLIB_PATH`) and `SECUDIR` environment variables, and change to the `exe` directory of SAP Host Agent:

```
export LD_LIBRARY_PATH=/usr/sap/hostctrl/exe/  
export SECUDIR=/usr/sap/hostctrl/exe/<CID>/hb/sec
```

i Note

Set up `SECUDIR` as an absolute path in order to avoid trouble with the `sapgenpse` tool.

3. Create the client PSE, `SAPSSLC.pse`, and the certificate signing request (CSR). Run the following command as the `sapadm` user so that the created files are owned by this user:

- (Linux)

```
sapgenpse gen_pse -p SAPSSLC.pse -x <password> -r  
/tmp/myhost-csr.p10 "CN=myhost.wdf.sap.corp, O=SAP AG, C=DE"
```

- (Windows)

```
%PROGRAMFILES%\SAP\hostctrl\exe\sapgenpse.exe gen_pse -p SAPSSLC.pse -x  
<password> -r myhost-csr.p10 "CN=myhost.wdf.sap.corp, O=SAP AG, C=DE"
```

i Note

Use the fully qualified domain name (FQDN) as your host name. For example:
`myhost.wdf.sap.corp`.

4. Configure the user with permission to start and access the Fault Manager without a password (the default user is `root`):

- (Linux)

```
sapgenpse seclogin -p SAPSSLC.pse -x <password> -O root
```

- (Windows)

```
%PROGRAMFILES%\SAP\hostctrl\exe\sapgenpse.exe seclogin -p SAPSSLC.pse -x  
<password> -O root
```

5. Get the certificate:

- a. Send the certificate signing request to an appropriate certifying authority (CA).
- b. If the CA replies to the request file with a CA response file (which contains the signed certificate in the PKCS#7 format), use the file as input for importing the signed certificate into the server PSE.

i Note

When the PKCS#7 format is used for signing the certificate, the default name of the certificate file is `myhost.p7b`.

6. Import the certificate:

- o (Linux)

```
sapgenpse import_own_cert -p SAPSSLC.pse -x <password> -c /tmp/myhost.p7b
```

- o (Windows)

```
%PROGRAMFILES%\SAP\hostctrl\exe\sapgenpse.exe import_own_cert -p  
SAPSSLC.pse -x  
<password> -c myhost.p7b
```

7. Verify the certificate chain:

- o (Linux)

```
sapgenpse get_my_name -p SAPSSLC.pse -x <password> -v
```

- o (Windows)

```
%PROGRAMFILES%\SAP\hostctrl\exe\sapgenpse.exe get_my_name -p SAPSSLC.pse -x  
<password> -v
```

8. Import both certificates that you generated when you created the server PSE on the Fault Manager host for the heartbeat client:

- o (Linux)

```
sapgenpse maintain_pk -p SAPSSLC.pse -x <password> -a /tmp/serverCA.cer
```

- o (Windows)

```
%PROGRAMFILES%\SAP\hostctrl\exe\sapgenpse.exe maintain_pk -p SAPSSLC.pse -x  
<password> -a serverCA.cer
```

9. Restart the SAP Host Agent on port 1129:

```
/usr/sap/hostctrl/exe/saphostexec -restart
```


5.3.6 Set Profile Parameters in the Fault Manager

Set the profile parameters for Fault Manager:

Procedure

1. Set the following profile parameters for SSL:

```
ha/syb/db_ssl=0
// Enable/disable SSL at the database-level (1 to enable, 0 to disable; in
this sample it is set to 1 (enabled)
ha/syb/db_ssl_certificate
// Path to the trusted.txt file
ha/syb/primary_ssl_dbport
// SSL-enabled Database port for primary
ha/syb/standby_ssl_dbport
// SSL-enabled Database port for standby
ha/syb/ssl=1
// SAP Host Agent is SSL-enabled. When set to 0, it is not SSL-enabled.
ha/syb/hb_ssl=1
// When set to 1, Heartbeat communication is SSL-enabled; when set to 0, it
is not SSL-enabled.
ha/syb/ssl_anon=1
// SAP Host Agent is SSL-enabled, but client certificate not sent/not
verified (anonymous)
ha/syb/hb_ssl_anon=1
// When set to 1, Heartbeat communication is SSL-enabled, but the client
certificate not sent/not verified (anonymous)
DIR LIBRARY = <work dir>
// Location of SAP crypto libraries on the Fault Manager host
DIR INSTANCE = <work dir>
// Location of the 'sec' directory on the Fault Manager host
ha/syb/secudir = <secudir>
// Path to the Server PSE on the Fault Manager host
```

2. (Optional) Set the custom path to the heartbeat client on the database hosts with the following parameter:

```
ha/syb/hb_secudir = <hb_secudir>
// Custom path to the Heartbeat client PSE on the database hosts
```

5.4 Encrypting Databases in an HADR Environment

You can encrypt databases running in an HADR environment.

Prerequisites

Data is replicated, but failover does not occur, if a database in the primary companion is encrypted but is not encrypted in the standby companion. If the primary database is encrypted, the database on the secondary companion must be encrypted with the same keys.

You can create new database encryption keys if your HADR system has numerous user databases for encryption, but you need only a single master key for the system.

i Note

For Business Suite, you have to set up encryption before setting up HADR as the setup of standby with Software Provisioning Manager (SWPM) will copy the same encryption keys to the standby host.

To encrypt a database:

Procedure

1. Configure the primary and secondary companions for encryption:

```
sp_configure 'enable encrypted columns', 1
```

2. Configure the primary and secondary companions for the number of worker processes:

```
sp_configure 'number of worker processes', 10
```

3. On the primary companion:

- a. Create a master key for all database encryption key encryption:

```
create encryption key master with passwd '<password>'
```

- b. Set the encryption key for the master key:

```
set encryption passwd '<password>' for key master
```

- c. Create an encryption key for the database encryption:

```
create encryption key <key_name> for database encryption
```

- d. Query `sysencryptkeys` for the master key (you will need the value of the `eksalt`, `value`, `status` columns, `eksalt 01000e7662f4d97ac74d01` below, for the secondary companion):

```
select eksalt,value,status from sysencryptkeys where id =
object_id('sybencrmasterkey')
-----
01000e7662f4d97ac74d01
321c4c49cd5869d2d947ce20b40953cf653dee0c2040ec4d9446ef8f858cb86ba1b5bc6a7c6
b8ecfb23ec05ed3942c3201 256
```

- e. Query `sysencryptkeys` for the database encryption key (you will need this information for the secondary companion, `0100908f61ef71ebe29c01` below):

```
select eksalt,value,status from sysencryptkeys where id =
object_id('<key_name>')
-----
0100908f61ef71ebe29c01
3ab16e1b0684d6b7b7b3916cb4fb839ee241ec30cfdd686a637af3934c4c8430b80bb6263c9
c4ad5e9d2b148fc50e37c01 2049
```

4. On the secondary companion:

- a. Create the encryption key for the master key using the `eksalt`, `value`, `status` values from the primary companion, 01000e7662f4d97ac74d01 below.

```
create encryption key master for AES with passwd 0x01000e7662f4d97ac74d01
init_vector random keyvalue
0x321c4c49cd5869d2d947ce20b40953cf653dee0c2040ec4d9446ef8f858cb86ba1b5bc6a7
c6b8ecfb23ec05ed3942c3201 keystatus 256
```

- b. Create the encryption key for database encryption using the `eksalt`, `value`, `status` values from the primary companion, 0100908f61ef71ebe29c01 below:

```
create encryption key dek_db1 for AES for database encryption with
keylength 256 passwd 0x0100908f61ef71ebe29c01 init_vector
random keyvalue
0x3ab16e1b0684d6b7b7b3916cb4fb839ee241ec30cfd686a637af3934c4c8430b80bb6263
c9c4ad5e9d2b148fc50e37c01 keystatus
2049
```

5. On the primary and secondary companions:

- a. Set the encryption key for the master key:

```
set encryption passwd '<password>' for key master
```

- b. Alter the database to use encryption:

```
alter database <database_name> encrypt with <key_name>
```

- c. Use the `dbencryption_status` function to check the status of the encryption (database encryption is a background process):

```
dbencryption_status('status',db_id('<database_ID>'))
```

6. Update the SAP ASE runserver file or Windows Service with the `--master-key-passwd=<password>` parameter for the database encryption key, or to configure the primary and secondary companions for encryption with `automatic_startup`.

i Note

SAP Business Suite product require an automatic startup.

To configure SAP ASE to automatically use encryption when it starts:

1. Enable SAP ASE to automatically access the master key:

```
sp_configure 'automatic master key access', 1
```

2. Alter the master key to use encryption when SAP ASE starts:

```
alter encryption key master with passwd <password> add encryption for
automatic_startup'
```

i Note

The database encryption key is named `sapdbkey` by default. See SAP Note 2224138: <https://launchpad.support.sap.com/#/notes/0002224138>.

6 Replicating Data In and Out of an Existing HADR System

You can replicate data, including stored procedures and SQL statements, from an existing HADR system to an external system, or into an HADR system from an external system.

6.1 Requirements and Restrictions

There are a number of requirements and restrictions when replicating data from an HADR system to an external system and vice versa.

The requirements are:

- The HADR system and the external SAP Replication Server must use the same platform.

6.2 Architectural Overview

HADR systems with external replication can replicate data from the external replication system to the HADR system, and vice versa.

An external replication system includes all components in a replication system except the current HADR system with which you are working. The system can be an SAP Replication Server that contacts either an HADR system or an SAP ASE server.

The HADR system contains a primary companion and a secondary companion, and each companion includes an SAP ASE and an SAP Replication Server.

Replicating Data into the HADR System

In this architecture, the external SAP Replication Server replicates data to the primary SAP ASE server in the HADR system. The primary server is available to the external system for replication even if all other servers in the HADR system are down, because the external system replicates into the primary server without going through the SAP Replication Server inside the HADR system.

Even though the primary SAP ASE server is exposed to the external replication system, the details of the HADR system, such as routes and subscriptions are all hidden from the external system. The failover steps within the HADR system do not change when the external replication system replicates data into this HADR system.

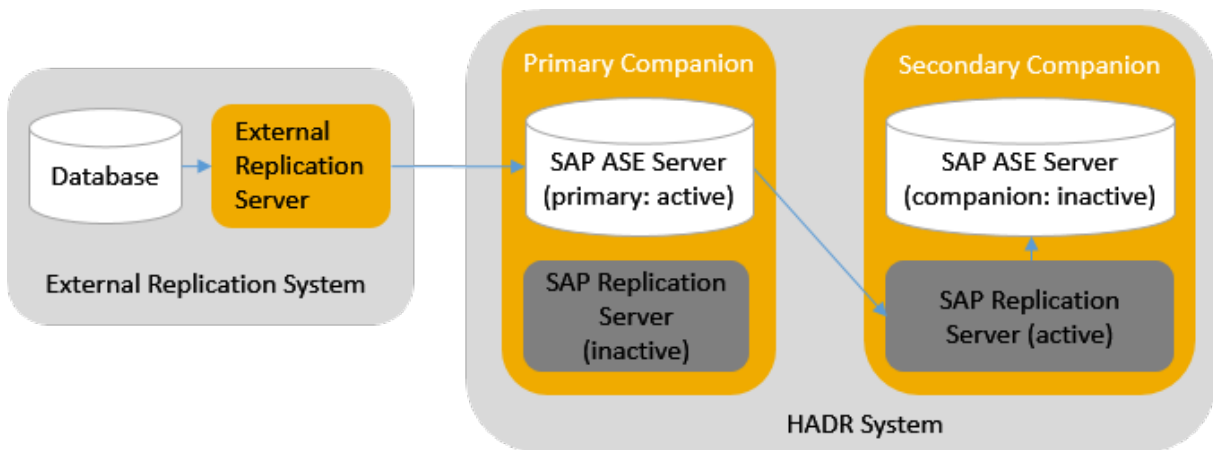


Figure 1: Replicating Data into an HADR System

The Data Server Interface (DSI) from the external Replication Server is perceived as a regular user in the primary companion, and is redirected to the current active companion during failover.

Replicating Data out from the HADR System into Replication Server

In this architecture, the active SAP Replication Server (the server on the secondary companion) replicates data to the external replication system by using the embedded SPQ Agent.

Even though the active Replication Server is exposed to external replication system for configuring the SPQ Agent, the details of the HADR system such as routes and subscriptions are all hidden from the external system. After a failover, drain the previous, active SPQ Agent to the external SAP Replication Server before starting a new active SPQ Agent.

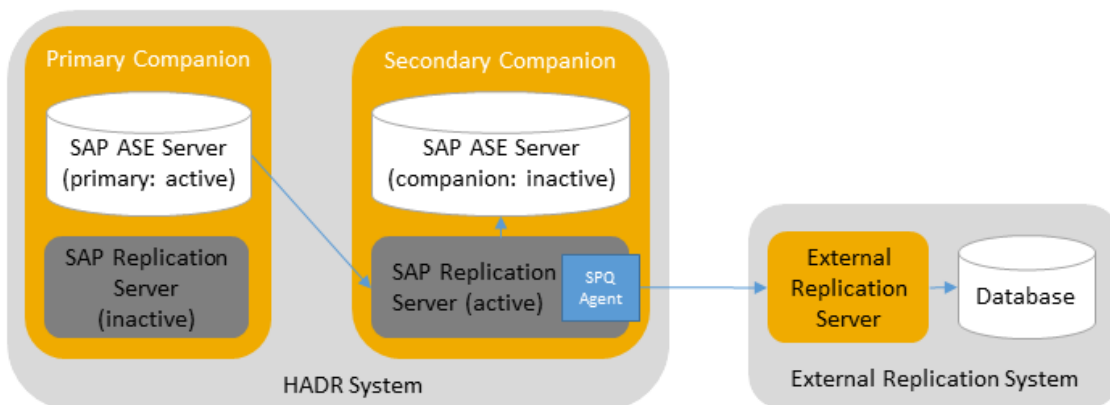


Figure 2: Replicating Data from an HADR System to an External System

6.3 Configuring HADR with an External System

The steps for configuring HADR with an external system are different depending on if you are replicating data from the external system to the HADR or from the HADR to an external system.

6.3.1 Requirements for Integrating an HADR System with External Replication

External replication includes these requirements and limitations.

- To support external replication:
 - SAP ASE must be version 16.0 SP02 PL05 or later.
 - SAP Replication Server must be version 15.7.1 SP305 or later.
 - SAP Replication Server on the primary and secondary companions must use the same platform.
- An HADR system database with external replication cannot participate in warm standby replication. Before you convert a system that includes warm standby databases to an HADR system, you must remove the warm standby implementation before you configure the system for HADR (without multi-site availability (MSA)). Additional standby servers outside of an HADR system must use MSA, and cannot be implemented with a warm standby using external replication with the HADR system as the active or standby server.

Requirements for replicating data into HADR:

- The data server interface (DSI) from the external Replication Server to the primary companion in HADR must enable `dsi_replication_ddl`. If it is disabled (set to off), the DSI applies DDLs using its original user, and if the original user has `sa` or `allow HADR login` permission, the connection may not be redirected to the primary companion.
With `dsi_replication_ddl` enabled, the maintenance user executes the DDL and login redirection ensures that the DDL is applied to the primary companion. Additionally, the DSI does not issue `set replication off` when applying DDLs as the maintenance user when it is connected to the primary companion, which is important if your site requires bidirectional replication for external replication. DDLs issued by the maintenance user are filtered out so that DDLs do not return to the site they were originally executed.

i Note

Using RMA to set up the HADR system enables `dsi_replication_ddl` in the HADR system automatically.

The maintenance user from the external Replication Server must have `set proxy authorization` permissions to replicate DDL with `dsi_replication_ddl` since the DSI connection uses `set user` to perform DDL operations.

- The maintenance user for DSI connections to HADR cannot use the same name as the maintenance user in the HADR system (for example, `DR_maint`, which is the default HADR maintenance user name).
- The maintenance user for DSI connections into HADR cannot be named `DR_admin`. HADR uses this name for its DR administrator.

- The external maintenance user cannot be aliased to `dbo`. The `rs_lastcommit` (and other `rs_*` system tables and system procedures) owned by `dbo` are used by HADR, so the external maintenance user must be different from `dbo` and have its own `rs_*` tables.
- The external maintenance user cannot have, or have inherited, the `sa_role`, `sso_role`, `allow HADR login` permission, or the `HADR admin role`.
- Granular permission must be enabled in the HADR system (requires the SECDIR license option if you are not running SAP ASE PE) so the external maintenance user can be granted `identity_insert any table` permission if the external system needs to replicate identity columns into HADR.

Requirements for replicating data out of HADR:

- After a failover, the SPQ Agent associated with the newly active Replication Server must wait for SPQ Agent from the previously active Replication Server to drain its data before it can start replicating. If a second failover occurs before the drain is completed during this wait period, external replication may suffer a data loss and must be rematerialized or resolved using Data Assurance (or using manual methods) in situations where rematerialization is not possible.
- Transactions made by `DR_admin` are not replicated to the external system. The external system cannot use the `DR_admin` user to create, drop, or alter subscriptions, and any marker issued by `DR_admin` is replicated only within the HADR system.
- The version of the external SAP Replication Server must be the same as or later than the version of the two SAP Replication Server servers and the two SAP ASE servers within the HADR system. Otherwise, replication out of HADR fails.
- If you set up HADR with a login other than the default HADR administrator user (`DR_admin`), use the `configure replication server` or `alter connection` command with the `cap_filter_dr_admin_name` parameter to change the default administrator user name to your own login name. This ensures that the external Replication Server can filter out all commands issued under your login name. Otherwise, data inconsistency may happen, causing DSI to go down.

Limitations for this release of HADR:

- This release of HADR replicates the effects of stored procedures, but does not replicate the execution of stored procedures.
- When replicating data into an HADR system, dropping the connection to the HADR system does not delete the system tables created by the external maintenance user in the primary companion automatically.

6.3.2 Configuring Replication In for an HADR System

You can use Replication Server to create a connection to the external database as a primary source.

Prerequisites

Replication Server is installed, and is managing the external database. See the Replication Server installation guide for your platform if you have not yet installed Replication Server.

Procedure

1. Add an entry to the Replication Server interfaces file for the primary and companion SAP ASE server in the HADR system. This example includes entries for the primary and companion servers:

```
SFSAP1
  query tcp ether SFMACHINE1 5000
  query tcp ether SJMACHINE2 5000
```

2. Create a maintenance user in the primary SAP ASE server in the HADR system. The maintenance user for DSI connections to HADR cannot use the same name as the maintenance user in the HADR system, or the name of the default user who executes RMA commands (DR_maint and DR_admin).
 - a. Create the login. This example creates the pubs2_maint login:

```
create login pubs2_maint with password 'SuperSecret123'
exempt inactive lock true
```

- b. Add the pubs2_maint user to the pubs2 database:

```
use pubs2
go
sp_adduser pubs2_maint
go
```

- c. To ensure this login is always redirected to the active companion, it cannot have the sa_role, sso_role, or the allow HADR login and HADR admin role permissions. Use sp_displayroles to display the role information. For example:

```
sp_displayroles pubs2_maint
Role Name
-----
sa_role
replication_role
replication_maint_role_gp
sap_maint_user_role
```

Use the revoke command to remove roles and permissions. For example:

```
revoke role sa_role from pubs2_maint
```

- d. Create a role for external replication. In this example, the role is named external_replication_role:

```
create role external_replication_role
```

- e. Grant set proxy permissions to the external_replication_role, but restrict it from switching to the sa_role, sso_role, and mon_role roles:

```
grant set proxy to external_replication_role
restrict role sa_role, sso_role, mon_role
```

- f. Grant the replication_role to the external_replication_role:

```
grant role replication_role to external_replication_role
```


- g. Alter the `pubs2_maint` login to automatically activate the `external_replication_role` role on login:

```
alter login pubs2_maint
add auto activated roles external_replication_role
```

- h. Since the maintenance user is not aliased as the database owner (dbo), manually grant it the following permissions on the tables:

Note

You need the SECDIRS license to grant the following permissions to the maintenance user.

```
grant delete any table to replication_role
go
grant create any table to replication_role
go
grant create any procedure to replication_role
go
grant execute any function to replication_role
go
grant execute any procedure to replication_role
go
grant identity_insert any table to replication_role
go
grant identity_update any table to replication_role
go
grant insert any table to replication_role
go
grant select any system catalog to replication_role
go
grant select any table to replication_role
go
grant truncate any table to replication_role
go
grant update any table to replication_role
go
```

3. Issue this command from the Replication Server that replicates into the HADR system:

```
create connection to <SAP_ASE_server.dbname>
set error class to rs_sqlserver_error_class
set function string class to rs_sqlserver_function_class
set username to <maintenance_user_name>
set password to <maintenance_user_password>
connect to hadr
```

For example:

```
create connection to SFSAP1.pubs2
set error class to rs_sqlserver_error_class
set function string class to rs_sqlserver_function_class
set username to pubs2_maint
set password to SuperSecret123
connect to hadr
```

4. Perform the following to replicate data from the external database to the primary companion's database:
- Create table- or database-level replication definitions (`repdefs`) using the primary as the external database. This example creates the `pubs_rep` replication definition for the `pubs2.publishers` table:

```
create replication definition pubs_rep
with primary at LASAP1.pubs2
```

```
with all tables named 'publishers'  
(pub_id char(4),  
 pub_name varchar(40),  
 city varchar(20),  
 state varchar(2))  
primary key (pub_id)
```

- Create subscriptions for the replication definition with replicates at the HADR database. This example creates the `pubs_sub` subscription:

```
Create subscription pubs_sub  
for pubs_rep  
with replicate at SFSAP1.pubs2
```

You can also replicate stored procedures and SQL statements into the HADR system. The procedure of replicating them in stream replication is the same as the procedure for Log Transfer Language (LTL) replication. See *Replication Server Administration Guide Volume 1 > Manage Replicated Functions > Use Replicated Functions* and *Replication Server Administration Guide Volume 2 > Performance Tuning > SQL Statement Replication* for details.

When you replicate a request stored procedure and the `select into` statement, grant the `execute` privilege on the following stored procedures to the original user who executes them in the primary companion's database:

- `<maintenance_user>.rs_update_last_commit`
- `<maintenance_user>.rs_get_lastcommit`
- `<maintenance_user>.rs_get_thread_seq`
- `<maintenance_user>.rs_initialize_threads`
- `<maintenance_user>.rs_syncup_lastcommit`
- `<maintenance_user>.rs_ticket_report`
- `<maintenance_user>.rs_update_threads`

The `<maintenance_user>` is the maintenance user for DSI connections to HADR, and not the maintenance user in the HADR system.

Next Steps

To configure the system to replicate bidirectionally, set up replication both into and out of the HADR cluster.

6.3.3 Configuring Replication Out From an HADR System

When you create a primary connection from an HADR system, SPQ Agent is enabled and configured on the active Replication Server. The SPQ Agent on the active Replication Server reads data from its SPQ and sends it to external system.

Prerequisites

Replication Server is installed, and is managing the external database. See the Replication Server installation guide for your platform if you have not yet installed Replication Server.

Procedure

1. Create a connection to the external SAP ASE database as the replicate target.
2. Create a maintenance user in the primary SAP ASE server in the HADR system. The maintenance user cannot be named `DR_maint` or `DR_admin`, which are default HADR maintenance user names.
 - a. Create the login. This example creates the `pubs2_maint` login:

```
create login pubs2_maint with password 'SuperSecret123'  
exempt inactive lock true
```

- b. To ensure this login is always redirected to the active companion, it cannot have the `sa_role`, `sso_role`, or the `allow HADR login` and `HADR admin role` permissions. Use `sp_displayroles` to display the role information. For example:

```
sp_displayroles pubs2_maint  
Role Name  
-----  
sa_role  
replication_role  
replication_maint_role_gp  
sap_maint_user_role
```

Use the `revoke` command to remove roles and permissions. For example:

```
revoke role sa_role from pubs2_maint
```

- c. Create a role for external replication. In this example, the role is named `external_replication_role`:

```
create role external_replication_role
```

- d. Grant `set proxy` permissions to the `external_replication_role`, but restrict it from switching to the `sa_role`, `sso_role`, and `mon_role` roles:

```
grant set proxy to external_replication_role  
restrict role sa_role, sso_role, mon_role
```

- e. Grant the `replication_role` to the `external_replication_role`:

```
grant role replication_role to external_replication_role
```

- f. Grant the `external_replication_role` to the `pubs2_maint` login.

```
grant role external_replication_role to pubs2_maint
```

- g. Alter the `pubs2_maint` login to automatically activate the `external_replication_role` role on login.

```
alter login pubs2_maint
add auto activated roles external_replication_role
```

- h. Grant permissions to the maintenance user on the tables; this user is aliased as the database owner. Issue commands similar to the following to grant the correct permissions to the maintenance user on your system.

Note

You need the SECDIRS license to grant the following permissions to the maintenance user.

```
grant delete any table to replication_role
go
grant execute any function to replication_role
go
grant execute any procedure to replication_role
go
grant identity_insert any table to replication_role
go
grant identity_update any table to replication_role
go
grant insert any table to replication_role
go
grant select any system catalog to replication_role
go
grant select any table to replication_role
go
grant truncate any table to replication_role
go
grant update any table to replication_role
go
```

3. Create the same maintenance user on the active and standby HADR Replication Server and grant `manage spq_agent` privileges to this user:

```
create user pubs2_maint set password SuperSecret123
go
grant manage spq_agent to pubs2_maint
go
```

4. Create the SPQ Agent user on the external Replication Server, and grant the `connect source` role to this user. HADR Replication Server uses this user to connect to the external Replication Server. For example:

```
create user spq_ra_user set password Sybase123
go
grant connect source to spq_ra_user
go
```

5. Add an entry to the external Replication Server interfaces file for the primary and companion SAP ASE servers in the HADR system. This example includes entries for the primary and companion servers:

```
SFSAP1
  query tcp ether SFMACHINE1 5000
  query tcp ether SJMACHINE2 5000
```

6. Issue this command from the external Replication Server to create a primary connection from HADR system:

```
create connection to <SAP_ASE_server.dbname>
set error class to rs_sqlserver_error_class
set function string class to rs_sqlserver_function_class
set username to <maintenance_user_name>
set password to <maintenance_user_password>
with primary only
connect to hadr
set spq_agent_username to <spq_agent_username>
set spq_agent_password to <password>
```

For example:

```
create connection to SFSAP1.pubs2
set error class to rs_sqlserver_error_class
set function string class to rs_sqlserver_function_class
set username to pubs2_maint
set password to SuperSecret123
with primary only
connect to hadr
set spq_agent_username to spq_ra_user
set spq_agent_password to krayzkat
```

i Note

When you issue `create connection`:

- If you specify `with primary only` parameter, the data server interface (DSI) is disabled, and outbound queues and their stable queue management (SQM) and DSI threads are not started.
- The command configures an SPQ Agent in the active Replication Server in the HADR system that reads data from the SPQ for the intended database, and forwards the data to the external Replication Server.
- When the command completes, the SPQ Agent connects to the external Replication Server.

7. Perform the following to replicate data from the HADR database to the external database:
 - Create table- or database-level replication definitions (`repdefs`) with the primary at the HADR database. This example creates the `pubs_rep` replication definition for the `pubs2.publishers` table:

```
create replication definition pubs_rep
with primary at SFSAP1.pubs2
with all tables named 'publishers'
(pub_id char(4),
pub_name varchar(40),
city varchar(20),
state varchar(2))
primary key (pub_id)
```

- Create subscriptions for the replication definition with replicates at the replicate database. This example creates the `pubs_sub` subscription:

```
Create subscription pubs_sub
```

```
for pubs_rep  
with replicate at LASAP1.pubs2
```

You can also replicate stored procedures and SQL statement from the HADR system to the external database. The procedure of replicating them in stream replication is the same as the procedure for LTL replication. See *Replication Server Administration Guide Volume 1 > Manage Replicated Functions > Use Replicated Functions* and *Replication Server Administration Guide Volume 2 > Performance Tuning > SQL Statement Replication* for details.

Next Steps

i Note

To configure the system to replicate bidirectionally, set up replication both in and out to the HADR cluster.

6.3.4 Migrating SAP ASE to an HADR System

Adding HADR to an primary SAP ASE with replication should not damage any existing replication. Although it may pause replication for a short time, it should not cause data loss or duplicates in the existing replication system.

6.3.4.1 Migrating the Primary SAP ASE to an HADR System

This task assumes a system that includes a Replication Server that is replicating data between two SAP ASE servers. The data is replicating from the primary to the target server. This task migrates the primary SAP ASE to an HADR system.

Prerequisites

These steps describe the migration from a dataserver `DS_P` that currently replicates data to target dataserver `DS_T` using Replication Server `EX_RS`. These steps add high availability capability to dataserver `DS_P` by migrating it to an HADR system. Since `DS_P` already replicates data to dataserver `DS_T`, it already has a connection to the Replication Server `EX_RS`, and assume the maintenance user for this connection is `pubs2_maint`, and that it has the replication definitions and subscriptions in Replication Server `EX_RS`.

Procedure

1. Shut down all applications connected to the primary SAP ASE (DS_P in this example). Do not shut down SAP ASE and Replication Server.
2. Drain the primary SAP ASE logs. Issue `rs_ticket` on the primary SAP ASE (DS_P) and then check it on the target SAP ASE (DS_T in this example):

```
exec rs_ticket 'my_ticket'
```

When the ticket reaches the target SAP ASE (DS_T), the logs are drained. See [Checking Latency with rs_ticket \[page 344\]](#).

3. Stop the Replication Agent from the primary SAP ASE (DS_P). Log in using `isql` and issue a command similar to (this stops the Replication Agent running on the `pubs2` database):

```
use pubs2
go
sp_stop_rep_agent pubs2
go
```

4. Check the maintenance user roles. It cannot have the `sa_role` or `sso_role` roles, or have the `allow HADR login` permission.

Use `rs_helpuser` to check if the maintenance user is aliased to `dbo`. For example:

```
rs_helpuser pubs2_maint
```

Use `sp_displayroles` to display the role information. For example, this checks the `pubs2_maint` role:

```
sp_displayroles pubs2_maint
Role Name
-----
sa_role
replication_role
replication_maint_role_gp
sap_maint_user_role
```

Use the `revoke` command to remove roles and permissions. For example:

```
revoke role sa_role from pubs2_maint
```

5. Use the SAP GUI or the `setuphadr` utility to migrate the primary SAP ASE (DS_P) to an HADR system. See [Installing HADR with an Existing System \[page 86\]](#).
6. Add an entry to the external Replication Server (EX_RS in this example) interfaces file for the companion SAP ASE server.

```
DS_P
  query tcp ether SFMACHINE1 5000
  query tcp ether SJMACHINE2 5000
```

7. Create the external maintenance user (`pubs2_maint` in this example) on the primary and standby Replication Servers in the HADR system, and grant this user `manage spq_agent` permissions.

```
create user pubs2_maint set password <password>
go
grant manage spq_agent to pubs2_maint
go
```

i Note

You need to use the password which is created when creating the connection between the primary SAP ASE and the external Replication Server.

8. Log in external Replication Server and create the SPQ Agent user and grant this user the `connect source` role. The Replication Servers use this user to connect to external Replication Server.

```
create user spq_ra_user set password krazykat
go
grant connect source to spq_ra_user
go
```

9. Issue this command on the external Replication Server (`EX_RS`) to enable replication to the external database (`pubs2` in this example):

```
alter connection to DS_P.pubs2
connect to hadr
set spq_agent_username to spq_ra_user
set spq_agent_password to krazykat
```

Results

When you finish migrating the primary SAP ASE server to an HADR system, data replication — including the replication of stored procedures and SQL statements, if available before the migration — continues from the HADR system to the replicate SAP ASE server.

6.3.4.2 Migrating the Target SAP ASE to an HADR System

This task assumes a system that includes a Replication Server that is replicating data between two SAP ASE servers. The data is replicating from the primary to the target server. This task migrates the target SAP ASE to an HADR system.

Prerequisites

These steps describe the migration from a dataserver `DS_P` that currently replicates data to target dataserver `DS_T` using Replication Server `EX_RS`. These steps add high availability capability to dataserver `DS_T` by migrating it to an HADR system. Since `DS_P` already replicates data to dataserver `DS_T`, it already has a connection to the Replication Server `EX_RS`, and assume the maintenance user for this connection is `pubs2_maint`, and that it has the replication definitions and subscriptions in Replication Server `EX_RS`.

Procedure

1. Shut down all applications.
2. Drain the primary SAP ASE logs. Issue `rs_ticket` on the primary SAP ASE (`DS_P` in this example) and then check it on the target SAP ASE (`DS_T` in this example):

```
exec rs_ticket 'my_ticket'
```

When the ticket reaches the target SAP ASE (`DS_T`), the logs are drained. See [Checking Latency with `rs_ticket` \[page 344\]](#).

3. Stop the Replication Agent from the primary SAP ASE (`DS_P`). Log in to the primary SAP ASE using `isql` and issue a command similar to (this stops the Replication Agent running on the `pubs2` database):

```
use pubs2
go
sp_stop_rep_agent pubs2
go
```

4. Suspend the DSI connection to the replicate database (`pubs2` in this example):

```
suspend connection to DS_T.pubs2
```

5. Log in as user `sa`, and drop the `pubs2_maint` alias as `dbo` in replicate database, and add it as a user `pubs2_maint` (this example assumes you are using the `pubs2` database):

```
use master
go
grant set proxy to pubs2_maint
go
use pubs2
go
sp_dropalias pubs2_maint
go
sp_adduser pubs2_maint
go
grant all to pubs2_maint
go
```

6. Load and run the `rs_install_replicate.sql` file as `pubs2_maint` to create replication system tables and procedures in the replicate database, such as `pubs2_maint.rs_lastcommit`, `pubs2_maint.rs_threads`, or `rs_update_lastcommit`.
7. Log in using `isql` and issue these commands as `pubs2_maint` in the replicate database to copy the data from `dbo.rs_lastcommit` into `pubs2_maint.rs_lastcommit`:

```
insert into rs_lastcommit select * from dbo.rs_lastcommit
go
insert into rs_threads select * from dbo.rs_threads
go
insert into rs_mat_status select * from dbo.rs_mat_status
go
insert into rs_dbversion select * from dbo.rs_dbversion
go
```

8. Use the SAP GUI or the `setuphadr` utility to migrate the target server to an HADR system. See [Installing HADR with an Existing System \[page 86\]](#).

9. Add an entry to the external Replication Server (EX_RS in this example) interfaces file for the companion SAP ASE server:

```
DS_T
  query tcp ether SFMACHINE1 5000
  query tcp ether SJMACHINE2 5000
```

10. Issue this command on the external Replication Server (EX_RS) to enable replication from the primary SAP ASE to the HADR system:

```
alter connection to <target_server>.<database>
connect to hadr
go
resume connection to <target_server>.<database>
go
```

For example:

```
alter connection to DS_T.pubs2
connect to hadr
go
resume connection to DS_T.pubs2
go
```

11. Start the Replication Agent on the primary SAP ASE (DS_P).

```
use pubs2
go
sp_start_rep_agent pubs2
go
```

12. (Optional) Data replication — including stored procedures and SQL statements, if available before the migration — continues from the primary SAP ASE server to the HADR system. However, to replicate a request stored procedure and the `select into` statement from the primary SAP ASE server to the HADR system, grant the `execute` privilege on these stored procedures to the original user who executes them in the primary companion's database:

- `<maintenance_user>.rs_update_last_commit`
- `<maintenance_user>.rs_get_lastcommit`
- `<maintenance_user>.rs_get_thread_seq`
- `<maintenance_user>.rs_initialize_threads`
- `<maintenance_user>.rs_syncup_lastcommit`
- `<maintenance_user>.rs_ticket_report`
- `<maintenance_user>.rs_update_threads`

The `<maintenance_user>` is the maintenance user for DSI connections to HADR, and not the maintenance user in the HADR system. For example, `<maintenance_user>` used in this task would be `pubs2_maint`.

6.4 Removing Replication

Removing external replication includes dropping all subscriptions, replication definitions (repdefs), and connections from the external Replication Server.

6.4.1 Removing External Replication

Dropping the connections also disables the SPQ Agent at the active HADR Replication Server in systems that replicate data out from an HADR system.

You must drop the subscriptions and replication definitions before you drop the connections. The SPQ Agent on the active HADR Replication Server is shutdown when you drop the connection for replicating data out.

To remove the external replication:

1. Drop the subscriptions and replication definitions for the external database. For example, this removes the subscriptions and replication definitions for the `pubs2` database in the external data server, `LASAP1`:

```
drop subscription pubs2_sub
  for database replication definition pubs2_rep
  with primary at SFSAP1.pubs2
  with replicate at LASAP1.pubs2
  without purge
```

This drops a replication definition named `pubs2_def` and any function strings that exist for it:

```
drop replication definition pubs2_def
```

2. Drop the connections. This example drops the connections to the `pubs2` database in the HADR system:

```
drop connection to SFSAP1.pubs2
```

6.4.2 Migrating an HADR System with Replication Out to Standalone SAP ASE

This task tears down an HADR system which replicates data to an external system, but maintains the replication.

Context

In an HADR system which replicates data out to external replication system, when the secondary companion is down, you can migrate the HADR system to a standalone SAP ASE and connect the external system with the Replication Agent thread on the standalone SAP ASE.

Procedure

1. Log into RMA and execute following:

```
sap_teardown
```

2. Remove companion SAP ASE host and port information from the interface file of the external Replication Server.
3. Run the `rs_install_primary.sql` script manually (located in `$SYBASE/DM/REP_16_0/scripts`) as `DR_admin` to install the `rs_ticket` command.
4. Log into primary SAP ASE and reconfigure the Replication Agent:

```
sp_config_rep_agent <database_name>, enable, <server_name>, <ext_main_user>,  
<password>  
go  
sp_config_rep_agent <database_name>, rs_servername,  
<external_rep_machine_name:port_number>  
go  
sp_config_rep_agent <database_name>, 'connect dataserver',  
<primary_server_name>  
go  
sp_config_rep_agent <database_name>, 'stream replication', 'true'  
go  
sp_config_rep_agent <database_name>, 'stream mode', 'sync'  
go  
sp_stop_rep_agent <database_name>  
go  
sp_start_rep_agent <database_name>  
go
```

5. Drain the data from the primary SAP ASE. Issue `rs_ticket` on the primary SAP ASE and then check it on the target SAP ASE:

```
exec rs_ticket 'my_ticket'
```

When the ticket reaches the target SAP ASE, the data is drained.

6. Alter the connection from HADR to the external Replication Server:

```
alter connection <server_name>.<database_name> not connect to hadr
```

7. Shut down the secondary companion in the HADR system. Log in with `isql` and issue:

```
shutdown
```

8. Shut down the inactive Replication Server (running on the original primary site) in the HADR system. Log in with `isql` and issue:

```
shutdown
```

9. Use `sap_teardown` to ensure that Replication Server is shut down. If Replication Server is not shut down, use the `kill` command to terminate the process.

10. Suspend and resume log transfer:

```
suspend log transfer from <primary_ASE.database>  
go  
suspend capture <primary_ASE.database>  
go  
resume capture <primary_ASE.database>
```

```
go
resume log transfer from <primary_ASE.database>
go
```

6.4.3 Migrating an HADR System with Replication In to a Standalone SAP ASE

This task tears down an HADR system into which an external system replicates data, but maintains the replication.

Procedure

1. Log into the RMA you are tearing down and execute:

```
sap_teardown
```

2. Remove companion SAP ASE host and port information from the interface file.
3. Alter the connection to no longer connect to the HADR system:

```
alter connection <server_name>.<database_name> not connect to hadr
```

4. Restart the DSI connection to the original HADR system. For example:

```
resume connection to <primary_ASE.database>
go
```

5. Shut down the secondary companion in the HADR system. Log in with `isql` and issue:

```
shutdown
```

6. Use `sap_teardown` to ensure that Replication Server is shut down. If Replication Server is not shut down, use the `kill` command to terminate the process.

6.5 Managing an HADR System with External Replication

Manage an external replication system by disabling or enabling replication, performing failover, monitoring the system, and so on.

6.5.1 Disabling and Enabling Replication to an External Replication Server

Disabling replication to an external Replication Server requires that RMA stops its SPQ Agent and disables the secondary truncation point.

Issue this command to disable replication from a specific database to the external Replication Server:

```
sap_disable_external_replication <dbname>
```

See [sap_disable_external_replication \[page 458\]](#).

This command stops the SPQ Agent in the active Replication Server and disables its second truncation point in SPQ so that it can be truncated properly. However, this command may result in data loss.

To enable replication from a specific database to the external Replication Server, enter this at the command line:

```
sap_enable_external_replication <dbname>
```

See [sap_enable_external_replication \[page 465\]](#).

You must manually rematerialize the databases after enabling external replication. See [Rematerializing Databases for External Replication \[page 262\]](#)

i Note

This topic only talks about disabling and enabling replication paths to the external Replication Server, for disabling and enabling replication paths inside an HADR system, refer to [Suspending, Resuming, Enabling, and Disabling Databases \[page 314\]](#)

6.5.2 Rematerializing Databases for External Replication

Under certain circumstances, databases may lose synchronization in the external replication system. You need to rematerialize the databases to resynchronize data.

This section describes how to rematerialize databases for the following scenarios:

- Replicating out from an HADR system
- Replicating in to an HADR system

To rematerialize databases inside an HADR system, see [Materializing and Rematerializing Databases \[page 306\]](#).

6.5.2.1 Rematerializing Databases when Replicating Out from an HADR System

When replicating data from an HADR system to the external replication system, perform the rematerialization steps manually to rematerialize databases in the external replication system.

Prerequisites

If you are rematerializing the databases from an HADR system to another HADR system, you need to shut down all applications before rematerializing the databases. Do not start the applications until the rematerialization is complete.

Context

You can use this method to rematerialize data only from a single database in the HADR system to a single database in the external replication system.

When HADR databases are replicating data to an external system, disabling the replication path from the primary SAP ASE server to the active SAP Replication Server disrupts the process, and also requires you to rematerialize the external databases.

Procedure

1. Log in to RMA on the active Replication Server in the HADR system and disable the replication from the databases to the external replication environment. To disable the replication for one specific database, specify the `<database_name>` in the command, to disable the replication for all databases, execute the command without the `<database_name>` variable:

```
sap_disable_external_replication [<database_name>]
go
```

2. Log in to RMA on the active Replication Server in the HADR system and enable the replication from the databases to the external replication environment. To enable the replication for one specific database, specify the `<database_name>` in the command, to enable the replication for all databases, execute the command without the `<database_name>` variable:

```
sap_enable_external_replication [<database_name>]
go
```

3. (Optional) Log in to RMA on the active Replication Server in the HADR system and check if the paths inside the HADR system are still activated:

```
sap_status path
go
```

4. Manually rematerialize the replication path from a specific database to the external replication system. If you want to rematerialize all external databases from the HADR system, repeat the configurations under this step to rematerialize each database one by one.

- a. (Optional) If the external database is an HADR system, log in to the RMA on the primary server of the HADR system in the external environment, then disable the replication path from the primary server to the companion server:

```
sap_disable_replication <primary_host_logical_name>,  
<companion_host_logical_name>, <database_name>  
go
```

- b. Log in to the external Replication Server and suspend log transfer from the HADR system.

```
suspend log transfer from <data_server_name>.<database_name>  
go
```

- c. Log in to the external Replication Server and hibernate on the external Replication Server:

```
sysadmin hibernate_on  
go
```

→ Tip

You can choose to run the `sysadmin sqm_purge_queue` command to purge queues, without necessarily hibernating on the Replication Server. Instead, you can suspend the appropriate modules in the Replication Server, and then purge queues as usual. Running `sysadmin sqm_purge_queue` with the `[, check_only]` parameter facilitates this scenario, as it checks and reports if the appropriate modules were suspended successfully (it **does not** purge queues), thus enabling you to make an informed decision before purging queues. Note that you can continue to purge queues like you did before – by hibernating on the Replication Server. For more information, see the *Usage* section under *SAP Replication Server Reference Manual > SAP Replication Server Commands > sysadmin sqm_purge_queue*.

- d. Log in to the external Replication Server and purge queues:

```
sysadmin spq_purge_queue, <data_server_name>, <database_name>  
go  
sysadmin sqm_purge_queue, <q_number>, <q_type>  
go
```

i Note

Use `sysadmin sqm_purge_queue` to purge both the inbound queue and the outbound queue. To purge the inbound queue, set `<q_type>` to 1. To purge the outbound queue, set `<q_type>` to 0.

- e. Log in to the external Replication Server and hibernate off the external Replication Server:

```
sysadmin hibernate_off  
go
```

- f. Resume the connection from the HADR system to the external Replication Server and from the external Replication Server to the external database:

```
resume log transfer from <data_server_name>.<database_name>  
go
```

```
resume connection to <data_server_name>.<database_name>
```



```
go
```

- g. Without purging, drop the existing subscriptions in the external Replication Server:

```
drop subscription <sub_name>
for {<table_rep_def> | database replication definition <db_repdef>}
with primary at <data_server_name>.<database_name>
with replicate at <data_server_name>.<database_name>
without purge
go
```

- h. Create a temporary user (<temp_remater_maint_user>) on the primary SAP ASE server in the HADR system, and grant it replication role and all permissions. This automatically creates the user on the companion SAP ASE server. The temporary user is used to define the subscription on the external Replication Server.

```
use master
go
create login <temp_remater_maint_user> with password <password>
go
grant role replication_role to <temp_remater_maint_user>
go
use <dbname>
go
sp_adduser <temp_remater_maint_user>
go
grant all to <temp_remater_maint_user>
go
```

- i. Create the temporary user (<temp_remater_maint_user>) on the external SAP Replication Server and grant it sa role.

```
create user <temp_remater_maint_user> set password <password>
go
grant sa to <temp_remater_maint_user>
go
```

- j. Log in to the external Replication Server as <temp_remater_maint_user> and define the subscription. The subscription is used to dump the database.

```
define subscription <sub_name>
for {<table_rep_def> | database replication definition <db_repdef>}
with primary at <data_server_name>.<database_name>
with replicate at <data_server_name>.<database_name>
use dump marker <label_name>
go
```

- k. Log in to the primary server in the HADR system as an sa user and dump the database.

```
dump database <database_name> to "<file_path_and_name>" at
<backup_server_name> with label = '<label_name>'
go
```

- l. Log in to the external Replication Server to check that the subscription is valid:

```
check subscription <sub_name>
for database replication definition <db_repdef>
with primary at <data_server_name>.<database_name>
with replicate at <data_server_name>.<database_name>
go
```

- m. In the database to be rematerialized, check if any processes are running:

```
use master
```

```
go
select spid from sysprocesses where dbid = db_id ('<database_name>')
go
```

If yes, kill the processes:

```
use master
go
kill <spid>
go
```

- n. Load the database to the active data server in the external replication system with the sa user:

```
load database <database_name> from "<file_name>"
go
```

- o. Online the rematerialized database:

```
online database <database_name>
go
```

- p. Truncate the following dbo system tables on the external database:

```
use <database_name>
go
truncate table dbo.rs_lastcommit
go
truncate table dbo.rs_threads
go
truncate table dbo.rs_ticket_history
go
truncate table dbo.rs_mat_status
go
truncate table dbo.rs_dbversion
go
```

- q. (Optional) If you used two different external maintenance users to create connections from the HADR system to the external Replication Server and from the external Replication Server to the external database. On the external database, drop the external maintenance user for the HADR system, then add the external maintenance user for the external database and grant all permissions to it.

```
use <database_name>
go
sp_dropuser <ext_maint_user_a>
go
sp_adduser <ext_maint_user_b>
go
grant all to <ext_maint_user_b>
go
```

- r. (Optional) When the external database is an HADR system, create the system tables and procedures for the external maintenance user. Log in to the primary server of the HADR system in the external environment and load the `rs_install_replicate.sql` file:

```
isql -U<ext_maint_user> -P<password> -S<host_name>:<port_number>
-D<database_name> -i /SRS_Suite/rs_install_replicate.sql
```

- s. The connection between the external SAP Replication Server and the external database becomes suspended after you dump the database. Resume the connection:

```
resume connection to <data_server_name>.<database_name>
go
```

- t. (Optional) If the external database is an HADR system, you also need to rematerialize the database on the companion server in this HADR system. See [Materializing and Rematerializing Databases \[page 306\]](#) for more information.

6.5.2.2 Rematerializing the Database When Replicating from SAP ASE to an HADR System

When replicating from an SAP ASE server to an HADR system, perform the rematerialization steps manually to rematerialize all databases in the SAP ASE server.

Prerequisites

Shut down all applications before rematerializing the databases. Do not start the applications until the rematerialization is complete.

Procedure

1. Log in to RMA on the active Replication Server in the HADR system and disable the replication from the databases to the HADR system. To disable the replication for one specific database, specify the `<database_name>` in the command, to disable the replication for all databases, execute the command without the `<database_name>` variable:

```
sap_disable_replication [<database_name>]
go
```

2. Log in to RMA on the active Replication Server in the HADR system and enable the replication from the databases to the HADR system. To enable the replication for one specific database, specify the `<database_name>` in the command, to enable the replication for all databases, execute the command without the `<database_name>` variable:

```
sap_enable_replication [<database_name>]
go
```

3. Manually rematerialize the replication path from a specific database to the HADR system. If you want to rematerialize all databases in the SAP ASE server, repeat the configurations under this step to rematerialize each database one by one.
 - a. Log in to the RMA on the primary server of the HADR system and disable the replication path from the primary server to the companion server:

```
sap_disable_replication <primary_host_logical_name>,
<companion_host_logical_name>, <database_name>
go
```

- b. Log in to the external Replication Server and suspend log transfer from the database.

```
suspend log transfer from <data_server_name>.<database_name>
```

```
go
```

- c. Log in to the external Replication Server and hibernate on the external Replication Server:

```
sysadmin hibernate_on  
go
```

→ Tip

You can choose to run the `sysadmin sqm_purge_queue` command to purge queues, without necessarily hibernating on the Replication Server. Instead, you can suspend the appropriate modules in the Replication Server, and then purge queues as usual. Running `sysadmin sqm_purge_queue` with the `[, check_only]` parameter facilitates this scenario, as it checks and reports if the appropriate modules were suspended successfully (it **does not** purge queues), thus enabling you to make an informed decision before purging queues. Note that you can continue to purge queues like you did before – by hibernating on the Replication Server. For more information, see the *Usage* section under *SAP Replication Server Reference Manual > SAP Replication Server Commands > sysadmin sqm_purge_queue*.

- d. Log in to the external Replication Server and purge queues:

```
sysadmin spq_purge_queue, <data_server_name>, <database_name>  
go  
sysadmin sqm_purge_queue, <q_number>, <q_type>  
go
```

i Note

Use `sysadmin sqm_purge_queue` to purge both the inbound queue and the outbound queue. To purge the inbound queue, set `<q_type>` to 1. To purge the outbound queue, set `<q_type>` to 0.

- e. Log in to the external Replication Server and hibernate off the external Replication Server:

```
sysadmin hibernate_off  
go
```

- f. Resume the connection from the database to the external Replication Server and from the external Replication Server to the HADR system:

```
resume log transfer from <data_server_name>.<database_name>  
go
```

```
resume connection to <data_server_name>.<database_name>  
go
```

- g. Without purging, drop the existing subscriptions in the external Replication Server:

```
drop subscription <sub_name>  
for {<table_rep_def> | database replication definition <db_repdef>}  
with primary at <data_server_name>.<database_name>  
with replicate at <data_server_name>.<database_name>  
without purge  
go
```

- h. Create a temporary user (<temp_remater_maint_user>) on the SAP ASE server, and grant it replication role and all permissions. The temporary user is used to define the subscription on the external Replication Server.

```
use master
go
create login <temp_remater_maint_user> with password <password>
go
grant role replication_role to <temp_remater_maint_user>
go
use <dbname>
go
sp_adduser <temp_remater_maint_user>
go
grant all to <temp_remater_maint_user>
go
```

- i. Create the temporary user (<temp_remater_maint_user>) on the external SAP Replication Server and grant it sa role.

```
create user <temp_remater_maint_user> set password <password>
go
grant sa to <temp_remater_maint_user>
go
```

- j. Log in to the external Replication Server as <temp_remater_maint_user> and define the subscription. The subscription is used to dump the database.

```
define subscription <sub_name>
for {<table_rep_def> | database replication definition <db_repdef>}
with primary at <data_server_name>.<database_name>
with replicate at <data_server_name>.<database_name>
use dump marker <label_name>
go
```

- k. Log in to the SAP ASE server as an sa user and dump the database.

```
dump database <database_name> to "<file_path_and_name>" at
<backup_server_name> with label = '<label_name>'
go
```

- l. Log in to the external Replication Server to check that the subscription is valid:

```
check subscription <sub_name>
for database replication definition <db_repdef>
with primary at <data_server_name>.<database_name>
with replicate at <data_server_name>.<database_name>
go
```

- m. In the database to be rematerialized, check if any processes are running:

```
use master
go
select spid from sysprocesses where dbid = db_id ('<database_name>')
go
```

If yes, kill the processes:

```
use master
go
kill <spid>
go
```

- n. Load the database to the active data server in the HADR system with the sa user:

```
load database <database_name> from "<file_name>"
go
```

- o. Online the rematerialized database:

```
online database <database_name>
go
```

- p. Truncate the following dbo system tables on the primary data server in the HADR system:

```
use <database_name>
go
truncate table dbo.rs_lastcommit
go
truncate table dbo.rs_threads
go
truncate table dbo.rs_ticket_history
go
truncate table dbo.rs_mat_status
go
truncate table dbo.rs_dbversion
go
```

- q. (Optional) If you used two different external maintenance users to create connections from the SAP ASE to the external Replication Server and from the external Replication Server to the HADR system. In the HADR system, drop the external maintenance user for the SAP ASE, then add the external maintenance user for the HADR system to the database and grant all permissions to it.

```
use <database_name>
go
sp_dropuser <ext_maint_user_a>
go
sp_adduser <ext_maint_user_b>
go
grant all to <ext_maint_user_b>
go
```

- r. Create the system tables and procedures for the external maintenance user. Log in to the primary server of the HADR system and load the `rs_install_replicate.sql` file:

```
isql -U<ext_maint_user> -P<password> -S<host_name>:<port_number>
-D<database_name> -i /SRS_Suite/rs_install_replicate.sql
```

- s. The connection between the external SAP Replication Server and the HADR system becomes suspended after you dump the database. Resume the connection:

```
resume connection to <data_server_name>.<database_name>
go
```

- t. Rematerialize the database on the companion server in the HADR system. See [Materializing and Rematerializing Databases \[page 306\]](#) for more information.

6.5.3 Monitoring the HADR System Using the SPQ Agent

The SPQ Agent is a Replication Server component that reads the Simple Persistent Queue (SPQ), and forwards the messages to the external system, thus acting as a Replication Agent to the external system.

Use the `sap_status spq_agent` command to monitor the status of the SPQ Agent. The `sap_status spq_agent` command lets you monitor the following:

- The SPQ Agent path and thread states. It indicates if the path is configured for the SPQ Agent.
- The SPQ Agent backlog size.

6.5.3.1 Using the `sap_status spq_agent` Command to Monitor the HADR Environment

Run the `sap_status spq_agent` command in an active HADR system that has external replication configured. Internally, for each registered host, the RMA uses these two Replication Server commands, and merges the results:

- `admin disk_space, mb, spqra`
- `admin who, spqra`

Note

The `sap_status spq_agent` command provides basic information on the SPQ Agent, such as SPQ Agent state and its backlog size. To get additional details, use internal Replication Server commands such as `admin who, spqra`.

The following is the sample output from `sap_status spq_agent` command:

Sample Code

```
1> sap_status spq_agent
2> go
HOST          INFO                STATE      BACKLOG
-----
HA            110 PI2_PR_R2.PI2  ACTIVE     0
PR            108 PI2_HA_R1.PI2  INACTIVE   0
```

If `BACKLOG` displays `N/A`, and `STATE` displays `INACTIVE`, they indicate that the SPQ Agent on this path is not configured, and there is no functional external replication on this path.

6.5.4 Managing Failovers within HADR with External Replication

Failover switches activities to the standby node when the primary node is unavailable.

You can perform planned and unplanned failovers within an HADR cluster with external replication. The failover happens between the primary node and the standby node within the HADR system, and does not impact the external replication system.

When data is replicating from an external replication system to HADR, the failover process is the same as that of an independent HADR cluster. See [Planned and Unplanned Failovers \[page 325\]](#) for details.

The process differs when data is replicating from HADR to an external replication system. After you run the `sap_failover` command to perform the failover, run the `sap_failover_drain_to_er` command to make sure that all backlogs are drained from the primary node to the external replication system to achieve zero data loss. When this command has finished successfully, replication is established from the former standby node to the external replication system.

This chapter focuses on how to perform a failover when the data is replicating from an HADR cluster to an external replication system, with examples using the following values:

- Logical host name for primary node: PR
- Logical host name for standby node: HA
- Logical host name for external system: ER

6.5.4.1 Performing a Planned Failover

A planned failover to the standby node allows you to perform regular maintenance work and scheduled system upgrades on the primary node without affecting replication to the external system.

Prerequisites

Suspend customer application activities against the primary database to ensure a clean transition to the standby site, sparing client applications from reacting to the server downtime.

Procedure

1. Log in to the primary RMA and run a command similar to the following. This example uses a timeout of 120 seconds:

```
sap_failover PR, HA, 120
go
```

`sap_failover` is an asynchronous command and must complete before you perform the next step. Use the `sap_status` command to check the failover status, indicated in the example output in bold text.

```
sap_status
go
```

TASKNAME	TYPE	VALUE
Status	Start Time	Wed Sep 07 12:00:32 UTC 2016
Status	Elapsed Time	00:00:06
Failover	Task Name	Failover
Failover	Task State	Completed
Failover	Short Description	Failover makes the current standby ASE as the primary server.


```

Failover Long Description      Failover from source 'PR' to target 'HA' is
complete. The target may be unquiesced.
Failover Additional Info      Please run command 'sap_failover_drain_to_er
<timeout>' to complete the data drain for external system.
Failover Current Task Number  14
Failover Total Number of Tasks 14
Failover Task Start           Wed Sep 07 12:00:32 UTC 2016
Failover Task End             Wed Sep 07 12:00:38 UTC 2016
Failover Hostname             site0

```

When `sap_failover` has finished, the SAP ASE server on the former standby node becomes the new primary server. It is only activated after all transaction backlogs on the former primary node are drained to the former standby node (current primary). Client applications can connect to the former standby node to continue business activities.

2. Run the following command to check whether all transaction backlogs are drained from the former primary node to the external replication system. The example uses a timeout of 120 seconds:

```

sap_failover_drain_to_er 120
go

```

Note

You may also disable the replication from the HADR cluster to a database or all databases on the external system by executing `sap_failover_drain_to_er skip [<dbName>]`, but doing so causes the external replicate databases to be out of sync with the HADR cluster.

This is an asynchronous command. Use the `sap_status` command to check the progress:

```

sap_status
go

TASKNAME          TYPE
VALUE
-----
-----
Status            Start Time      Wed Sep 07 12:01:00 UTC
2016
Status            Elapsed Time
00:00:36

FailoverDrainToER Task Name      Failover drain to
ER.

FailoverDrainToER Task State
Completed

FailoverDrainToER Short Description  Failover drain to ER deactivate old
replication path and activate new replication path for external replication
system.
FailoverDrainToER Long Description  Failover drain to ER completed.
Update mark.
FailoverDrainToER Additional Info   Please run command
'sap_host_available' to complete the failover
operation.
FailoverDrainToER Current Task Number
5

FailoverDrainToER Total Number of Tasks
5

```

```
FailoverDrainToER Task Start      Wed Sep 07 12:01:00 UTC
2016
FailoverDrainToER Task End      Wed Sep 07 12:01:36 UTC
2016
FailoverDrainToER Hostname
site0
```

When `sap_failover_drain_to_er` has successfully finished, all backlogs are drained to the external system and replication is established from the new primary node to the external replication system.

- (Optional) If the Fault Manager is configured to restart SAP ASE, SAP Replication Server, and RMA, stop it before you perform any maintenance activity on the former primary node. From the `<install_directory>/FaultManager` directory, issue:

```
<Fault_Manager_install_dir>/FaultManager/bin/sybdbfm stop
```

Note

If you configured a floating IP in the Fault Manager (`ha/syb/support_floating_ip = 1`) and stopped the Fault Manager before issuing the `sap_failover` command, you need to move the floating IP manually. See SAP Note [2709394](#) for more information.

- When the former primary node is ready to rejoin the replication system, run a similar command in RMA:

```
sap_host_available PR
go
```

The system displays the following information when the command has finished successfully.

```
TASKNAME      TYPE      VALUE
-----
-----
HostAvailable Start Time      Wed Sep 07 12:02:06 UTC 2016
HostAvailable Elapsed Time    00:02:43
HostAvailable Task Name      HostAvailable
HostAvailable Task State    Completed
HostAvailable Short Description Resets the original source logical host
when it is available after failover.
HostAvailable Long Description Completed the reset process of logical
host 'PR' receiving replication from logical host 'HA'.
HostAvailable Current Task Number 10
HostAvailable Total Number of Tasks 10
HostAvailable Task Start      Wed Sep 07 12:02:06 UTC 2016
HostAvailable Task End      Wed Sep 07 12:04:49 UTC 2016
HostAvailable Hostname      site0
```

Replication from the former standby node to the former primary node is established.

- Verify replication path status.
 - To verify the replication paths within the HADR cluster, run `sap_status path` in RMA.

```
sap_status path
go
```

```
PATH      NAME      VALUE
INFO
-----
-----
```

```

Start Time 2016-09-07 12:21:13.446
Time command started executing.
Elapsed Time 00:00:00
Command execution time.
PR Hostname site0
Logical host name.
PR HADR Status Standby : Inactive
Identify the primary and standby sites.
PR Synchronization Mode Synchronous
The configured Synchronization Mode value.
PR Synchronization State Inactive
Synchronization Mode in which replication is currently operating.
PR Distribution Mode Remote
Configured value for the distribution mode replication model property.
PR Replication Server Status Active
The status of Replication Server.
HA Hostname HA
Logical host name.
HA HADR Status Primary : Active
Identify the primary and standby sites.
HA Synchronization Mode Synchronous
The configured Synchronization Mode value.
HA Synchronization State Synchronous
Synchronization Mode in which replication is currently operating.
HA Distribution Mode Remote
Configured value for the distribution mode replication model property.
HA Replication Server Status Active
The status of Replication Server.
PA.HA.ERP State Suspended
Path is suspended (Replication Agent Thread). Transactions are not being replicated.
PA.HA.ERP Latency Time Unknown
No latency information for database 'ERP'.
PA.HA.ERP Latency Unknown
No latency information for database 'ERP'.
PA.HA.ERP Commit Time Unknown
No last commit time for the database 'ERP'.
PA.HA.ERP Distribution Path HA
The path of Replication Server through which transactions travel.
PA.HA.ERP Drain Status Unknown
The drain status of the transaction logs of the primary database server.
PA.HA.ERP_1 State Suspended
Path is suspended (Replication Agent Thread). Transactions are not being replicated.
PA.HA.ERP_1 Latency Time Unknown
No latency information for database 'ERP_1'.
PA.HA.ERP_1 Latency Unknown
No latency information for database 'ERP_1'.
PA.HA.ERP_1 Commit Time Unknown
No last commit time for the database 'ERP_1'.
PA.HA.ERP_1 Distribution Path HA
The path of Replication Server through which transactions travel.
PA.HA.ERP_1 Drain Status Unknown
The drain status of the transaction logs of the primary database server.
PA.HA.master State Suspended
Path is suspended (Replication Agent Thread). Transactions are not being replicated.
PA.HA.master Latency Time Unknown
No latency information for database 'master'.
PA.HA.master Latency Unknown
No latency information for database 'master'.
PA.HA.master Commit Time Unknown
No last commit time for the database 'master'.
PA.HA.master Distribution Path HA
The path of Replication Server through which transactions travel.
PA.HA.master Drain Status Unknown
The drain status of the transaction logs of the primary database server.

```

```

HA.PR.ERP           State           Active
Path is active and replication can occur.
HA.PR.ERP           Latency Time           2016-09-07 12:20:12.340
Time latency last calculated
HA.PR.ERP           Latency                 205560
Latency (ms)
HA.PR.ERP           Commit Time            2016-09-07 12:20:12.346
Time last commit replicated
HA.PR.ERP           Distribution Path      PR
The path of Replication Server through which transactions travel.
HA.PR.ERP           Drain Status           Not Drained
The drain status of the transaction logs of the primary database server.
HA.PR.ERP_1        State           Active
Path is active and replication can occur.
HA.PR.ERP_1        Latency Time           2016-09-07 12:20:26.080
Time latency last calculated
HA.PR.ERP_1        Latency                 219430
Latency (ms)
HA.PR.ERP_1        Commit Time            2016-09-07 12:20:26.080
Time last commit replicated
HA.PR.ERP_1        Distribution Path      PR
The path of Replication Server through which transactions travel.
HA.PR.ERP_1        Drain Status           Not Drained
The drain status of the transaction logs of the primary database server.
HA.PR.master      State           Active
Path is active and replication can occur.
HA.PR.master       Latency Time           2016-09-07 12:20:19.506
Time latency last calculated
HA.PR.master       Latency                 213143
Latency (ms)
HA.PR.master       Commit Time            2016-09-07 12:20:19.506
Time last commit replicated
HA.PR.master       Distribution Path      PR
The path of Replication Server through which transactions travel.
HA.PR.master       Drain Status           Not Drained
The drain status of the transaction logs of the primary database server.

```

i Note

If the replication data load is low, the synchronization state may not update to *Synchronous* after you run the `sap_host_available` command to establish replication. To refresh its value, run the `sap_send_trace <primary_host_name>` command, then re-run the `sap_status path` command.

- To verify the replication to the external system, connect to SAP Replication Server on that site and run the `sysadmin path_check` command:

```

sysadmin path_check
go

Primary Database Replicate Database Status
-----
HADR_PR.ERP      HADR_2.ABC      Up

```

6. (Optional) If the Fault Manager is stopped, restart it:

```
<Fault_Manager_install_dir>/FaultManager/sydbfm_<CID>
```

6.5.4.2 Performing an Unplanned Failover

When the primary SAP ASE server is down or lost, perform an unplanned failover so client applications can continue to work on the SAP ASE server configured on standby node.

Context

Use the `sap_failover` command with the `unplanned` option to perform an unplanned failover from the primary node to the standby node.

i Note

If the Fault Manager is configured, do not set the `ha/syb/set_standby_available_after_failover` Fault Manager profile parameter to 1 (default is 0).

Procedure

1. If the Fault Manager is not configured, in the RMA, run the `sap_status path` command to verify that the synchronization state of the primary node is synchronous.

```

sap_status path
go

```

PATH	NAME	VALUE	
INFO			
	Start Time	2016-09-07 12:16:27.511	Time
command started	Elapsed Time	00:00:00	Command
execution time.	Hostname	mo-d34ece715.mo.sap.corp	Logical
PR	HADR Status	Unknown	Identify
the primary and standby sites.	Synchronization Mode	Synchronous	The
PR	configured Synchronization Mode value.		
PR	Synchronization State	Synchronous	
Synchronization Mode in which replication is currently operating.	Distribution Mode	Remote	
PR	Configured value for the distribution_mode replication model property.		
PR	Replication Server Status	Active	The
status of Replication Server.	Hostname	mo-b01b38e58.mo.sap.corp	Logical
HA	HADR Status	Standby : Inactive	Identify
the primary and standby sites.	Synchronization Mode	Synchronous	The
HA	configured Synchronization Mode value.		
HA	Synchronization State	Inactive	
Synchronization Mode in which replication is currently operating.	Distribution Mode	Remote	
HA	Configured value for the distribution_mode replication model property.		

```
HA          Replication Server Status Active          The
status of Replication Server.
...
```

A status of synchronous means there is no data loss between the primary and standby SAP ASE servers. After the failover, client applications can directly connect to the former standby server and resume business.

A status of asynchronous means there may be some data loss on the standby SAP ASE, in which case, make sure the data loss is acceptable before you perform an unplanned failover. Otherwise, failover is not recommended.

2. If the Fault Manager is not configured, enter an `sap_failover` command with the `unplanned` option to initiate the unplanned failover. The example uses a deactivation timeout of 120 seconds:

```
sap_failover PR, HA, 120 unplanned
go
```

Be sure to use the `unplanned` option. Or else the command fails and you get a warning message asking you to rerun the command with the `unplanned` option.

Use the `sap_status` command to check the progress, and proceed only after `sap_failover` has finished.

```
sap_status
go

TASKNAME      TYPE
VALUE

-----
-----
-----
Status        Start Time          Wed Sep 07 12:16:40 UTC
2016

Status        Elapsed Time
00:00:06

Failover      Task Name
Failover

Failover      Task State
Completed

Failover      Short Description    Failover makes the current standby ASE as
the primary server.
Failover      Long Description     Failover from source 'PR' to target 'HA' is
complete. The target may be unquiesced.
Failover      Additional Info      Please run command 'sap_failover_drain_to_er
<timeout>' to complete the data drain for external system.
Failover      Additional Info 2    The primary Replication Server 'site1:14635'
is configured for synchronization mode and was found running in
synchronization mode.
Failover      Current Task Number
12

Failover      Total Number of Tasks
12

Failover      Task Start          Wed Sep 07 12:16:40 UTC
2016
```

```

Failover Task End Wed Sep 07 12:16:46 UTC
2016

Failover Hostname
site0

```

When `sap_failover` has finished successfully, the SAP ASE server on the former standby node becomes the new primary server. It is activated only after all transaction backlogs on the former primary node are drained to the former standby node. Client applications can connect to the former standby node to continue activities.

3. Run the `sap_failover_drain_to_er` command to check whether all transaction backlogs are drained from the former primary node to the external replication system. This example uses a timeout of 120 seconds:

```

sap_failover_drain_to_er 120
go

```

Note

You may also disable replication from the HADR cluster to a database or all databases on the external system by executing `sap_failover_drain_to_er skip [<dbName>]`, but doing so causes the external replicate databases to be out of sync with the HADR cluster.

Use the `sap_status` command to check the progress:

```

sap_status
go

TASKNAME          TYPE
VALUE
-----
-----
Status            Start Time      Wed Sep 07 12:17:12 UTC
2016
Status            Elapsed Time
00:00:34

FailoverDrainToER Task Name      Failover drain to
ER.

FailoverDrainToER Task State
Completed

FailoverDrainToER Short Description  Failover drain to ER deactivate old
replication path and activate new replication path for external replication
system.
FailoverDrainToER Long Description  Failover drain to ER completed.
Update mark.
FailoverDrainToER Additional Info    Please run command
'sap_host_available' to complete the failover
operation.
FailoverDrainToER Current Task Number
5

FailoverDrainToER Total Number of Tasks
5

FailoverDrainToER Task Start      Wed Sep 07 12:17:12 UTC
2016

```

```
FailoverDrainToER Task End          Wed Sep 07 12:17:46 UTC
2016
FailoverDrainToER Hostname
site0
```

When `sap_failover_drain_to_er` has finished successfully, all backlogs are drained to the external replication system and replication is established from the new primary node to the external system.

4. Restore the SAP ASE server on the former primary node. When the node is ready to rejoin the replication system, run a similar command in RMA:

```
sap_host_available PR
go
```

The system displays the following information when the command has finished successfully.

```
TASKNAME          TYPE          VALUE
-----
-----
HostAvailable Start Time          Wed Sep 07 12:18:43 UTC 2016
HostAvailable Elapsed Time        00:02:13
HostAvailable Task Name           HostAvailable
HostAvailable Task State         Completed
HostAvailable Short Description    Resets the original source logical host
when it is available after failover.
HostAvailable Long Description   Completed the reset process of logical
host 'PR' receiving replication from logical host 'HA'.
HostAvailable Current Task Number  10
HostAvailable Total Number of Tasks 10
HostAvailable Task Start          Wed Sep 07 12:18:43 UTC 2016
HostAvailable Task End            Wed Sep 07 12:20:56 UTC 2016
HostAvailable Hostname            site0
```

Replication from the former standby node to the former primary node is established.

5. Verify replication path status.
 - o To verify the replication paths within the HADR cluster, run the `sap_status path` command in RMA.

```
sap_status path
go

PATH          NAME          VALUE
INFO
-----
-----
Time command started executing.      Start Time          2016-09-07 12:21:13.446
Command execution time.              Elapsed Time        00:00:00
PR                                     Hostname            site0
Logical host name.                   HADR Status         Standby : Inactive
Identify the primary and standby sites. Synchronization Mode Synchronous
PR                                     Synchronization State Inactive
The configured Synchronization Mode value.
PR                                     Distribution Mode    Remote
Synchronization Mode in which replication is currently operating.
Configured value for the distribution_mode replication model property.
PR                                     Replication Server Status Active
The status of Replication Server.
```



```

HA                               Hostname                               HA
Logical host name.
HA                               HADR Status                             Primary : Active
Identify the primary and standby sites.
HA                               Synchronization Mode                     Synchronous
The configured Synchronization Mode value.
HA                               Synchronization State                     Synchronous
Synchronization Mode in which replication is currently operating.
HA                               Distribution Mode                           Remote
Configured value for the distribution_mode replication model property.
HA                               Replication Server Status                 Active
The status of Replication Server.
PA.HA.ERP                       State                                   Suspended
Path is suspended (Replication Agent Thread). Transactions are not being replicated.
PA.HA.ERP                       Latency Time                             Unknown
No latency information for database 'ERP'.
PA.HA.ERP                       Latency                                   Unknown
No latency information for database 'ERP'.
PA.HA.ERP                       Commit Time                               Unknown
No last commit time for the database 'ERP'.
PA.HA.ERP                       Distribution Path                           HA
The path of Replication Server through which transactions travel.
PA.HA.ERP                       Drain Status                               Unknown
The drain status of the transaction logs of the primary database server.
PA.HA.ERP_1                     State                                   Suspended
Path is suspended (Replication Agent Thread). Transactions are not being replicated.
PA.HA.ERP_1                     Latency Time                             Unknown
No latency information for database 'ERP_1'.
PA.HA.ERP_1                     Latency                                   Unknown
No latency information for database 'ERP_1'.
PA.HA.ERP_1                     Commit Time                               Unknown
No last commit time for the database 'ERP_1'.
PA.HA.ERP_1                     Distribution Path                           HA
The path of Replication Server through which transactions travel.
PA.HA.ERP_1                     Drain Status                               Unknown
The drain status of the transaction logs of the primary database server.
PA.HA.master                   State                                   Suspended
Path is suspended (Replication Agent Thread). Transactions are not being replicated.
PA.HA.master                    Latency Time                             Unknown
No latency information for database 'master'.
PA.HA.master                    Latency                                   Unknown
No latency information for database 'master'.
PA.HA.master                    Commit Time                               Unknown
No last commit time for the database 'master'.
PA.HA.master                    Distribution Path                           HA
The path of Replication Server through which transactions travel.
PA.HA.master                    Drain Status                               Unknown
The drain status of the transaction logs of the primary database server.
HA.PR.ERP                       State                                   Active
Path is active and replication can occur.
HA.PR.ERP                       Latency Time                             2016-09-07 12:20:12.340
Time latency last calculated
HA.PR.ERP                       Latency                                   205560
Latency (ms)
HA.PR.ERP                       Commit Time                               2016-09-07 12:20:12.346
Time last commit replicated
HA.PR.ERP                       Distribution Path                           PR
The path of Replication Server through which transactions travel.
HA.PR.ERP                       Drain Status                               Not Drained
The drain status of the transaction logs of the primary database server.
HA.PR.ERP_1                     State                                   Active
Path is active and replication can occur.
HA.PR.ERP_1                     Latency Time                             2016-09-07 12:20:26.080
Time latency last calculated

```

```

HA.PR.ERP_1          Latency          219430
Latency (ms)
HA.PR.ERP_1          Commit Time      2016-09-07 12:20:26.080
Time last commit replicated
HA.PR.ERP_1          Distribution Path PR
The path of Replication Server through which transactions travel.
HA.PR.ERP_1          Drain Status     Not Drained
The drain status of the transaction logs of the primary database server.
HA.PR.master        State          Active
Path is active and replication can occur.
HA.PR.master         Latency Time     2016-09-07 12:20:19.506
Time latency last calculated
HA.PR.master         Latency          213143
Latency (ms)
HA.PR.master         Commit Time      2016-09-07 12:20:19.506
Time last commit replicated
HA.PR.master         Distribution Path PR
The path of Replication Server through which transactions travel.
HA.PR.master         Drain Status     Not Drained
The drain status of the transaction logs of the primary database server.

```

i Note

If the replication data load is low, the synchronization state may not update to *Synchronous* after you run the `sap_host_available` command to establish replication. To refresh its value, run the `sap_send_trace <primary_host_name>` command, then re-run the `sap_status path` command.

- o To verify the replication to the external system, connect to SAP Replication Server on that site and run the `sysadmin path_check` command:

```

sysadmin path_check
go

```

```

Primary Database Replicate Database Status
-----
HADR_PR.ERP      HADR_2.ABC      Up

```

6.5.5 Configuring a Primary SAP ASE to Replicate Data to an External System When the Standby Host is Down

When the standby host in an external replication system goes down, configure the primary SAP ASE server temporarily to replicate data to the external SAP Replication Server directly. After the standby host is restored, you can then configure the primary SAP ASE server to connect to the standby SAP Replication Server once again.

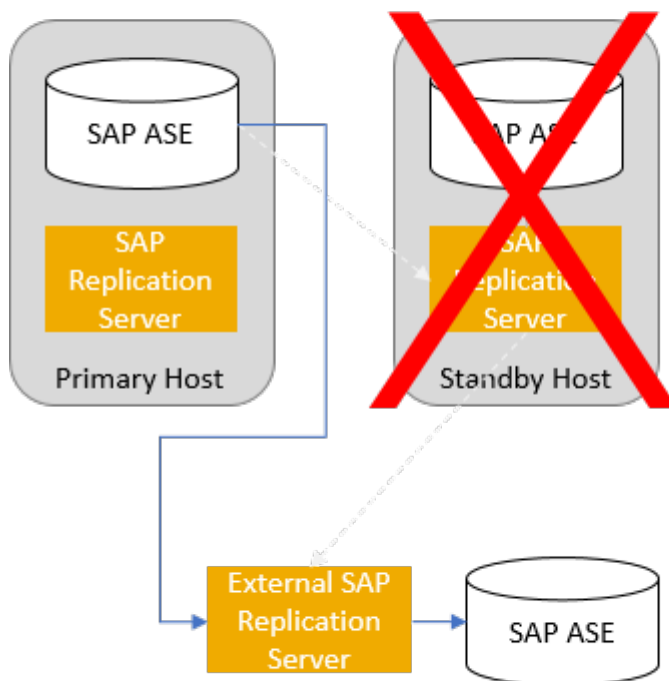
In an external replication system, the SAP Replication Server instance on the standby host receives data from the primary SAP ASE server, then replicates data to the external SAP Replication Server. Replication to the external system stops when the standby host goes down. To resume data replication, configure the primary SAP ASE server to bypass the standby SAP Replication Server, and replicate data to the external SAP Replication Server directly. After you restore the standby host, reconfigure the primary SAP ASE server to connect to the standby SAP Replication Server and rematerialize the standby databases.

i Note

If you choose to configure the primary SAP ASE server to connect to the external SAP Replication Server when the standby host is down, you must then rematerialize the standby databases after reconfiguring the

connection back to the standby SAP Replication Server. If you choose not to materialize them, do not configure the primary SAP ASE server to connect to the external SAP Replication Server when the standby host is down; wait until the standby host is restored before resuming the replication.

The following diagram shows the data flow after you configure the primary SAP ASE server to connect to the external SAP Replication Server:



Choose either configuration method:

- Automatic – the preferred method. See [Automatically Configuring Primary SAP ASE to Replicate Data to External System When the Standby Host is Down](#) [page 284].
- Manual – use this method only if the automatic method does not work. See [Manually Configuring Primary SAP ASE to Replicate Data to External System When the Standby Host is Down](#) [page 286].

6.5.5.1 Automatically Configuring Primary SAP ASE to Replicate Data to External System When the Standby Host is Down

When the standby host goes down, you can configure the primary SAP ASE server to replicate data to the external SAP Replication Server, then configure it to connect to the standby SAP Replication Server again after the standby host goes up automatically.

Perform the following steps:

1. Shut down the Fault Manager. See [Shutting Down the Fault Manager \[page 284\]](#).
2. Configure RepAgent to connect to the external system. See [Configuring RepAgent to Connect to the External System \[page 285\]](#).
3. Reconfigure RepAgent after restoring the standby host. See [Reconfiguring RepAgent After Standby Host is Restored \[page 285\]](#).

Related Information

[Manually Configuring Primary SAP ASE to Replicate Data to External System When the Standby Host is Down \[page 286\]](#)

6.5.5.1.1 Shutting Down the Fault Manager

Shut down the Fault Manager so that the primary SAP ASE is not demoted to standby SAP ASE when the standby host goes down.

Context

Perform these steps to manually shut down Fault Manager when you cannot stop it gracefully.

Procedure

1. Manually kill the heartbeat process on the primary host if you cannot stop it by executing `sybdbfm stop`:

```
$ ps -aef | grep sybdbfm
root <hb_pid> 1 0 20:12 ? 00:00:02 /usr/sap/hostctrl/exe/PDS/hb/sybdbfm_PDS
hahb pf=/usr/sap/hostctrl/exe/PDS/hb/SYBHA_PDS.PFL
sudo kill -9 <hb_pid>
```

2. Kill the Fault Manager process on the host on which the Fault Manager is running:

```
kill -9 <fm_pid>
```

3. Disable the configuration parameter `FM enabled` on the primary SAP ASE:

```
sp_configure 'FM enabled', 0
```

6.5.5.1.2 Configuring RepAgent to Connect to the External System

When the standby host goes down, configure RepAgent to connect to the external SAP Replication Server to enable data replication from the primary SAP ASE to the external system.

Procedure

Execute the `sap_configure_rat` command with the `redirect_to_er` parameter to redirect the primary RepAgent to connect to the external SAP Replication Server:

```
sap_configure_rat redirect_to_er, {<database> | All}, <ER admin user>, <ER admin password>
```

- `<database> | All`
Specify `<database>` to redirect the connection for a specific database, and specify `All` to redirect the connection for the whole HADR environment.
- `<ER admin user>, <ER admin password>`
Enter the admin user and password of the external SAP Replication Server to allow RMA to connect to it.

6.5.5.1.3 Reconfiguring RepAgent After Standby Host is Restored

After you restore the standby host but before you start any HADR components, reconfigure RepAgent by configuring the primary SAP ASE server to connect to the standby SAP Replication Server again.

Procedure

Execute the `sap_configure_rat` command with the `redirect_to_ha` parameter to redirect the primary RepAgent to connect to the standby SAP Replication Server:

```
sap_configure_rat redirect_to_ha, {<database> | All}, <ER admin user>, <ER admin password>
```

- `<database>` | All
Specify `<database>` to redirect the connection for a specific database, and specify `ALL` to redirect the connection for the whole HADR environment.
- `<ER admin user>`, `<ER admin password>`
Enter the admin user and password of the external SAP Replication Server to allow RMA to connect to it.

6.5.5.2 Manually Configuring Primary SAP ASE to Replicate Data to External System When the Standby Host is Down

When the standby host goes down, you can configure the primary SAP ASE server to replicate data to the external SAP Replication Server, then configure it to connect to the standby SAP Replication Server again after the standby host goes up manually.

Perform the following steps:

1. Shut down the Fault Manager. See [Shutting Down the Fault Manager \[page 286\]](#).
2. Collect RepAgent configuration parameters. See [Collecting RepAgent Configuration Parameters \[page 287\]](#).
3. Configure RepAgent to connect to the external system. See [Configuring RepAgent to Connect to the External System \[page 288\]](#).
4. Reconfigure RepAgent after the standby site becomes available. See [Reconfiguring RepAgent After Standby Host is Restored \[page 290\]](#).

Related Information

[Manually Configuring Primary SAP ASE to Replicate Data to External System When the Standby Host is Down \[page 286\]](#)

6.5.5.2.1 Shutting Down the Fault Manager

Shut down the Fault Manager so that the primary SAP ASE is not demoted to standby SAP ASE when the standby host goes down.

Context

Perform these steps to manually shut down Fault Manager when you cannot stop it gracefully.

Procedure

1. Manually kill the heartbeat process on the primary host if you cannot stop it by executing `sybdbfm stop`:

```
$ ps -aef | grep sybdbfm
root <hb_pid> 1 0 20:12 ? 00:00:02 /usr/sap/hostctrl/exe/PDS/hb/sybdbfm_PDS
hahb pf=/usr/sap/hostctrl/exe/PDS/hb/SYBHA_PDS.PFL
sudo kill -9 <hb_pid>
```

2. Kill the Fault Manager process on the host on which the Fault Manager is running:

```
kill -9 <fm_pid>
```

3. Disable the configuration parameter `FM enabled` on the primary SAP ASE:

```
sp_configure 'FM enabled', 0
```

6.5.5.2.2 Collecting RepAgent Configuration Parameters

Run `sp_config_rep_agent` on the primary SAP ASE to collect RepAgent configuration parameters, which are needed to reconfigure RepAgent when the standby host is back up.

Procedure

1. Run `sp_config_rep_agent`:

```
sp_config_rep_agent <user_database>
go
```

2. Collect the information for the following configuration parameters:
 - `rs servername` – the server name of the standby SAP Replication Server to which the RepAgent on the primary SAP ASE server connects.
 - `rs username` – the user name created on the standby SAP Replication Server used for creating the connection from the RepAgent on the primary SAP ASE server.
 - `connect dataserver` – the server name of the SAP ASE server on which the RepAgent connects to the standby SAP Replication Server, and which was indicated when connecting to the standby host. This value may differ depending on whether you connected to the standby host or external replication system.
 - `stream mode` – the stream replication mode, `sync` or `async`.

6.5.5.2.3 Configuring RepAgent to Connect to the External System

When the standby host goes down, configure RepAgent to connect to the external SAP Replication Server to enable data replication from the primary SAP ASE to the external system.

Procedure

1. Stop RepAgent on the primary database:

```
sp_stop_rep_agent <user_database>
go
```

2. Log into the external SAP Replication Server instance and suspend the log transfer:

```
suspend log transfer from <data_server>.<user_database>
go
```

3. Log into the primary SAP ASE to configure RepAgent to connect to the external SAP Replication Server:

```
sp_config_rep_agent <user_database>,"rs servername","<external_rs_servername>
| <external_rs_host> : <external_rs_port>"
go
sp_config_rep_agent <user_database>,"rs username",<spq_ra_user>
go
sp_config_rep_agent <user_database>,"rs password",<password>
go
sp_config_rep_agent <user_database>,'connect dataserver', <server_name>
go
sp_config_rep_agent <user_database>,'stream replication','true'
go
sp_config_rep_agent <user_database>,'stream mode','async'
go
```

Note

- The value of `connect dataserver` may differ when connecting to the standby host and the external system. The server name you need to provide here is the server name that was indicated when creating the connection to the external system. See [Configuring Replication Out From an HADR System \[page 251\]](#).
- The `rs username` is the user you created on the external SAP Replication Server when establishing the connection from the HADR system to the external replication system. The user was created for SPQ Agent to connect to the external SAP Replication Server. See [Configuring Replication Out From an HADR System \[page 251\]](#).

4. Log into the external SAP Replication Server and resume the log transfer:

```
resume log transfer from <data_server>.<user_database>
go
```

5. Log into the primary SAP ASE and start the RepAgent:

```
sp_start_rep_agent <user_database>
go
```


- (Optional) If the standby node is down for a long time, you may also need to stop the RepAgent and temporarily disable the replication paths for the databases that do not participate in external replication, such as master and CID, to avoid filling the logs:


```
sp_stop_rep_agent <database_name>
go
sp_config_rep_agent <database_name> 'disable'
go
```

6.5.5.2.3.1 Preventing Data Loss When Connecting Primary SAP ASE to External SAP Replication Server

By default, SAP Replication Server ensures that there is no data loss when you switch the primary SAP ASE to replicate data to the external SAP Replication Server.


In stream replication, when simple persistent queue (SPQ) readers Capture and SPQ Agent read data from the SPQ, they mark their truncation points in different spots in the SPQ due to different data processing speeds. SAP Replication Server sends the truncation point (TP) to RepAgent as its secondary truncation point (STP) according to two different mechanisms:

- Mechanism A sends the truncation point of the faster SPQ reader to RepAgent as its secondary truncation point:



In mechanism A, SAP Replication Server sends the truncation point of the faster SPQ reader (TP2) to RepAgent as the secondary truncation point. RepAgent truncates all data before TP2 in the primary SAP ASE log. In this situation, if replication to the external system is slower than the standby host and the standby host is down, switching the primary SAP ASE to connect to the external SAP Replication Server causes data loss between TP1 and TP2.

- Mechanism B sends the truncation point of the slower SPQ to RepAgent as its secondary truncation point:



SAP Replication Server uses mechanism B (sending the truncation point of the slower SPQ reader, that is TP1, to RepAgent as its secondary truncation point) by default to make sure that no data is lost after switching the connection of RepAgent from the standby to the external SAP Replication Server. In this mechanism, RepAgent always truncates data before TP1 in the primary SAP ASE log. Even if the external system replicates data slower than the standby host, RepAgent does not truncate data that has not been replicated to the external SAP Replication Server.

One side effect of mechanism B is that if replication to the external system is slow, data that has not been replicated to the external system accumulates in SAP ASE log files. To change the default behavior to prevent the accumulation of data in the ASE logs, turn on the trace `CI_MOVE_TP_BY_GTP_OFF`:

```
trace 'on', ci, CI_MOVE_TP_BY_GTP_OFF
go
```

When the trace is turned on, SAP Replication Server uses mechanism A and sends the truncation point of the faster SPQ reader to RepAgent as its secondary truncation point. The secondary truncation point moves faster and the logs are truncated faster in the primary SAP ASE. However, in this situation, there may be data loss when the standby host is down and you switch the primary SAP ASE server to connect to the external SAP Replication Server. To prevent data loss, turn off the trace `CI_MOVE_TP_BY_GTP_OFF`:

```
trace 'off', ci, CI_MOVE_TP_BY_GTP_OFF
go
```

6.5.5.2.4 Reconfiguring RepAgent After Standby Host is Restored

After you restore the standby host, configure the primary SAP ASE server to connect to the standby SAP Replication Server again, by reconfiguring RepAgent before starting any HADR components.

Procedure

1. Log into the primary SAP ASE and stop RepAgent on the primary database:

```
sp_stop_rep_agent <user_database>
go
```

2. Log into the external SAP Replication Server to suspend the log transfer:

```
suspend log transfer from <data_server>.<user_database>
go
```

3. Log into the primary SAP ASE and reconfigure RepAgent on the primary database to connect to the standby SAP Replication Server:

```
sp_config_rep_agent <user_database>,"rs servername","<standby_rs_servername>
| <standby_rs_host> : <standby_rs_port>"
go
sp_config_rep_agent <user_database>,"rs username", <rs_username>
go
sp_config_rep_agent <user_database>,"rs password", <password>
go
sp_config_rep_agent <user_database>,"connect dataserver", <dataserver_name>
go
sp_config_rep_agent <user_database>,"stream replication", 'true'
go
sp_config_rep_agent <user_database>,"stream mode", 'sync'
go
```

i Note

The values you provide here for the following configuration parameters should be the values you have collected that are used for the connection between the primary SAP ASE server and the standby SAP Replication Server. See [Collecting RepAgent Configuration Parameters \[page 287\]](#):

- rs servername
- rs username
- rs password
- connect dataserver

4. Start the SAP ASE server on the standby host.
5. Start the SAP ASE backup server on the standby host.
6. Start the SAP Replication Server instance on the standby host.
7. (Optional) Start RMA on both the primary and the standby sites if they are not already started.
8. Log into the external SAP Replication Server and resume log transfer:

```
resume log transfer from <data_server>.<user_database>
go
```

9. Log into the primary SAP ASE and start the RepAgent:

```
sp_start_rep_agent <user_database>
go
```

10. Rematerialize standby SAP ASE databases that are replicated to the external system. See [Rematerializing Databases for External Replication \[page 262\]](#).
11. (Optional) Rematerialize the databases that do not participate in external replication, such as master and CID, if you disable their data replication when the standby site is down. See [Rematerializing Databases for External Replication \[page 262\]](#).

6.6 Tuning the SPQ Agent for External Replication

Tuning the performance of an external Replication Server involves configuring two SPQ Agent parameters, `ci_pool_size` and `ci_package_size`.

- `ci_pool_size` – Specifies the size of Component Interface (CI) buffer pool. For data replication from the simple persistent queue (SPQ) to the SPQ Agent, use the `alter connection` command to change the value of `ci_pool_size` parameter. For example:

```
alter connection ds.db set ci_pool_size to '100'
```
- `ci_package_size` – Specifies the size of a CI package. Each CI package in CI buffer pool shares the same size configured by `ci_package_size` parameter. Use the same value for `ci_package_size` as you used in the Rep Agent.

6.6.1 System Hardware, Operating System, and Networking Configurations

This section lists the recommended hardware, operating system, and networking configurations for an external replication system.

- System Hardware
 - Simple Persistent Queue (SPQ) must be on a dedicated solid-state drive (SSD).
 - Use 8 to 32 GB of memory and 4 to 12 cores, depending on the size on the host for HADR components.
- Operating System
 - Tune file system for SPQ and Replication Server partitions.
 - For SPQ, consider using PCIe solid-state drive (SSD) technology over host bus adapters (HBAs). Depending on sizing and volume, inbound queue (IBQ) and outbound queue (OBQ) may share PCIe SSD with SPQ.
 - Separate log and data. The file systems hosting logs and data should be on different mount points and different volumes.
- Network Configurations
 - You should use 10 gigabit Ethernet, or a separate subnet and NICs for high-availability components.
 - Set network TCP send and receive buffer.
 - Tune NIC interface for queue size.

7 Testing Your HADR Configuration

There are a number of commands you can run to confirm that your HADR system is running correctly.

To verify the HADR functionality and configuration:

- Log into the primary and companion servers and use the `hadr_mode`, `hadr_state`, and the `asehostname` functions to confirm that:
 - The HADR mode is "primary"
 - The HADR state is "active" on the primary server
 - The HADR mode is "standby"
 - The HADR state is "inactive" on the companion server
 - The host names are correct

This example confirms that the HADR mode is "Primary" and the HADR state is "Active" on the primary site, mo-ae3a62265:

```
select asehostname(), hadr_mode(), hadr_state()
-----
mo-ae3a62265      Primary      Active
(1 affected)
```

This example confirms that the HADR mode is "Standby" and the HADR state is "Inactive" on the companion site, mo-4a63cdeba:

```
select asehostname(), hadr_mode(), hadr_state()
-----
mo-4a63cdeba     Standby     Inactive
(1 row affected)
```

- Connect to RMA on one of the hosts (the primary, mo-ae3a62265, in this example) and execute the `sap_status_path` command to confirm that the synchronization state shows "Synchronous" and the paths are "Active" for all replicated databases (see the bold text):

```
isql -UDR_admin -Psybase123 -Smo-ae3a62265:7001 -w1000
1> sap_status_path
2> go
  PATH                NAME                VALUE
-----
-----
command started executing.
                        Start Time                2015-08-07 13:02:50.707 Time
                        Elapsed Time                00:00:00
Command execution
time.
Bangalore             Hostname                mo-4a63cdeba
Logical host
name.
Bangalore             HADR Status            Standby : Inactive
Identify the primary and standby
sites.
Bangalore             Synchronization Mode    Synchronous           The
configured Synchronization Mode
value.
```

Bangalore	Synchronization State	Inactive	
Synchronization Mode in which replication is currently operating.			
Bangalore	Distribution Mode	Remote	
Configured value for the distribution_mode replication model property.			
Bangalore	Replication Server Status	Active	The
status of Replication Server.			
Pune	Hostname	mo-ae3a62265	
Logical host name.			
Pune	HADR Status	Primary : Active	
Identify the primary and standby sites.			
Pune	Synchronization Mode	Synchronous	The
configured Synchronization Mode value.			
Pune	Synchronization State	Synchronous	
Synchronization Mode in which replication is currently operating.			
Pune	Distribution Mode	Remote	
Configured value for the distribution_mode replication model property.			
Pune	Replication Server Status	Active	The
status of Replication Server.			
Bangalore.Pune.IND	State	Suspended	Path
is suspended (Replication Agent Thread). Transactions are not being replicated.			
Bangalore.Pune.IND	Latency	Unknown	No
latency information for database 'IND'.			
Bangalore.Pune.IND	Commit Time	Unknown	No
last commit time for the database 'IND'.			
Bangalore.Pune.IND	Distribution Path	Pune	The
path of Replication Server through which transactions travel.			
Bangalore.Pune.master	State	Suspended	Path
is suspended (Replication Agent Thread). Transactions are not being replicated.			
Bangalore.Pune.master	Latency	Unknown	No
latency information for database 'master'.			
Bangalore.Pune.master	Commit Time	Unknown	No
last commit time for the database 'master'.			
Bangalore.Pune.master	Distribution Path	Pune	The
path of Replication Server through which transactions travel.			
Pune.Bangalore.IND	State	Active	Path is
active and replication can occur.			
Pune.Bangalore.IND	Latency Time	2015-08-07 09:28:12.144	Time
latency last calculated			
Pune.Bangalore.IND	Latency	47	
Latency (ms)			
Pune.Bangalore.IND	Commit Time	2015-08-07 09:28:13.440	Time
last commit replicated			
Pune.Bangalore.IND	Distribution Path	Bangalore	The
path of Replication Server through which transactions travel.			
Pune.Bangalore.master	State	Active	Path is
active and replication can occur.			
Pune.Bangalore.master	Latency Time	2015-08-06 15:36:54.952	Time
latency last calculated			
Pune.Bangalore.master	Latency	763	
Latency (ms)			

```
Pune.Bangalore.master Commit Time          2015-08-06 15:36:54.960 Time
last commit replicated
Pune.Bangalore.master Distribution Path      Bangalore           The
path of Replication Server through which transactions
travel.
```

- Insert an `rs_ticket` on the primary server, and verify that it appears on the companion server to confirm that replication is functioning correctly:

```
use IND
go
rs_ticket "Testing_HADR_Configuration"
go
(return status = 0)
```

On the companion server, query the `rs_ticket_history` table in the same database to find the inserted ticket.

- Verify the Fault Manager is running correctly by issuing `sybdbfm status` from the installation directory on the host running the Fault Manager to confirm that the Fault Manager is monitoring the HADR configuration, and that the command displays the replication status as `SYNC_OK`. For example:

```
fault manager running, pid = 34534, fault manager overall status = OK,
currently executing in mode PAUSING
*** sanity check report (8720)***.
node 1: server mo-ae3a62265.mo.sap.corp, site hasite0.db host status: OK.db
status OK hadr status PRIMARY.
node 2: server mo-4a63cdeba.mo.sap.corp, site hasite1.db host status: OK.db
status OK hadr status STANDBY.
replication status: SYNC_OK.
```

Also, verify that the HADR tab in the SAP ASE Cockpit shows the Fault Manager status under [Service Components](#) as *UP*.

8 HADR Administration

Administering the HADR system includes adding users and databases, materializing and rematerializing databases, performing planned failover, recovering from unplanned failover, and so on.

8.1 Using isql to Log in to SAP ASE, Replication Server, and RMA

Many of the administrative tasks require that you log into SAP ASE, Replication Server, and RMA at the `isql` command line.

Set the SAP ASE environment variables by sourcing the `$$SYBASE/SYBASE.csh` (for SAP ASE) and `$$SYBASE/DM/SYBASE.csh` files (for Replication Server and RMA). You can view the Replication Server port numbers in the `$$SYBASE/DM/interfaces` file.

Use this syntax to log in to SAP ASE:

```
$$SYBASE/$$SYBASE_OCS/bin/isql -U<login_name> -P<password> -S<server_name>
```

For example:

```
/work/SAP1/OCS-16_0/bin/isql -Usa -Psecret396 -SSFSAP1
```

Use this syntax to log in to Replication Server:

```
$$SYBASE/OCS-16_0/bin/isql -U<login_name> -P<password> -S<host_name>:<Rep_Server_port>
```

For example:

```
/work/SAP1/OCS-16_0/bin/isql -Usa -Psecret396 -SSF_MACHINE1:5005
```

Use this syntax to log in to RMA:

```
isql -U<login_name> -P<password> -S<host_name>:<RMA_port_number>
```

For example:

```
/work/SAP1/OCS-16_0/bin/isql -UDR_admin -Psecret396 -SSF_MACHINE1:7001
```

See the SAP ASE *Utility Guide* for more information about `isql`.

8.2 Logins and Roles

Create the SAP ASE maintenance user to apply activity to the target system (this login is created automatically when you use `setuphadr` utility to set up HADR).

Context

The maintenance user requires a unique SAP ASE login. Do not use an existing SAP ASE user as the maintenance user.

i Note

Protect the Replication Server maintenance user's password. See the *Replication Server Administration Guide > Manage Database Connections > Manage the Maintenance User*.

Replication Server applies changes to the standby database using the unique maintenance user login.

8.2.1 Creating the Maintenance Login

Replication Server uses the maintenance login to log into SAP ASE servers in the HADR system.

Context

The maintenance user requires a unique SAP ASE login. Do not use an existing SAP ASE user as the maintenance user.

i Note

Protect the Replication Server maintenance user's password. See the *Replication Server Administration Guide > Manage Database Connections > Manage the Maintenance User*.

Replication Server applies changes to the standby database using the unique maintenance user login.

To add the maintenance user, perform these steps on both the primary and standby server:

- Create the maintenance login on SAP ASE servers in the HADR system.
- The maintenance login name is `<SID_name>_maint`. For example, if the SID of the HADR system is SAP1, the name of the maintenance user is SAP1_maint.
- The maintenance login SUID is the same on all SAP ASE instances (the SAP installer sets the maintenance SUID to 1001 by default).

Procedure

1. Create the maintenance user. This example creates a maintenance user named `D01_maint`:

```
use master
go
create login D01_maint with password Sybase123
go
```

2. Grant the `replication_role` to the maintenance user, enabling it to replicate the `truncate table` command:

```
grant role replication_role to D01_maint
go
```

3. Use `sp_addalias` to alias the maintenance user to the database owner on the master and user databases, allowing this user to update tables that use `IDENTITY` columns:

```
sp_alias D01_maint, dbo
go
```

4. Grant the `sa_role` to the maintenance user so it can replicate `insert`, `update`, and `delete` operations on all tables:

```
grant role sa_role to D01_maint
```

5. Create a new role named `sap_maint_user_role`:

```
create role sap_maint_user_role
```

6. Grant `set session authorization` permissions to the maintenance user, allowing it to become another user when applying DDL commands to the replicate database. Grant the permission to a user or a role, not a login:

```
grant set session authorization to sap_maint_user_role
```

7. Automatically activate the maintenance user role upon login:

- a. Grant the `sap_maint_user_role` to the maintenance user:

```
sp_role "grant", sap_maint_user_role, D01_maint
```

- b. Automatically activate the maintenance user role at login:

```
alter login D01_maint add auto activated roles sap_maint_user_role
```

8.3 Adding Databases from the Command Line After Installation

You can add new or existing databases to the HADR system.

The HADR system maintains a list of databases that are replicated (issue the `sap_set` RMA command to see this list). However, user databases created by a `create database` command are not automatically added to the HADR participating database list.

Before you add a database to the HADR system:

- Create the new database on both the primary and standby servers with similar physical size configurations.

i Note

Creating the database on the standby server with the `for_load` parameter skips the page-clearing step and reduces operation time.

- The database administrator needs proper permissions to create and add the database.

To add a new database into the list:

1. Issue `sap_set` from the RMA command line to determine if the database is already added to the HADR system. If the `sap_set` output lists the database you are adding in the `participating_databases` line, this database is already included in the HADR replication system. For example, this indicates that the `pubs2` database is not included in the HADR replication system (see the bold text):

```
sap_set
go
PROPERTY                                VALUE
-----                                -
maintenance_user                        DR_maint
sap_sid                                  AS1
installation_mode                        nonBS
participating_databases                 [master,AS1]
connection_timeout                       5
connection_alloc_once                    true
. . .
```

2. Create the new database on the primary and standby servers, making sure the databases use appropriate sizing for the data and log devices. For example, if you create the `pubs2` database on the primary companion, create it on the secondary companion as well.
3. Add the `DR_maint` login as the database owner (dbo) alias on the primary and standby database for the newly created database. For example, on the newly created `pubs2` database:

```
use pubs2
go
sp_addalias DR_maint,dbo
go
```

4. Issue `sap_status path` to verify that the paths for all database are active. The output of `sap_status path` is very long. The lines that indicates the active paths should look similar to (see the bold text):

```
sap_status path
go
PATH          NAME          VALUE
INFO
```



```

go
TASKNAME          TYPE
      VALUE
-----
-----
Materialize      Start Time
                Thu Nov 19 21:27:24 UTC 2015
Materialize      Elapsed Time
                00:00:00
DRExecutorImpl  Task Name
                Materialize
DRExecutorImpl  Task State
                Running
DRExecutorImpl  Short Description
                Materialize database
DRExecutorImpl  Long Description
                Started task 'Materialize' asynchronously.
DRExecutorImpl  Additional Info
                Please execute command 'sap_status task' to determine when task
'Materi
                alize' is complete.
Materialize      Task Name
                Materialize
Materialize      Task State
                Running
Materialize      Short Description
                Materialize database
Materialize      Hostname
                SFMACHINE1

```

This command may take a long time to complete. Issue the `sap_status` command until it includes the line below in bold, which indicates the `sap_materialize` command is successful:

```

sap_status
go
TASKNAME          TYPE
      VALUE
-----
-----
. . .
Materialize Long Description
Completed automatic materialization of database 'pubs2' from source 'SF
HADR1' to target 'SJHADR2'.
. . .

```

Issue `sap_status path` to verify the HADR system is healthy and the `pubs2` database is added to the HADR system. The output of `sap_status path` is very long. The lines that indicates the active paths should look similar to the lines in bold below (which indicate that the `pubs2` database is active):

```

sap_status path
go
PATH              NAME              VALUE
      INFO
-----
-----
                Start Time              2015-11-19 20:21:41.185
Time command started executing.
                Elapsed Time              00:00:00
Command execution time.
. . .

```

```
SFHADR1.SJHADR2.pubs2 State Active
Path is active and replication can occur.
```

8.4 Loading from an External Dump

You can load data from an external dump using the RMA command line or from the SAP ASE Cockpit. Typically, external dumps are used to refresh local SAP ASE databases with the latest data dumped from another database system, or to migrate your existing SAP ASE database into an SAP ASE HADR-enabled database system.

To load from an external dump:

- The HADR system is set up and running.
- The database is included in the participating database list.
- The source database dump is compatible with the HADR database into which it is being loaded.
- The database administrator has the sa role and the DR_admin account (to log into RMA).

Loading an External Dump from RMA

This example describes loading a dump of the pubs2 database.

1. Issue `sap_status path` to verify the HADR system is healthy and the pubs2 database is added to the HADR system. The output of `sap_status path` is very long. The lines that indicates the active paths should look similar to the lines in bold below (which indicate that the pubs2 database is active):

```
sap_status path
go
PATH          INFO          NAME          VALUE
-----
-----
                Start Time          2015-11-19 20:21:41.185
Time command started executing.
                Elapsed Time          00:00:00
Command execution time.
. . .
SFHADR1.SJHADR2.pubs2 State Active
Path is active and replication can occur.
```

2. Log into the RMA on the primary host and disable replication for the database (this disables RMA for that database):

```
sap_disable_replication SFHADR1, pubs2
go
TASKNAME      TYPE
-----
Disable Replication Start Time
Fri Nov 20 01:06:01 UTC 2015
```

```

Disable Replication Elapsed Time
00:00:06
DisableReplication Task Name
Disable Replication
DisableReplication Task State
Completed
DisableReplication Short Description
Disable the flow of Replication
DisableReplication Long Description
Successfully disabled Replication for database 'pubs2'.
DisableReplication Task Start
Fri Nov 20 01:06:01 UTC 2015
DisableReplication Task End
Fri Nov 20 01:06:07 UTC 2015
DisableReplication Hostname
SFMACHINE1

```

3. Load the database dump by logging into the SAP ASE running on the primary site and executing a command similar to:

```
load database pubs2 from '/work/SAP1/pubs2_dump'
```

4. Bring the database online by logging into the SAP ASE running on the primary site and executing a command similar to:

```
online database pubs2
go
```

5. If necessary, add the DR_maint user as database owner for the database by logging into the SAP ASE running on the primary site and executing a command similar to:

```
use pubs2
go
sp_addalias DR_maint, dbo
go
```

6. Enable replication for the database by logging into the RMA running on the primary site and executing an RMA command similar to:

```
sap_enable_replication SFHADR1, pubs2
```

7. Issue `sap_status path` to verify that the paths for the database are active. The output of `sap_status path` is very long. The lines that indicates the active paths should look similar to (see the bold text):

```

sap_status path
go
PATH                NAME                VALUE
INFO
-----
-----
                Start Time                2015-11-19 20:21:41.185
Time command started
executing.
                Elapsed Time                00:00:00
Command execution time.
. . .
SFHADR1.SJHADR2.pubs2 State      Active
Path is active and replication can

```

occur.

8. Materialize the databases from the primary site to the standby site using either the automatic or manual method. This example uses the automatic method:

```
sap_materialize auto SFHADR1 SJHADR2 pubs2
go

TYPE
VALUE

-----
-----
Materialize      Start Time              Wed Nov 19 20:31:13 EST
2015
Materialize      Elapsed Time
00:00:00

DRExecutorImpl  Task Name
Materialize

DRExecutorImpl  Task State
Running

DRExecutorImpl  Short Description           Materialize
database
DRExecutorImpl  Long Description            Started task 'Materialize'
asynchronously.
DRExecutorImpl  Additional Info             Please execute command
'Materialize
Materialize      Task Name
Materialize

Materialize      Task State
Running

Materialize      Short Description           Materialize
database
Materialize      Long Description            Executing ASE dump and load
task for database 'pubs2'.
Materialize      Hostname
site0

ASEDBDumpAndLoad Task Name              Ase Dump and
Load
ASEDBDumpAndLoad Task State
Running

ASEDBDumpAndLoad Short Description       Dump and Load an Adaptive
Server Enterprise database
ASEDBDumpAndLoad Long Description       Dumping Adaptive Server
Enterprise database 'pubs2'.
ASEDBDumpAndLoad Failing Command Error Message The creation of dump directory
failed, it has already existed.
ASEDBDumpAndLoad Current Task Number
1

ASEDBDumpAndLoad Total Number of Tasks
2

ASEDBDumpAndLoad Hostname              SFMACHINE
```


This command may take a long time to complete. Issue the `sap_status` command until it includes the line below in bold, which indicates the `sap_materialize` command is successful:

```
sap_status
go
TASKNAME          TYPE
  VALUE
-----
-----
. . .
Materialize Long Description
Completed automatic materialization of database 'pubs2' from source 'SF
HADR1' to target 'SJHADR2'.
. . .
UpdateReplication Hostname
                  SFMACHINE1
```

- Issue `sap_status path` to verify that the paths for the database are active. The output of `sap_status path` is very long. The lines that indicates the active paths should look similar to the following (see the bold text), with the replication path from SFHADR1 to SJHADR2 for the `pubs2` database is in Active state:

```
sap_status path
go
PATH              NAME              VALUE
INFO
-----
-----
                          Start Time          2015-11-19 20:21:41.185
Time command started
executing.
                          Elapsed Time          00:00:00
Command execution time.
. . .
SFHADR1.SJHADR2.pubs2 State          Active
Path is active and replication can
occur.
```

Loading an External Dump from the SAP ASE Cockpit

- Log into RMA on the primary site and execute this RMA command:

```
sap_disable_replication <primary_logical_site>, <database_name>
```

- Load the database. See *Manage SAP ASE > Backup and Restore > Restoring (Loading) a Database > Generating a Database Load Sequence* in the SAP ASE Cockpit documentation.
- Enable replication by logging into RMA on the primary site and executing this RMA command:

```
sap_enable_replication <primary_logical_site>, <database_name>
```

- Rematerialize the database. See *Manage SAP ASE > Always-On (HADR) Option > Rematerialize Databases* in the SAP ASE Cockpit documentation.

8.5 Materializing and Rematerializing Databases

Materialize a database with the `sap_materialize` command.

i Note

Materializing a database is a resource-intensive process. Do not run `sap_materialize` on more than one database at a time.

Materializing a Database Using `sap_materialize`

Log into the RMA and use the `sap_materialize` command to materialize the databases from the primary site to the companion site using the automatic or manual method:

- Automatic – Log into RMA on the primary site and issue this RMA command:

```
sap_materialize auto, <primary_logical_site>, <companion_logical_site>,
<database_name>
```

For example, to materialize the `pubs2` database from the SFHADR1 site to the SJHADR2, issue:

```
sap_materialize auto, SFHADR_SITE1, SJHADR_SITE2, pubs2
```

See [sap_materialize \[page 479\]](#).

- Manual
 1. Log into the RMA running on the primary site and issue this RMA command:

```
sap_materialize start <primary_logical_site>, <companion_logical_site>,
<database_name>
```

For example:

```
sap_materialize start SFHADR1, SJHADR, pubs2
```

This command is asynchronous. Run this RMA command until it shows that the `sap_materialize` command succeeded:

```
sap_status
go
TASKNAME      TYPE
VALUE
-----
-----
-----
-----
-----
Status       Start Time      Wed Sep 28 21:56:47 EDT
2016
Status       Elapsed Time
00:00:03
```

```

Materialize Task Name
Materialize

Materialize Task State
Completed

Materialize Short Description Materialize
database

Materialize Long Description The prerequisite work for manually dumping
and loading database pubs2 is finished. You can use "dump database pubs2
to ... with compression = 1, label = 'RMA_DUMP_LABEL'" to dump the
database. Be sure to use the label option named 'RMA_DUMP_LABEL'.
Materialize Task Start          Wed Sep 28 21:56:47 EDT
2016

Materialize Task End            Wed Sep 28 21:56:50 EDT
2016

Materialize Hostname
SFMACHINE1

(9 rows affected)

```

The RMA prompts you to use the label, `RMA_DUMP_LABEL`, to distinguish the database dump that you issue on the primary server from other scheduled system dumps. Replication resumes when the dump marker with `RMA_DUMP_LABEL` arrives at SAP Replication Server. When you materialize the replication database automatically using `sap_materialize auto`, RMA dumps the database internally with the specified label. You can modify the label in the `/DM/RMA-/instances/AgentContainer/config/bootstrap.prop` or `/DM/RMA-/config/bootstrap.prop` files. For more information, see [sap_materialize \[page 479\]](#).

2. Dump the database from the primary server using `RMA_DUMP_LABEL` and load it onto the companion server without loading the source dump onto the companion.

```

dump database <database_name>
with label = '<label_name>'

```

For example,

```

dump database pubs2
with compression = 1,
    label = 'RMA_DUMP_LABEL'

```

Note

Run only one `dump database` command between each `sap_materialize start` and `sap_materialize finish` commands you perform. Any modifications you make to the master database between the time of making the dump and issuing `sap_materialization finish` are not applied to the materialization. This restriction applies to the master database only, not user databases.

3. Log into the RMA running on the primary site and issue this RMA command:

```

sap_materialize finish <primary_logical_site>, <companion_logical_site>,
<database_name>

```

For example:

```

sap_materialize finish SFHADR, SJHADR, pubs2

```

This command is asynchronous. Run this RMA command until it shows that the `sap_materialized` command succeeded:

```
sap_status task
```

See [sap_materialize \[page 479\]](#) and [sap_status \[page 520\]](#).

Rematerializing Databases

Rematerialization reactivates the replication paths that are inactive due to replication problems (such as row count mismatch). Replication paths are critical for keeping the primary and standby databases in sync. Inactive replication paths lead to data inconsistency between the primary and standby HADR databases. Rematerialization resolves these inconsistencies.

To rematerialize a database:

- The database administrator has the sa role and the DR_admin account (to log into RMA).
- The primary and standby database are online and accessible.
- The primary and standby Replication Servers are online.

The steps below describe rematerializing a database from the RMA command line. You can also rematerialize databases from the SAP ASE Cockpit. See *Manage SAP ASE > Always-On (HADR) Option > Rematerializing Databases* in the SAP ASE Cockpit documentation for information.

This example rematerializes the `pubs2` database.

1. Disable replication to the database to be rematerialized. You can perform this task with replication to the standby database in a suspended state. From RMA issue:

```
sap_disable_replication SFHADR1, SJHADR2, pubs2
go
TASKNAME                TYPE
VALUE
-----
Disable Replication Start Time      Wed Nov 04 03:06:03 EST
2015
Disable Replication Elapsed Time    00:00:03
Disable Replication Task Name       Disable
Replication
Disable Replication Task State      Completed
Disable Replication Short Description Disable the flow of
Replication
Disable Replication Long Description Successfully disabled Replication for
database 'pubs2'.
Disable Replication Task Start      Wed Nov 04 03:06:03 EST
2015
Disable Replication Task End        Wed Nov 04 03:06:06 EST
2015
Disable Replication Hostname        SFHADR1
```

2. Enable replication to the database to be rematerialized. From RMA issue:

```
sap_enable_replication SFHADR1, SJHADR2, pubs2
go
```

```

TASKNAME          TYPE
VALUE
-----
-----
Enable Replication Start Time      Wed Nov 04 03:44:52 EST
2015
Enable Replication Elapsed Time
00:01:42
Enable Replication Task Name      Enable
Replication
Enable Replication Task State
Completed
Enable Replication Short Description Enable the flow of
Replication
Enable Replication Long Description Successfully enabled Replication for
database 'pubs2'.
Enable Replication Task Start      Wed Nov 04 03:44:52 EST
2015
Enable Replication Task End        Wed Nov 04 03:46:34 EST
2015
Enable Replication Hostname
SFHADR1

```

- Issue `sap_status path` to verify the HADR system is healthy and the `pubs2` database is added to the HADR system. The output of `sap_status path` is very long. The lines that indicates the active paths should look similar to the lines in bold below (which indicate that the `pubs2` database is defined):

```

sap_status path
go
PATH          INFO          NAME          VALUE
-----
-----
-----
Start Time      2015-11-19 20:21:41.185
Time command started executing.
Elapsed Time    00:00:00
Command execution time.
. . .
SFHADR1.SJHADR2.pubs2 State      Defined
Path is defined and ready for replication.

```

- Materialize the newly added database so it synchronizes with the corresponding database on the standby server. This example uses automatic materialization:

```

sap_materialize auto SFHADR1, SJHADR2, pubs2
go
TASKNAME          TYPE
VALUE
-----
-----
Materialize      Start Time      Wed Nov 04 04:00:12 EST
2015
Materialize      Elapsed Time
00:00:00

DRExecutorImpl  Task Name
Materialize

DRExecutorImpl  Task State
Running

```

```

DRExecutorImpl Short Description Materialize
database
DRExecutorImpl Long Description Started task 'Materialize'
asynchronously.
DRExecutorImpl Additional Info Please execute command
'sap_status task' to determine when task 'Materialize' is complete.
Materialize Task Name
Materialize

Materialize Task State
Running

Materialize Short Description Materialize
database
Materialize Long Description Executing ASE dump and load
task for database 'pubs2'.
Materialize Hostname
site0

ASEDBDumpAndLoad Task Name Ase Dump and
Load
ASEDBDumpAndLoad Task State
Running

ASEDBDumpAndLoad Short Description Dump and Load an Adaptive
Server Enterprise database
ASEDBDumpAndLoad Long Description Dumping Adaptive Server
Enterprise database 'db1'.
ASEDBDumpAndLoad Failing Command Error Message The creation of dump directory
failed, it has already existed.
ASEDBDumpAndLoad Current Task Number
1

ASEDBDumpAndLoad Total Number of Tasks
2

ASEDBDumpAndLoad Hostname
SFHADR1

```

This command may take a long time to complete. Issue the `sap_status` command until it includes the line below in bold, which indicates that the `sap_materialize` command succeeded:

```

sap_status task
go
TASKNAME          TYPE
      VALUE
-----
-----
. . .
Materialize Long Description
Completed automatic materialization of database 'pubs2' from source 'SF
HADR1' to target 'SJHADR2'.
. . .
UpdateReplication Hostname
                  SFMACHINE1

```

- Issue `sap_status path` to verify the HADR system is healthy and the `pubs2` database is added to the HADR system. The output of `sap_status path` is very long. The lines that indicates the active paths should look similar to the lines in bold below (which indicate that the `pubs2` database is active):

```

sap_status path
go
PATH              NAME          VALUE
      INFO
-----
-----

```

```

-----
-----
                Start Time                2015-11-19 20:21:41.185
Time command started executing.
                Elapsed Time                00:00:00
Command execution time.

SFHADR1.SJHADR2.pubs2  State      Active
Path is active and replication can occur.

```

8.6 Using SAP ASE Cockpit to Manage and Monitor the HADR System

The SAP Adaptive Server Enterprise Cockpit (SAP ASE Cockpit) is a Web-based tool for monitoring the status and availability of SAP ASE servers.

SAP ASE Cockpit provides availability monitoring, historical performance monitoring, and administration capabilities in a scalable Web application that are integrated with management modules for other SAP products, including the always-on option. The cockpit offers management of alerts that provide state- and threshold-based notifications about availability and performance in real-time, and intelligent tools for spotting performance and usage trends, all via a thin-client, rich Internet application delivered through your Web browser.

In an HADR environment, the cockpit offers a visual display of the status for SAP ASE and Replication Server, the modes in which they are currently running, and how efficient the connections are. This section provides an overview of the [MONITOR](#), [EXPLORING](#), and [ALERTS](#) tabs.

Along with the usual user logins, you can log into the SAP ASE Cockpit as the sa or the DR_admin. The DR_admin user is more restricted in its scope than sa, and is sufficient for most tasks in the SAP ASE Cockpit.


The SAP ASE Cockpit includes extensive online help. To view an online version, select [Help](#) at the top of the screen, or go to <http://help.sap.com/adaptive-server-enterprise>.

Monitoring the HADR System with the SAP ASE Cockpit

To see an overview of the HADR system:

1. In SAP ASE Cockpit, click the [MONITOR](#) tab.
2. In the left pane, click [HADR](#).

The SAP ASE Cockpit displays the primary and standby machines graphically as boxes, with a red, gray, or green line connecting the boxes. The server on the left is always the SAP ASE server on which you are focusing. The primary server is in light yellow color, and the standby server is in light gray color. A green line indicates that the systems are successfully connected. The site name, HADR status, synchronization mode, and synchronization state are indicated in the boxes. The logical host name, host name, and port number on which the server is run are shown in the inner, colored box.



A red line indicates that replication has stopped. A gray line indicates that replication is suspended.



Click the center button to display and close the *Replication Paths* table, which includes information about:

- *State* – describes the status of the replication paths:
 - Green arrow – Active
 - Gray arrow – Inactive
 - Yellow parallel bars – Suspended
 - Yellow warning sign with exclamation mark – Warning
 - Red warning sign with exclamation mark – Error
- The *Log Records*, *Backlog*, *Latency* columns indicate KPI alerts:
 - Green arrow – Good; no alerts
 - Yellow warning sign with exclamation mark – Warning; a KPI hit its low threshold
 - Red warning sign with exclamation mark – Error; a KPI hit its high threshold

The *Log Records* column indicates the % of Free Transaction Log KPI. The *Backlog* column reflects the ASE Backlog, Primary RS Backlog, and the Remote RS Backlog KPIs. The *Latency* column reflects the Total Latency KPI.

Note

The status of these columns is updated during the HADR statistics collection cycle. Mouse over the column headings to determine when they are updated. The column headings and the tool tip indicate the collection time and interval.

Click *Collapse* to remove the table.

The *Service Component Status* and the *Replication Paths Status* screens summarize the status of all components in the *Service Component* panel and the summary status in the middle Replication Paths table, respectively. The screens use the following labels and icons for different health conditions:

- *Service Component Status* – The green icon with an “Active” label indicates that there are no warnings or errors in any of the service components. The yellow icon with “Warning” label indicates that warnings exist in one of the service components. The red icon with “Error” label indicates that errors exist in one of the service components.
- *Replication Paths Status* –
 - The green icon with “Active” label indicates that there are no warnings or errors in any of the State, Log Records, Backlog and Latency columns. For the State column, Inactive and Suspended indicates warning; Down, Incomplete and Unknown means error; otherwise Active indicates active or good.
 - The yellow icon with “Warning” label indicates that warning exists in one of the State, Log Records, Backlog and Latency columns.

- The red icon with “Error” label indicates that an error exists in one of the columns.

Below the *Replication Paths* table is a table with the following tabs:

- *Service Components* – Shows the status of the local and remote RMA, Fault Manager, and the local and remote Replication Server. RMA is considered active if either the local or remote RMA is running.
- *Fault Manager Message* – Displays up to 100 Fault Manager messages, with the most recent on top. Click the header tab to change the default sort order. High severity messages are shown in red, low severity messages are shown in black, and Recovered messages are shown in green. When SAP ASE receives a new Fault Manager message, it automatically switches to the Service Components screen to display the message.
- *Log Records* – Displays information based on the % of Free Transaction Log, Log Records Scanned, and Log Records Processed KPIs.
- *Throughput* – Displays information based on the level of throughput.
- *Backlog* – Displays information based on the ASE Log Backlog, Primary RS Backlog, and Remote RS Backlog KPIs.
- *Latency* – Displays information based on the PDA to EXEC Latency, EXEC to DIST Latency, DIST to DSI Latency, and DSI to RDB Latency KPIs.

SAP ASE Cockpit collects these alerts for each RMA agent or server:

- State of primary Replication Server
- State of remote Replication Server
- State of Fault Manager
- State of RMA
- Failover initiated
- Synchronous to asynchronous replication
- Suspended replication
- Down replication
- Fault Manager hibernate
- Primary SAP Host Agent cannot be contacted
- Standby SAP Host Agent cannot be contacted

SAP ASE Cockpit collects these alerts for each replication path:

- Overall Ticket Latency (ms)
- Percentage of Free Transaction Log Space
- Primary Replication Queue Backlog (MB)
- Percentage of Device Usage in Primary RS
- Remote Replication Queue Backlog (MB)
- Percentage of Device Usage in Remote RS
- State of a replication path

Explore the HADR System with the SAP ASE Cockpit

The *Explore* tab allows you to administer objects in SAP ASE.

Select **ASE Servers** > `<server_name>` > **Manage Disaster Recovery** from the drop-down list to view the options for administering the HADR System.

See the online help for more information.

Viewing Alerts in the HADR System with the SAP ASE Cockpit

The **Alerts** tab is a visual display of current and previous alerts issued against the system.

From this screen, you can also configure the SAP ASE Cockpit to notify you of incoming alerts, configure how often the SAP ASE Cockpit scans for alerts, and set the alert thresholds.

See the online help for more information.

8.7 Suspending, Resuming, Enabling, and Disabling Databases

Use the RMA command to suspend, resume, enable, and disable databases.

For example, use them for:

- Long-term site failure
- Reestablishing a connection to HADR

Suspending and Resuming Databases

See [sap_suspend_replication \[page 544\]](#) and [sap_resume_replication \[page 492\]](#). However, these commands only work for planned suspend and resume activities (for example, when using the `sap_configure_rs` command). If a DSI connection suspends due to an error, you may need to use the native Replication Server `resume` because `sap_resume` does not support system transactions (DDL) and skip capabilities.

Disabling and Enabling Databases

Disabling replication in an HADR environment requires that RMA stops the Replication Agent and disables the secondary truncation point.

Issue this command to disable replication from the active Replication Server to the companion server inside HADR system:

```
sap_disable_replication <primary_logical_host_name>,  
<companion_logical_host_name>, [<dbname>]
```

Indicates the database name if you want to disable the replication only for a specific database. Removes the database name if you want to disable the replication for the whole server. See [sap_disable_replication \[page 459\]](#).

This command stops the `capture` in the active Replication Server and disables its truncation point in SPQ so that it can be truncated properly. However, this command may result in data loss, and you must rematerialize the standby database before `capture` restarts.

To enable replication in an HADR environment, enter this at the command line:

```
sap_enable_replication <primary_logical_host_name>,  
<companion_logical_host_name>, <dbname>
```

Indicates the database name if you want to enable the replication only for a specific database. Removes the database name if you want to enable the replication for the whole server. See [sap_enable_replication \[page 466\]](#)

See [Adding Databases from the Command Line After Installation \[page 299\]](#), or see *Manage SAP ASE > Always-On (HADR) Option > Suspend Replication and Resume Replication* in the SAP ASE Cockpit documentation for information about performing this task in the SAP ASE Cockpit.

8.8 Starting and Stopping the HADR System

The components in the HADR system must be shutdown and started in a ordered sequence.

Note

Shutting down the standby server may require you to restart the data server interface (DSI) thread in Replication Server.

Shut down the HADR applications in this sequence:

1. Fault Manager. On the host running the Fault Manager, source the `<installation_directory>/SYBASE.csh` (`SYBASE.sh` on the Korn shell or `SYBASE.bat` on Windows) file, and issue:
 - (UNIX) – `<Fault_Manager_install_dir>/FaultManager/bin/sybdbfm stop`
 - (Windows) – `<Fault_Manager_install_dir>\FaultManager\bin\sybdbfm.exe stop`
2. SAP ASE Cockpit on both hosts. If the cockpit is running, issue:
 - In the foreground – At the `cockpit>` prompt, execute:

```
shutdown
```

- In the background:
 - (UNIX) – `$SYBASE/COCKPIT-4/bin/cockpit.sh --stop`
 - (Windows) – `net stop "Cockpit 4.0"`
- 3. Backup Server on both hosts. Log into SAP ASE on both hosts using `isql` and issue :

```
shutdown SYB_BACKUP
```

4. Primary SAP ASE. Log into SAP ASE using `isql` and issue:

```
sp_hadr_admin deactivate, '30', '<timeout_period>'  
go  
shutdown  
go
```

If the server can't be deactivated or has undrained transaction logs, gracefully shut it down:

```
shutdown with wait nowait_hadr  
go
```

5. Active (companion site) Replication Server. Log into Replication Server using `isql` and issue:

```
shutdown
go
```

6. RMA on both hosts. Log into the RMA using `isql` and issue:

```
shutdown
go
```

7. Standby SAP ASE. Log into SAP ASE using `isql` and issue:

```
shutdown
go
```

8. Inactive (primary site) Replication Server. Log into Replication Server using `isql` and issue:

```
shutdown
go
```

Start up the HADR applications in this sequence (all commands are issued from the command line):

1. Standby SAP ASE. Issue:

- (UNIX) – `$(SYBASE)/$(SYBASE_ASE)/install/RUN_<server_name>`

For example:

```
$(SYBASE)/$(SYBASE_ASE)/install/RUN_SJSAP2
```

- (Windows) – `$(SYBASE%\SYBASE_ASE%\install\RUN_<server_name>.bat`

For example:

```
$(SYBASE%\SYBASE_ASE%\install\RUN_SJSAP2.bat
```

i Note

If SAP ASE is configured as a Windows service, start the service instead.

2. Active (companion site) Replication Server. Issue:

- (UNIX) – `$(SYBASE)/DM/<cid>_REP_<logical_site_name>/`

`RUN_<cid>_REP_<logical_site_name>.sh`

For example:

```
$(SYBASE)/DM/AS1_REP_SJHADR2/RUN_AS1_REP_SJHADR2.sh
```

- (Windows) – `$(SYBASE%\DM\cid_REP_logical_site_name\<cid>_REP_<logical_site_name>\`
`RUN_<cid>_REP_<logical_site_name>.bat`

For example:

```
$(SYBASE%\DM\AS1_REP_SJHADR2\RUN_AS1_REP_SJHADR2.bat
```

i Note

If SAP ASE is configured as a Windows service, start the service instead.

3. Primary SAP ASE. Issue:

- (UNIX) – `$(SYBASE)/$(SYBASE_ASE)/install/RUN_<server_name>`

For example:

```
$SYBASE/$SYBASE_ASE/install/RUN_SFSAP1
```

- (Windows) – %SYBASE%\%SYBASE_ASE%\install\RUN_<server_name>.bat

For example:

```
%SYBASE%\%SYBASE_ASE%\install\RUN_SFSAP1.bat
```

i Note

If SAP ASE is configured as a Windows service, start the service instead.

```
$SYBASE/$SYBASE_ASE/install/RUN_<server_name>
```

4. Backup Server on both hosts. Issue:

- (UNIX) – \$SYBASE/\$SYBASE_ASE/install/RUN_<server_name>_BS

For example:

```
$SYBASE/$SYBASE_ASE/install/RUN_SFSAP1_BS
```

- (Windows) – %SYBASE%\%SYBASE_ASE%\install\RUN_<Backup_server_name>.bat

For example:

```
%SYBASE%\%SYBASE_ASE%\install\RUN_SFSAP1_BS.bat
```

5. Inactive (primary site) Replication Server. Issue:

- (UNIX) – \$SYBASE/DM/<cid>_REP_<logical_site_name>/
RUN_<cid>_REP_<logical_site_name>.sh

For example:

```
$SYBASE/DM/AS1_REP_SJHADR2/RUN_AS1_REP_SFHDR1.sh
```

- (Windows) – %SYBASE%\DM\cid_REP_logical_site_name\<cid>_REP_<logical_site_name>
\RUN_<cid>_REP_<logical_site_name>.bat

For example:

```
%SYBASE%\DM\AS1_REP_SJHADR2\RUN_AS1_REP_SFHDR1.bat
```

i Note

If SAP ASE is configured as a Windows service, start the service instead.

6. RMA on both hosts. Issue:

- (UNIX) – \$SYBASE/\$SYBASE_ASE/bin/rma
- (Windows) – Start the RMA Windows service by either method, where <cluster_ID> is the ID of the cluster:
 - Start *Sybase DR Agent* - <cluster_ID> from the Services panel, or
 - Issue this command:

```
net start SybaseDRAgent_<cluster_ID>
```

7. SAP ASE Cockpit on both hosts. From \$SYBASE/Cockpit-4/bin, issue:

- In the foreground – from the command line, issue:

- (UNIX) `-$SYBASE/COCKPIT-4/bin/cockpit.sh`
- (Windows) `-%SYBASE%\COCKPIT-4\bin\cockpit.bat`
- In the background – At the UNIX command line. From the Bourne shell (sh) or Bash, issue:

```
nohup ./cockpit.sh 2>&1 > cockpit-console.out &
```

From the C shell, issue:

```
nohup ./cockpit.sh >& cockpit-console.out &
```

8. Fault Manager. On the host running the Fault Manager, issue:

- (UNIX) `<Fault_Manager_install_dir>`

For example:

```
/work/FaultManager/sybdbfm_AS1
```

- (Windows) – `<Fault_Manager_install_dir>\FaultManager\sybdbfm_<CID>.bat`

For example:

```
<Fault_Manager_install_dir>\FaultManager\sybdbfm_AS1.bat
```

9. Reconnect to the HADR system using the Cockpit. For example:

```
https://SFMACHINE1:4283\cockpit\#
```

8.9 Unplanned Failover

An unplanned failover is required when the primary SAP ASE is unavailable. Unplanned failover steps are similar to the planned failover steps. The main difference is timing and planning for the possibility of data loss.

The database administrator (DBA) must make a conscious decision about when to fail over, how to handle any lost transactions, and what happens to these transactions when the original primary site comes back online.

To ensure that failover to the standby site can proceed when the primary site SAP ASE is not available, use the unplanned option with the `sap_failover` command:

```
sap_failover <primary_logical_host_name>, <standby_logical_host_name>,  
<deactivate_timeout> [,unplanned | 'null' [,<drain_timeout>]]  
go
```

Where:

- `deactivate_timeout` – is the number of seconds `sap_failover` waits while the failover process drains the transaction log and SAP Replication Server queues, and waits for all in-flight data to finish replicating. If the timeout is reached, the process terminates. You cannot specify the `force` option in an unplanned failover because the primary SAP ASE is not available and cannot be deactivated.
- `drain_timeout` – (optional) is the number of seconds the process waits while draining the transaction log from the primary SAP ASE to Replication Server. If the timeout is reached, the process terminates. If not set, the timeout equals to the value of `deactivate_timeout` by default.

Once `sap_failover` completes successfully, applications can operate with the former standby database which now runs as the new primary database.

When an unplanned failover occurs, the former standby SAP ASE becomes the new primary SAP ASE. However, depending on the synchronization state, the former standby SAP ASE may or may not contain the same data the former primary SAP ASE contained. If, at the time of failover, the environment is in the:

- Asynchronous replication state, or if the primary SAP Replication Server also failed during the event – the former primary SAP ASE data is lost and you must rematerialize the former primary SAP ASE to match the content of the new primary (former standby) SAP ASE.
- Synchronous replication state – the primary and standby SAP ASE contain the same data and you need not rematerialize the former primary SAP ASE.

To determine the replication synchronization mode and synchronization state of the replication path, execute the `sap_status path` command.

i Note

During an unplanned failover, it is important to know the synchronization state of the environment at the time the failover is performed. It is the state and not the requested mode of synchronization that determines whether data loss is possible if the failover is executed. See [sap_status path \[page 525\]](#) for a description of synchronization mode and synchronization state.

8.9.1 Performing an Unplanned Failover in the Synchronous Replication State

If the environment was in the synchronous replication state at the time of the failover, you need not rematerialize the former primary SAP ASE, as the primary and standby SAP ASE contain the same data and therefore, you can reconnect the former primary SAP ASE as the standby SAP ASE.

Procedure

1. Verify the replication synchronization state of the replication path at the time failure occurred is the synchronous replication state.

Enter:

```
sap_status path
go
```

If the replication path is in a synchronous or near-synchronous replication state, you see an output similar to this:

PATH	NAME	VALUE	INFO
Site1	HADR Status	Primary : Active	Identify the primary and standby sites.
Site1	Synchronization Mode	Synchronous	The configured Synchronization Mode value.
Site1	Synchronization State	Synchronous	Synchronization Mode in which replication is currently

```
operating.
```

If the Synchronization State is Asynchronous and a failover occurs, there is a risk of data loss because not all data from the primary is guaranteed to have reached the standby SAP ASE. To guarantee the databases are synchronized when the primary SAP ASE returns to service, rematerialize the primary SAP ASE.

If the Synchronization State is Synchronous, all data from the primary SAP ASE should have been applied to the standby SAP ASE. Rematerialization is not required.

2. Execute `sap_failover` with the `unplanned` option:

```
sap_failover <primary_logical_host_name>, <standby_logical_host_name>,  
<timeout>, unplanned  
go
```

i Note

The SAP ASE running on the standby site is now the primary companion, and applications can connect to it.

3. Wait for both the former primary SAP ASE and the former primary SAP Replication Server to start and become available and ensure that all servers in the HADR system are available for replication.
4. Reconfigure the former primary SAP ASE database as the new standby for the activity occurring at the former standby SAP ASE database site:

```
sap_host_available <primary_logical_host_name>  
go
```

8.9.2 Performing an Unplanned Failover in the Asynchronous Replication State

If the environment was in the asynchronous replication state at the time of the failover or the primary SAP ASE terminated, there may be data loss, so you may need to rematerialize the former primary SAP ASE as the primary and standby SAP ASE do not contain the same data. Make a careful, planned decision for failover because there is potential for data loss.

Procedure

1. Verify the replication synchronization state of the replication path at the time failure occurred is the asynchronous replication state.

Enter:

```
sap_status path  
go
```

If the replication path is in the asynchronous replication state, you see an output similar to this:

PATH	NAME	VALUE	INFO
-----	-----	-----	-----
-----	-----	-----	-----

Site1 HADR Status standby	Primary : Active	Identify the primary and standby sites.
Site1 Synchronization Mode	Synchronous	The configured Synchronization Mode value.
Site1 Synchronization State	Asynchronous	Synchronization Mode in which replication is currently operating.

If the Synchronization State is Synchronous or Near Synchronous, all data from the primary SAP ASE should have been applied to the standby SAP ASE. Rematerialization is not required.

If the Synchronization State is Asynchronous and a failover occurs, there is a risk of data loss because not all data from the primary is guaranteed to have reached the standby SAP ASE. To guarantee the databases are synchronized when the primary SAP ASE returns to service, you must rematerialize the primary SAP ASE.

- Execute `sap_failover` with the `unplanned` option:

```
sap_failover <primary_logical_host_name>, <standby_logical_host_name>,
<timeout>, unplanned
go
```

- Wait for both the former primary SAP ASE and the former primary SAP Replication Server to start and become available and ensure that all servers in the HADR system are available for replication.
- Reconfigure the former primary SAP ASE database as the new standby for the activity occurring at the former standby SAP ASE database site:

```
sap_host_available <primary_logical_host_name>
go
```

- Stop replication from the former standby SAP ASE (current primary) to the former primary SAP ASE (current standby):

```
sap_disable_replication <standby_logical_host_name>
go
```

- Reset replication from the former standby SAP ASE (current primary) to the former primary SAP ASE (current standby):

```
sap_enable_replication <standby_logical_host_name>
go
```

This is necessary to prepare for rematerialization of the former primary SAP ASE.

- Rematerialize the databases from the current primary site to the former primary site.

Ensure that you materialize from the current primary SAP ASE, which you defined earlier during HADR system setup as `<standby_logical_host_name>`, to the former primary SAP ASE, which you defined earlier during HADR system setup as `<primary_logical_host_name>`.

8.9.3 Recovering the Primary Data Server If SAP Replication Server is Unavailable

If SAP Replication Server is unavailable during an SAP ASE startup after an unplanned failover, use SAP ASE commands to recover a database that is enabled for synchronous replication, and make it accessible online.

Context

If the replication mode is synchronous for the primary data server and SAP Replication Server is unavailable during SAP ASE startup after an unplanned failover, SAP ASE cannot recover the original primary data server and make it assume the role of a standby data server, since SAP ASE cannot connect to SAP Replication Server to obtain information about the last transaction that arrived at SAP Replication Server. For example, if the database name is D01 and <dbid> represents the database ID, in the SAP ASE error log you see:

```
Error: 9696, Severity: 17, State: 1
Recovery failed to connect to the SAP Replication Server to get the last oqid for
database 'D01'.
Database 'D01' (dbid <dbid>): Recovery failed.
Check the ASE errorlog for further information as to the cause.
```

Procedure

1. Check the SAP ASE error log to see if the latest attempt to connect to the SAP Replication Server failed.
2. Verify that the original primary database has not been recovered.
For example, if the database name is D01, log in to `isql` and enter:

```
use D01
go
```

If the database has not been recovered, you see:

```
Msg 921, Level 14, State 1:
Server 'paris', Line 1:
Database 'D01' has not been recovered yet - please wait and try again.
```

3. Use `isql` to connect to SAP Replication Server and execute `admin who` to verify that SAP Replication Server has restarted and is receiving connections.
4. In SAP ASE, enable trace flag 3604 to log all events and any errors that occur during database recovery:

```
dbcc traceon(3604)
go
```

5. Recover D01 and make it available for online access:

```
dbcc dbrecover(D01)
go
```

The recovery is successful and the database is accessible online if you see the events logged by the trace flag ending with:

```
...
Recovery complete.
Database 'D01' is now online.
DBCC execution completed. If DBCC printed error messages, contact a user with
System Administrator (SA) role.
```

6. Verify that the database is recovered and can be accessed:

```
use D01
go
```

8.9.4 (Optional) Restarting the Primary Data Server Without Synchronization

You can use the `--recover-syncrep-no-connect` option to restart the primary data server without synchronization to SAP Replication Server if you cannot restart SAP Replication Server during SAP ASE startup after an unplanned failover.

Context

When you use the `--recover-syncrep-no-connect` option for the SAP ASE `dataserver` executable, SAP ASE starts and tries to connect to the SAP Replication Server during recovery. If the connection attempts to SAP Replication Server fail, there is no 9696 error. SAP ASE recovers the databases but the primary and standby databases may not be synchronized. Without synchronized replication between the databases and SAP Replication Server, you cannot recover from an unplanned failover with the assurance of no data loss that synchronized replication provides.

Procedure

1. Insert the `dataserver` executable with the `--recover-syncrep-no-connect` option in the `run_server` file to restart the primary data server without synchronization:

```
...
dataserver --recover-syncrep-no-connect
...
```

2. Execute the `run_server` file.

8.9.5 Recovery for MultiDatabase Transactions

After an unplanned failover for a database that is enabled for synchronous replication that participates in a multidatabase transaction, the recovery process may not apply changes to all the databases involved in the multidatabase transaction.

A multidatabase transaction, also known as a cross-database transaction, is a single transaction that affects tables in different databases where the coordinating database is the database where the transaction started, and the subordinate databases are the other databases affected by the transaction.

During recovery after an unplanned failover in an HADR system where the replication synchronization mode is synchronous, only the primary SAP ASE database, which is configured for synchronous replication, rolls back transactions that have not been stored in the simple persistent queue (SPQ) of the SAP Replication Server. However, if there is a multidatabase transaction, only the database that is enabled for synchronous replication rolls back the transaction for a multidatabase transaction. The other databases participating in a multidatabase transaction do not roll back the transaction. Therefore after recovery, the changes due to the multidatabase transaction will not have been applied across all the participating databases.

During recovery after an unplanned failover, SAP Replication Server assumes the coordinator role that was previously performed by the coordinating database where the multidatabase transaction started.

After recovery from an unplanned failover where multidatabase transactions have been applied, the status of the multidatabase transactions depends on whether the replication mode is synchronous for the coordinating database or one of the subordinate databases. See the SAP ASE error log for the status of the transaction and decide if you want to manually apply or roll back changes in one of the databases.

- If the coordinating database replication synchronization mode is not synchronous and the subordinate database replication synchronization mode is:
 - Synchronous – SAP ASE may roll back the changes in the subordinate database during recovery if the changes are newer than the last transaction received by SAP Replication Server and written to SPQ. You see this information in the SAP ASE error log:

```
Synchronously replicated multidatabase subordinate transaction (<page>,
<row>)
  in database '<dbname>' forcibly rolled back; however, the coordinating
  transaction may have committed.
```

- Not synchronous – SAP ASE does not roll back the changes in the subordinate database since the replication mode is not synchronous, and therefore the transaction behavior is the same as in a non-synchronous replication mode-enabled system with multidatabase transaction.
- If the coordinating database replication synchronization mode is synchronous, SAP ASE may roll back the changes in the coordinating database during recovery if SAP ASE did not receive acknowledgment that SAP Replication Server has received and written the transaction to the SPQ. However, the rollback of transactions in the subordinate databases depends on the order in which the databases are recovered. When the HADR system recovers the coordinating database:

- Before the subordinate – SAP ASE also rolls back the changes in the subordinate database.
- After the subordinate – SAP ASE does not roll back the changes in the subordinate database. You see this information in the error log:

```
Synchronously replicated multidatabase coordinating transaction (<page>,
<row>)
  in database '<dbname>' forcibly rolled back; however, subordinate
  transaction
  branches may have committed.
```

To analyze the events that occurred during recovery for an HADR system that is enabled for synchronous replication, check for the "SyncRep undone" status in the ENDXACT log records that SAP ASE writes if there is any transaction rollback during recovery after failover. For example:

```
ENDXACT          (1822,24)          sessionid=1822,20
attcnt=1 rno=24 op=30 padlen=4 len=32
odc_stat=0x0000 (0x0000)
loh_status: 0x0 (0x00000000)
BEGINXACT=(1822,20) endstat=XEND_ABORT time=Nov 29 2013 2:16:17:636PM
xstat=0x42 [Written by recovery, SyncRep undone]
```

8.10 The Split-Brain Check: Preventing Multiple Primary Servers

Only one server in the HADR group should perform client transactions at any one time. If more than one server assumes the role of the primary server, the databases on the HADR servers can no longer be synchronized, and the system enters a "split-brain" situation.

The HADR system provides a check against this, which is performed either at start-up if the SAP ASE configuration file instructs the server to start as a primary, or when you use `sp_hadr_admin primary` to manually promote a standby server to the primary server.

The check connects to and queries each configured HADR member. If a remote HADR member in the group is identified as an existing primary server, the check does not allow the local server to be promoted to the primary server. Generally, you cannot override this check.

If the check fails to connect to one or more remote HADR member, it assumes that the unreachable member may be a primary server, and refuses to promote the local server to primary. In this situation, you can use the `force` parameter to override the split-brain check:

```
sp_hadr_admin primary, force
```

Before using the `force` parameter, verify that there is no other primary server present in the group.

8.11 Planned and Unplanned Failovers

The HADR system includes planned manual or automatic unplanned failovers.

i Note

If your site is configured for Fault Manager, it handles primary SAP ASE failure, and automatic failover is triggered when it is safe to failover. If the Fault Manager detects potential data loss when fail over is triggered, you must manually intervene to restore the old primary site or accept data loss and promote the companion SAP ASE as the new primary SAP ASE. The steps described here are applicable if Fault Manger is not configured and the database administrator must decide how to recover from an unplanned failover.

If failover fails before the new primary SAP ASE is activated, RMA attempts to set the old primary as primary again. A failure after this point requires you to manually activate the new primary, start Replication Agent on the new primary, then execute `sap_host_available` there when the new standby running.

Manual Failover

A planned failover occurs when you intend to perform a task that requires a node to be brought down. You can perform manual failover from the command line or from SAP ASE Cockpit.

From the command line

1. Connect to the primary or companion RMA and issue `sap_failover`. This example uses a deactivation timeout of 60 seconds:

```
sap_failover SFHADR1,SJHADR2,60
```

RMA issues messages similar to:

```
TASKNAME      TYPE
VALUE
-----
-----
Failover      Start Time          Thu Dec 03 20:03:14 UTC
2015
Failover      Elapsed Time
00:00:02
DRExecutorImpl Task Name
Failover
DRExecutorImpl Task State
Running
DRExecutorImpl Short Description      Failover makes the current standby
ASE as the primary server.
DRExecutorImpl Long Description      Started task 'Failover'
asynchronously.
DRExecutorImpl Additional Info       Please execute command 'sap_status
task' to determine when task 'Failover' is complete.
Failover      Task Name
Failover
Failover      Task State
Running
Failover      Short Description      Failover makes the current standby
ASE as the primary server.
Failover      Long Description      Waiting 3 seconds: Waiting for the
end of data marker for database 'master' to be received.
Failover      Current Task Number
6
Failover      Total Number of Tasks
18
Failover      Task Start          Thu Dec 03 20:03:14 UTC
2015
Failover      Hostname            SFMACHINE1
```

`sap_failover` is an asynchronous command, and must complete before you perform the next step. You cannot perform two `sap_failover` commands in parallel. That is, the first `sap_failover` command must complete before you issue a second.

2. Connect to the primary or companion RMA and issue `sap_status` to check the status of the `sap_failover` command:

```
sap_status task
```

RMA issues messages similar to this when the failover task is finished (see the bold text):

```
TASKNAME      TYPE
VALUE
-----
-----
Status        Start Time          Thu Dec 03 20:03:14 UTC
2015

Status        Elapsed Time
00:00:09

Failover      Task Name
Failover

Failover      Task State
Completed

Failover      Short Description    Failover makes the current standby ASE as
the primary
server.

Failover      Long Description      Failover from source 'SFHADR1' to target
'SJHADR2' is complete. The target may be
unquiesced.
Failover      Additional Info      Please run command 'sap_host_available
SFHADR1' to complete disabling replication from the old source, now that the
target 'SJHADR2' is the new primary.
Failover      Current Task Number
14

Failover      Total Number of Tasks
14

Failover      Task Start           Thu Dec 03 20:03:14 UTC
2015

Failover      Task End             Thu Dec 03 20:03:23 UTC
2015

Failover      Hostname             SFMACHINE1
```

3. Once `sap_failover` completes, connect to the SAP ASE at the new primary site and use the `hadr_mode` and `hadr_state` functions to confirm that its HADR mode and state is now primary-active:

```
select hadr_mode(), hadr_state()
go
-----
Primary          Active
```

Log in to the old primary server and verify that its mode and state is standby inactive:

```
select hadr_mode(), hadr_state()
go
-----
Standby         Inactive
```

Alternatively, you can connect to the RMA on the primary companion and issue (see the bold text):

```

sap_status path
PATH          NAME          VALUE
INFO
-----
. . .

SJHADR2      HADR Status      Primary : Active      Identify
the primary and standby sites.
SJHADR2      Synchronization Mode Synchronous      The
configured Synchronization Mode value.
. . .
SFHADR1      HADR Status      Standby : Inactive    Identify
the primary and standby sites.
. . .

```

- (Optional) Stop the Fault Manager if it is configured to restart SAP ASE, Replication Server and RMA. Configuring the Fault Manager for unplanned failover and a subsequent automatic restart of these components can trigger actions that are undesirable during planned failover. Consequently, you should stop the Fault Manager during any planned activity. From the <install_directory>/FaultManager directory, issue:

```
<Fault_Manager_install_dir>/FaultManager/bin/sybdbfm stop
```

i Note

If you configured a floating IP in the Fault Manager (ha/syb/support_floating_ip = 1) and stopped the Fault Manager before issuing the `sap_failover` command, you need to move the floating IP manually. See SAP Note [2709394](#) for more information.

- After `sap_failover` successfully completes, it prints a message indicating that you must run `sap_host_available`. Issue this command from RMA to clean and disable the old replication path and activate the new direction for the replication path:

```

sap_host_available SFHADR1
TASKNAME          TYPE          VALUE
-----
-----
HostAvailable     Start Time      Thu Dec 03 23:48:34 UTC 2015
HostAvailable     Elapsed Time    00:01:24
HostAvailable     Task Name       HostAvailable
HostAvailable     Task State      Completed
HostAvailable     Short Description Resets the original source logical
host when it is available after fail over.
HostAvailable     Long Description Completed the reset process of
logical host 'SFHADR1' receiving replication from logical host 'SJHADR2'.
HostAvailable     Current Task Number 8
HostAvailable     Total Number of Tasks 8
HostAvailable     Task Start      Thu Dec 03 23:48:34 UTC 2015
HostAvailable     Task End        Thu Dec 03 23:49:58 UTC 2015
HostAvailable     Hostname        SFMACHINE1

```

- Confirm that replication is active from the SAP ASE Cockpit, or from the RMA by issuing:

```

sap_status path
PATH          NAME          VALUE
INFO
-----
-----

```



```

. . .
SJHADR2.SFHADR1.AS1      State              Active
                          Path is active and replication can occur.
. . .

```

Alternatively, you can confirm the direction of replication from the SAP ASE Cockpit.

7. (Optional) If the Fault Manager is stopped, restart it. From the `<install_directory>/FaultManager` directory, issue:

```
<Fault_Manager_install_dir>/FaultManager/sybdbfm_<CID>
```

From SAP ASE Cockpit

Follow the instructions in *Manage SAP ASE > Always-On (HADR) Option > Performing a Planned Failover* in the SAP ASE Cockpit documentation.

Automatic Failover

An unplanned, automatic failover occurs when an event causes a node to go down. The standby server is automatically promoted to the primary position with an internally executed `sap_failover` command.

1. Check SAP ASE Cockpit for alerts indicating fault detection, failover initiation, and failover completion.
2. Connect to the standby server and issue the `hadr_mode` and `hadr_state` functions to confirm that its HADR mode and state is now primary-active:

```

select hadr_mode(), hadr_state()
go
-----
Primary              Active

```

Alternatively, you can connect to the RMA on the primary companion and issue:

```
sap_status path
```

You can also use `sap_status task` from RMA to display the progress of the `sap_failover` command. Once the failover is complete, the SAP ASE Cockpit indicates that the SAP ASE running on the first site is now the primary server.

3. Once the server on the primary site is back online, issue this command to clean and disable the old replication path and activate the new direction for the replication path:

```
sap_host_available SFHADR1
```

4. Confirm that replication is active from the SAP ASE Cockpit, or from the RMA by issuing:

```

sap_status path
PATH          NAME          VALUE
INFO
-----
. . .
SetupReplication  Task          Long Description          Setting up
replication between the two ASE hosts 'SFHADR1' and 'SJHADR2' completed
successfully. Databases on 'SJHADR2' are now ready to be materialized.
. . .

```

8.11.1 Verifying the Environment After a Failover

Use the SAP ASE Cockpit Monitoring dashboard to verify the HADR environment after failover.

Check that:

- The standby server becomes the primary server (with a yellow-colored box).
- The primary server becomes the standby server (with a gray-colored box).
- A green connector joins the primary and standby server. If the connector is not green, check for an error condition.
- The text and icon for Service Component Status are green and *Active*.
- The text and icon for Replication Paths Status are green and *Active*.

8.12 Member Modes and States

Each server in an HADR configuration has an external mode that is visible to and known by other HADR members, as well as an internal state that is known only by the member.

External modes are:

- Primary – the member of an HADR configuration on which active transaction processing by user applications is allowed to take place.
- Standby – the member of an HADR configuration that contains copies of specific databases that originate on the primary member, and is available to take over transaction processing if the primary member fails. Replication Server replicates database changes on the primary that are marked for replication to standby members.
- Disabled – HADR is disabled on this member.
- Unreachable – the local member (the server from which you enter commands) cannot reach this remote HADR member.
- Starting – HADR member is starting.

Internal states are:

- Active – the normal state of the primary server. Transaction activity by user applications occurs in this state without any of the restrictions imposed by other states. The primary server is always in the active state after a restart, provided the split-brain check is successful.
- Inactive – restricts all activity by unprivileged connections, but imposes no restrictions on activity by privileged connections (including starting new transactions). Deactivating a primary server places it into the inactive state. Standby members are always in the inactive state.
- Deactivating – intermediate state between active and inactive states. The HADR system gracefully terminates transactions by unprivileged connections during the deactivating state. Use the `sp_hadr_admin timeout` parameter to specify a time limit for this state. The HADR system waits for the time specified by the `timeout` parameter for the transactions to complete. When the `timeout` period ends, the internal state is either rolled back to active or is advanced to inactive, depending on whether you included the `force` parameter with the `sp_hadr_admin` command.

i Note

When you include the `force` parameter, the HADR system forcibly terminates all the transactions started by privileged and unprivileged connections. Unprivileged connections cannot start new transactions when the server is in the deactivating state.

The internal state of a primary server is not preserved across restarts. However, the external mode is saved across restarts using the `HADR_mode` configuration parameter.

Use the `HADR_primary_check_frequency` configuration parameter to determine how often the standby server checks the primary server's mode and state.

If the standby server detects that the other server is not in the primary mode and an active state, it introduces a delay before sending the address list used for connection redirection. The length of this delay is determined by the `HADR_login_stall_time` configuration parameter. See the *Reference Manual: Configuration Parameters*.

8.12.1 Determining the Member Mode and State

There are a number of ways to determine the member's mode and state.

- Use the `hadr_mode` function and the `<@@hadr_mode>` global variable to determine the member mode. The return values for `<@@hadr_mode>` and `hadr_mode` are:

Return Value:	Description
-1	HADR is disabled.
0	HADR is enabled. This is a standby server.
1	HADR is enabled. This is a primary server.
2	HADR is enabled, but the server is unreachable. This value is not seen by the local server.
3	HADR is enabled, and the server is ready for initialization.

- Use the `hadr_state` function and the `<@@hadr_state>` global variable to determine the member state. The return values for `<@@hadr_state>` and `hadr_state` are:

Return Value:	Description
0	HADR is in an unknown state. This value is typically returned when HADR is in disabled mode.
1	The primary server allows transaction processing from user applications.
2	The server is inactive, and does not allow transaction processing from user applications.
3	The server is changing from the active to the inactive state, and the log is being drained. Eventually, the mode should transition to the inactive state. If deactivation times out, the mode may switch back to the active state.

- You can include a return value (-1, 0, 1, 2, or 3) as an input parameter with `hadr_mode` and `hadr_state` functions to determine the state this return value represents (this is the same verbose information that `<@@hadr_mode>` and `<@@hadr_state>` return). For example:

```
select hadr_mode(1)
-----
Primary
```

- Issuing `hadr_mode` and `hadr_state` functions without arguments returns the mode and state of the server, respectively:

```
select hadr_mode(), hadr_state()
-----
Primary          Active
(1 row affected)
```

- Issue the `HADR mode` configuration parameter to determine the current mode of the server (the server below is in non-HADR mode):

```
sp_configure 'HADR mode'
Parameter Name      Memory Used      Default      Config Value      Type
Run Value          Unit
-----
HADR mode          0                -1           -1                dynamic
(1 row affected)
(return status = 0)
```

`HADR mode` returns these values:

- 1 – server is not configured for HADR.
- 0 – server is configured as a standby.
- 1 – server is configured as a primary.
- 5 – server is configured as a disaster recovery standby.

See the *Reference Manual: Configuration Parameters*.

- You can also use the `sp_hadr_admin mode` and `sp_hadr_admin state` parameters to determine the server's mode and state. For example, this shows the server's mode as "Primary" and its state as "Active":

```
sp_hadr_admin mode
HADR Mode
-----
Primary
(1 row affected)
Possible values for HADR Mode are: 'NoHADR (-1)', 'Primary (1)', 'Standby
(0)', 'Unreachable (2)' and 'Starting (3)'
(return status = 0)
sp_hadr_admin state
HADR State
-----
Active
(1 row affected)
Possible values for HADR State are: 'Unknown (0)', 'Active (1)', 'Inactive (2)'
and 'Deactivating (3)'
(return status = 0)
```

Changing a Server's Mode

Use `sap_failover` to change a server's mode. See [sap_failover](#) [page 468].

8.12.2 Forcing a Server to Primary Mode

If the start-up sequence described in the previous section is not followed, and if the primary SAP ASE is started before the standby SAP ASE, it starts in standby mode due to split-brain check.

If this occurs, you can connect to the SAP ASE using the privileged login and change the state to primary. Make sure that the server you intend to promote to primary was indeed primary earlier by checking its SAP ASE log file.

To promote the server to primary, connect to it and issue:

```
sp_hadr_admin primary
go
sp_hadr_admin activate
go
```

Issue the `sp_start_rep_agent` system procedure in the master, CID, and in each database that participates in HADR. For example:

```
use master
go
sp_start_rep_agent master
go
use <CID>
go
sp_start_rep_agent <CID>
go
use <user_database_1>
go
sp_start_rep_agent <user_database_1>
go
use <user_database_2>
go
sp_start_rep_agent <user_database_2>
go
. . .
```

8.13 Checks to Perform After an Unplanned Failover

Managing data loss and viewing the Fault Manager alerts are important checks to perform after an unplanned failover.

8.13.1 Managing Data Loss in an Unplanned Failover Operation

There are a number of steps you perform to manage data loss after an unplanned failover.

1. Verify that the replication synchronization state of the replication path, at the time failure occurred, is the synchronous replication state:

```
sap_status path
```

If the replication path is in the synchronous replication state, you see output similar to:

PATH	NAME	VALUE	INFO
SJHADR2	HADR Status	Primary : Active	Identify the primary and standby sites.
SJHADR2	Synchronization Mode	Synchronous	The configured Synchronization Mode value.
SJHADR2	Synchronization State	Synchronous	Synchronization Mode in which replication is currently operating.

If the Synchronization Mode or Synchronization State is Asynchronous and a failover occurs, there is a risk of data loss because not all data from the primary is guaranteed to have reached the standby SAP ASE. To guarantee the databases are synchronized when the primary SAP ASE returns to service, rematerialize the primary SAP ASE.

If the Synchronization State is Synchronous, all data from the primary SAP ASE should have been applied to the standby SAP ASE. Rematerialization is not required.

2. Execute `sap_failover` with the unplanned option:

```
sap_failover <primary_logical_host_name>, <standby_logical_host_name>,  
<timeout>, unplanned
```

3. Wait for both the former primary SAP ASE and the former primary SAP Replication Server to start and become available, and ensure that all servers in the HADR system are available for replication.
4. Reconfigure the former primary SAP ASE database as the new standby for the activity occurring at the former standby SAP ASE database site:

```
sap_host_available <primary_logical_host_name>
```

8.13.2 Managing Fault Manager Alerts

Fault Manager sends alerts for events requiring database administrator attention. After an unplanned failover, there are a number of alerts you check to determine if the alert notifications are still active.

These alerts include:

- Failover initiated from '<site_1>' to '<site_2>'
- Replication is changed from SYNC to ASYNC

If the first alert is still active, wait for it to clear; until it is cleared, failover has not successfully completed. If the second alert is still active after unplanned failover, performing an unplanned failover might cause data loss.

Check for messages in the *Fault Manager Messages* table in the SAP ASE Cockpit that indicate the alerts are cleared.

Go to the *Service Components* page of the HADR monitoring dashboard:

- For the first alert, watch for this message (in green), which indicates the alert stating Failover initiated from 'site1' to 'site2' is cleared:

```
Failover succeeded from '<site_1>' to '<site_2>'
```

- For the second alert, watch for this message (in green), which indicates this alert is cleared:

```
Replication is changed from ASYNC to SYNC
```

See the *SAP ASE Cockpit > Alerts in SAP ASE* for more information.

8.14 Checking the System Status

You can check the status of Replication System using the RMA commands as well as the SAP ASE Cockpit. This topic provides the information on the RMA commands required to monitor replication system status. For example, `sap_status path`, `sap_status route`, and `sap_status resource`.

- **Check Replication Path Status** – Replication paths are used to deliver data changes between the primary and the standby ASE databases. Each pair of databases (between the primary and the standby SAP ASE servers) has two replication paths defined. Check the replication path status from the primary SAP ASE server to the standby SAP ASE server to ensure that the paths are in active state. Use the `sap_status path` command to check the status of the path. See [sap_status \[page 520\]](#).
- **Check Replication Path Sub Components Status** – Each replication path consists of servers (SAP ASE and Replication Server), threads (SAP ASE RepAgent thread and Replication Server internal threads), and queues (inbound, outbound, and SPQs). `sap_status route` allows you to collect the status of these components. See [sap_status route \[page 534\]](#).
- **Check Replication Queues Information** – Use the `sap_status resource` command to check the size and device buffer usage of device buffer and SPQs. See [sap_status resource \[page 530\]](#).

8.14.1 Tracing Latency

Latency is the amount of time it takes for a transaction committed on the primary server to be applied on the standby server. Send a trace to initiate latency monitoring. You should always initiate a trace from the primary host. For example, if you are replicating from site1 to site2, initiate the trace at site1 (the primary).

Procedure

Execute the `sap_send_trace` command. If you do not specify a database name, a trace is sent to all databases for that host: master and ERP (if it exists):

```
sap_send_trace <primary_logical_host_name> [,<database_name>]
go
```

This command inserts an `rs_ticket` into the source database or databases. Latency is calculated from the most recent entry in the target database's `rs_ticket_history` table.

The `sap_status path` command calculates latency based on the most recent trace received at the standby database. For example:

```
adsite1.site2.HA2 State Active Path is active and replication can occur.
adsite1.site2.HA2 Latency Time 2015-05-16 02:11:51.106 Time latency last
calculated
adsite1.site2.HA2 Latency 70 Latency (ms)
adsite1.site2.HA2 Commit Time 2015-05-16 02:11:52.008 Time last commit replicated
```

If there is a backlog of data, the trace element reflected by `sap_status path` results may not be the most recent trace element requested. Verify that `Time latency last calculated` is the current time, and not reflective of the trace element that was executed earlier.

8.14.2 Determining the Backlog

The amount of the SAP ASE transaction log remaining to be processed by Replication Server is the backlog.

These commands indicate the amount of backlog:

- `sap_status route` – see [sap_status \[page 520\]](#)
- `sap_status resource` – see [sap_status resource \[page 530\]](#)

SAP ASE determines the amount of backlog according to these criteria:

- If the second truncation point is not set, SAP ASE returns the backlog from the beginning of the log.
- If the second truncation point is set, execute this command on the primary server to view the number of backlog pages:

```
select rep_log_pages((select db_id('<database_name>')))
```

For example:

```
select rep_log_pages((select db_id('pubs2')))
```


- Calculate the amount of SAP ASE backlog (in megabytes) with this equation:

```
backlogpage X page_size/1024/1024
```

View the SAP ASE page size with the `<@@maxpagesize>` global variable (this server uses a 4K page size):

```
select @@maxpagesize
-----
         4096
```

8.15 Evaluating the Health of an HADR Cluster

Query SAP ASE, Replication Server, RMA, and so on to determine the health of the HADR system.

Evaluating the Health of the Primary and Companion Servers

Generally, you can look at the following to determine the health of the primary and companion servers:

- Log space – issue this in the target database to determine the available space for the logs:

```
sp_spaceused syslogs
```

The `loginfo` and `lct_admin` functions also display log information. See the *SAP ASE Reference Manual: Building Blocks*.

- Replication Agent status – issue this in the target replicated database:

```
sp_help_rep_agent <database_name>, 'process'
```

- SAP ASE error log messages – review messages in the error logs.
The default location of the SAP ASE error log is `$$SYBASE/$SYBASE_ASE/install`.
- (On the companion server) Monitoring latency – use the `rs_ticket` command to view detailed information regarding latency among the replication server internal threads. See [Checking Latency with rs_ticket \[page 344\]](#).
- Query RMA – RMA collects all the information about the status of the system. For example, this query indicates the state of the primary SAP ASE Replication Agent:

```
select case
    when (Status = "sleeping" and SleepStatus = "opening stream") or
         (Status = "sleeping" and SleepStatus = "stream sleep") or
         (Status = "sleeping" and SleepStatus = " sleeping on log
full" ) then "suspect"
    else "active"
end
from master..monRepScanners
where DBID = <dbid>
union
select case when max(Status) is NULL then "down" end from
master..monRepScanners where DBID = <dbid>
go
```

Evaluating the Health of Replication Server

Generally, you can look at the following to determine the health of Replication Server:

- SPQ size and capacity – issue this command to see the SPQ size:

```
admin disk_space, spq
```

See [Troubleshooting Data That is Not Replicating \[page 420\]](#) > *SPQ is Full* for information about fixing a full SPQ and [admin disk_space, mb, spq \[page 592\]](#).

- Replication Server queue size and capacity – issue this command to determine the Replication Server queue size and capacity:

```
admin disk_space
```

Use the `sap_set simple_persistent_queue_size` command to change the queue size. See [sap_set simple_persistent_queue_size \[page 501\]](#).

- DSI status – issue this to determine the status of the DSI threads:

```
admin who
```

8.15.1 Common Failure Scenarios

There are a few situations during which failures may typically occur in the HADR system.

Situations include:

- Failing to add instances to the HADR system – see [Troubleshooting the HADR System \[page 386\]](#) > *Failure to Add an Instance*.
- RMA commands fail to run – see [Troubleshooting the HADR System \[page 386\]](#) > *RMA Command Failure*.
- Failure during setup – see [Recovering from a Failed Setup \[page 401\]](#).

8.16 Read-Only Support from Companion Node in an HADR Cluster

The HADR system prevents data changes from occurring on a companion node by ensuring that nonprivileged users cannot log into the companion node.

Only users with `allow_hadr_login` privilege can login to the companion node. Users without this who attempt to log in to the companion node are rejected or redirected to the primary node, depending on the `redirection` property on the connection. If the `redirection` property is set, the connection is redirected; if it is not set, the connection is rejected. Users with the `sso_role` and `sa_role` are granted `allow_hadr_login` privilege by default, so administrators can log in to the companion node to perform administrative tasks.

The companion node can be used for read-only access or reporting. The following example describes how to set up read-only access. Generally, you should have a separate read-only user who is allowed to log into the companion node but does not have permission to modify data.

i Note

The read-only user should not own any objects on either the primary or the standby server because object owners implicitly have all permissions on their objects, so you cannot restrict them to read-only permissions.

In this example, objects in the `pubs2` database are owned by DBO. User `pubs2user` is created with privileges to modify `pubs2`, while user `pubs2rouser` is a read-only user.

1. Create the logins:

```
use master
go
create login pubs2user with password "Sybase123"
go
create login pubs2rouser with password "Sybase123"
go
```

2. Create the roles and grant privileges (if you created a default profile, see *Manage Login Profiles* in the SAP ASE Cockpit help for additional steps):

```
create role pubs2_readonly
go
grant allow hadr login to pubs2_readonly
go
grant role pubs2_readonly to pubs2rouser
go
```

3. Alter the `pubs2rouser` to be automatically activated on login:

```
alter login pubs2rouser add auto activated roles pubs2_readonly
go
```

4. Add the `pubs2user` and `pubs2ruser` to the `pubs2` database:

```
use pubs2
go
sp_adduser pubs2user
go
New user added.
(return status = 0)
sp_adduser pubs2rouser
go
New user added.
(return status = 0)
```

5. Grant appropriate permissions on the `authors` and `publishers` tables:

```
grant all on authors to pubs2user
go
grant all on publishers to pubs2user
go
...
grant select on authors to pubs2_readonly
go
grant select on publishers to pubs2_readonly
go
```

6. (Optional) If you have granular permission enabled, issue:

```
grant select any table to pubs2_readonly
```

When logging into the server, pubs2rouser can now view and alter data:

```
% $SYBASE/$SYBASE_OCS/bin/isql64 -Q -Upubs2user -PSybase123 -w 400
use pubs2
go
select au_fname, au_lname, state, postalcode from authors where state = "UT"
go
  au_fname          au_lname          state postalcode
-----
  Anne             Ringer             UT      84152
  Albert           Ringer             UT      84152
(2 rows affected)
update authors set postalcode = "84156" where state = "UT" and au_fname = "Anne"
go
(1 row affected)
select au_fname, au_lname, state, postalcode from authors where state = "UT"
go
  au_fname          au_lname          state postalcode
-----
  Anne             Ringer             UT      84156
  Albert           Ringer             UT      84152
(2 rows affected)
update authors set postalcode = "84152" where state = "UT" and au_fname = "Anne"
go
(1 row affected)
quit
```

However, when logging into the database, pubs2ruser can view, but not alter, the data:

```
% $SYBASE/$SYBASE_OCS/bin/isql64 -Q -Upubs2ruser -PSybase123 -w 400
use pubs2
go
set role pubs2_readonly on
go
update authors set postalcode = "84156" where state = "UT" and au_fname = "Anne"
go
Msg 10330, Level 14, State 1:
Line 1:
UPDATE permission denied on object authors, database pubs2, owner dbo
select au_fname, au_lname, state, postalcode from authors where state = "UT"
go
  au_fname          au_lname          state postalcode
-----
  Anne             Ringer             UT      84152
  Albert           Ringer             UT      84152
(2 rows affected)
```

8.17 Adding SAP Replication Server Device Space

Add device space to the replication paths.

Procedure

Execute this from the RMA to add 1 GB of additional storage to replication on the primary site:

```
sap_add_device <primary_logical_host_name>, part2, 1000  
go
```

Results

The `sap_add_device` command issues an `add partition` command to the underlying SAP Replication Server defined for the requested logical host.

i Note

When you enable stream replication, Replication Server automatically creates a simple persistent queue. By default, this queue consists of two 1000 MB files, but can extend to a maximum of one hundred files (100 GB of disk space). Use the `sap_set simple_persistent_queue_max_size` command to restrict and adjust the maximum amount of disk space allocated for the simple persistent queue.

8.17.1 Scenarios That May Require Additional Device Space

Replication Servers are configured with file system space that buffer and forward replication data between the primary and standby sites. In cases of high volume or system outages, you may need to increase the space used by replication.

Typical scenarios may include:

- If the replication throughput cannot meet the current primary system demand, the SAP ASE transaction log may become full while waiting for replication. Adding buffering space to the Replication Servers allows the SAP ASE logs to truncate more frequently, pushing the data out of the SAP ASE transaction log and into replication queue storage to remove the risk of a full log in the primary SAP ASE server.
- Primary SAP ASE transaction log space may fill if the standby site is unavailable (due to either planned or unplanned downtime). The SAP ASE transaction log may become full while waiting for the standby to return to service. Adding buffering space to the Replication Servers can allow the SAP ASE logs to truncate more frequently, pushing the data instead into Replication Server queues for storage until the standby server returns to service and can accept the replication backlog.

8.18 Tuning the HADR Components

Use the `sap_tune_rs` RMA command to tune and configure the HADR components (such as the primary and standby SAP ASE, Replication Agent, and Replication Server).

See [sap_tune_rs \[page 548\]](#).

The syntax is:

```
sap_tune_rs <memory>, <number of cpus>
```

Select the `<memory>` and `<number of cpus>` values based on the primary transaction log generation rate. For example, if the primary log generation rate is 3.5 GB per hour, you can tune Replication Server to have 4 GB or memory and 2 CPUs so that latency can be less than 5 seconds. If the primary log generation rate is high at any time (greater than 5 GB but less than 12 GB), then setting memory to 8 GB and CPUs to 4 should keep the latency to less than 5 seconds.

8.18.1 Tuning Replication Agent and Replication Server

There are several ways to tune Replication Agent and SAP Replication Server to improve their overall performance.

Tuning for Peak Transaction Rate

In the Replication Agent, the default value for `peak_transaction_threshold` is 5, and that for `peak_transaction_timer` is 300 (in seconds). This means that Replication Agent switches to “async” mode if some commits take longer than 10 seconds (which is the default value for `max_commit_wait_time`) for a total of five times over a period of 300-seconds.

Such mode switches occur if the disks on which SPQ is stored have slower-capacity I/O-writes than other disks in the system. In this case, set the value of `peak_transaction_threshold` to 50, and `peak_transaction_timer` to 120. These changes should prevent the mode from switching.

- Replication impact on primary database:
 - Understanding the cost of HADR synchronous replication on source databases.
 - Setting and modifying the `max_commit_wait_time` configuration parameter.
 - Failover time requirements and their relevance to Replication Server DSI performance (limiting the backlog that needs to be drained).
- Using `rs_ticket` to collect end-to-end latency information in HADR environment. See [Checking Latency with rs_ticket \[page 344\]](#) and the *Replication Server Reference Manual > Adaptive Server Stored Procedures*.
- Analysis of end-to-end replication performance by identifying bottlenecks in:
 - RMA.
 - Replication Server SQP and Capture.

- Replication Server data server interface (DSI) and high volume adaptive replication (HVAR). See *Replication Server Administration Guide Volume 2 > Performance Tuning > Advanced Services Option > High Volume Adaptive Replication > HVAR Performance Tuning*.
- I/O and network.
- Large transaction handling (see section below).
- Corrective action for bottlenecks:
 - Tuning SAP ASE (active and standby servers. See *SAP ASE Performance and Tuning Series*).
 - Tuning Replication Server (multiple modules).
 - Large Transactions (see next section).
- Using Replication Server monitors and counters:
 - Critical SAP ASE and Replication Server counters.
 - Enabling, collection, and analyzing SAP ASE counters.
 - Enabling, collecting and analyzing Replication Server counters. See *Replication Server Administration Guide Volume 2 > Monitor Performance Using Counters*.
- Configuring Replication Server SPQ and partitions:
 - Identifying when additional SPQ and partition space is needed. See [Troubleshooting Data That is Not Replicating \[page 420\]](#).
 - Managing SAP ASE dumps. See [Loading from an External Dump \[page 302\]](#) and the *SAP ASE System Administration Guide Volume 2 > Backing Up and Restoring User Databases and Restoring the System Databases*.

Tuning for Large Transactions Handling

Use the `alter connection` command to enable the early dispatch mechanism to handle large transactions in SAP Replication Server. Tweak values of the following parameters in the command to achieve this:

- Set the value of the parameter `parallel_dist` to `on` to enable early dispatch in the SAP Replication Server. This configuration change automatically suspends and resumes the distributor. However, the distributor component waits until the `SQT` database is flushed of any transactions. The following example enables early dispatch for the `pubs2` table on the `SFSAP1` server:

```
alter connection to SFSAP1.pubs2 set parallel_dist to 'on'
```

The Replication Server error log includes messages similar to this when the distributor thread starts incorrectly:

```
some transactions are present in SQM , Distributor will not be restarted
until the transactions are flushed
```

- Set the value of the `dsi_num_threads` to 5. For example:

```
alter connection to SFSAP1.pubs2
set dsi_num_threads to '5'
```

- Set the value of the `dsi_num_large_xact_threads` parameter to 2 (to handle two large transactions received in parallel).

For example:

```
alter connection to SFSAP1.pubs2
set dsi_num_large_xact_threads to '2'
```

Note

- The value for `dsi_large_xact_threads` should be equal to the number of expected large transactions that occur in parallel.
- The value for `dsi_num_threads` is equal to the value of `dsi_large_xact_threads` plus 3. That is:

```
dsi_num_threads = dsi_large_xact_threads+3
```

- Set the value for the `dsi_xact_threads` parameter to 2. For example:

```
alter connection to SFSAP1.pubs2 set dsi_num_large_xact_threads to '2'
```

Tuning for Memory Allocation

See the *Tuning Memory Allocation* chapter in *SAP Replication Server Administration Guide Volume 2* for more details.

8.18.2 Checking Latency with `rs_ticket`

Although the `sap_status path` command provides the latency information for each active replication path, you can also use the `rs_ticket` command to view detailed information regarding latency among the replication server internal threads.

To check the replication latency using `rs_ticket`

1. Log into the primary SAP ASE server with `isql` as the `DR_admin` login.
2. Switch to the database for which you need to investigate the replication latency.
3. Issue:

```
exec rs_ticket 'my_ticket'
```

4. Log into the standby SAP ASE server with `isql` as the `DR_admin` login.
5. Switch to the database for which you are investigating the replication latency.
6. Issue:

```
select pdb_t, exec_t, dist_t, dsi_t, rdb_t from rs_ticket_history where  
h1='my_ticket'
```

This produces output similar to:

```
  pdb_t          exec_t          dist_t          dsi_t          rdb_t  
-----  
  Aug 20 2015 11:05PM      Aug 20 2015 11:05PM      Aug 20 2015 11:05PM      Aug 20  
2015 11:05PM      Aug 20 2015 11:05PM  
(1 row affected)
```

This output shows the following latency information:

- Latency from the primary ASE server to the standby ASE server:
 - Calculate by `rdb_t - pdb_t`
- Latency from the primary ASE server to Replication Server internal threads:
 - Latency to EXEC thread: `exec_t - pdb_t`
 - Latency to DIST thread: `dist_t - pdb_t`
 - Latency to DSI thread: `dsi_t - pdb_t`

i Note

`rs_ticket` command can also be used to check the health of the Replication Server. This mechanism is referred to as "heartbeat." Write a loop in which `rs_ticket` is sent from the primary server after a specific time interval (for example, every 10 minutes), and then continue checking the `rs_ticket_history` table on the target or the standby server to verify if the ticket is received. A successfully received ticket indicates that the Replication Server is functional.

8.19 Customizing File Directories

Both the HADR pair and DR node support up to five customized directories for database, transfer log, log, configuration, and backup files.

i Note

This functionality applies to the Linux platforms only.

You can configure the following parameters in the response file to customize the file directories.

Parameter	Description	Sample Value
<code>dm_database_file_directory</code>	The database file directory	<code>\$SYBASE/<SID>/database</code>
<code>dm_translog_file_directory</code>	The translog file directory	<code>\$SYBASE/<SID>/translog</code>
<code>dm_log_file_directory</code>	The log file directory	<code>\$SYBASE/<SID>/log</code>
<code>dm_config_file_directory</code>	The configuration file directory	<code>\$SYBASE/<SID>/cfg</code>
<code>dm_backup_file_directory_for_database</code>	The backup file directory for the database	<code>\$SYBASE/<SID>/backup</code>

The following tables list SAP Replication Server and Replication Management Agent instance files that are saved in their corresponding customized directories.

Table 5: SAP Replication Server Instance Files

File Name	Description	Directory
<code>\$SYBASE/<SID>/log/<SID>_REP_PR/<SID>_REP_PR.log</code>	The log file for SAP Replication Server server, <code><SID>_REP_PR</code> .	The log file directory

File Name	Description	Directory
<code>\$SYBASE/<SID>/log/ <SID>_REP_PR/ <SID>_REP_PR.stderr</code>	Standard error output for SAP Replication Server server, <SID>_REP_PR.	The log file directory
<code>\$SYBASE/<SID>/log/ <SID>_REP_PR/ <SID>_REP_PR.stdout</code>	Standard output for SAP Replication Server server, <SID>_REP_PR.	The log file directory
<code>\$SYBASE/<SID>/log/ <SID>_REP_PR/ <SID>_REP_PR_RSSD_ra.out</code>	The output file for the db1tm process in the ERSSD.	The log file directory
<code>\$SYBASE/<SID>/cfg/ <SID>_REP_PR/ <SID>_REP_PR.cfg</code>	The configuration file for SAP Replication Server server, <SID>_REP_PR, created by <code>rs_init</code> .	The configuration file directory
<code>\$SYBASE/<SID>/cfg/ <SID>_REP_PR/ RUN_<SID>_REP_PR.sh</code>	The run file for SAP Replication Server server, <SID>_REP_PR.	The configuration file directory
<code>\$SYBASE/<SID>/cfg/ <SID>_REP_PR/ <SID>_REP_PR_RSSD_ra.cfg</code>	The onfiguration file for the db1tm process in the ERSSD.	The configuration file directory

Table 6: SAP Replication Server ERSSD Files

File Name	Description	Directory
<code>\$SYBASE/<SID>/database/ <SID>_REP_PR_RSSD/db/ <SID>_REP_PR_RSSD.db</code>	The database file for the <code>dbsrv17</code> process in the ERSSD.	The atabase file directory
<code>\$SYBASE/<SID>/log/ <SID>_REP_PR_RSSD/ <SID>_REP_PR_RSSD.out</code>	The output file for the <code>dbsrv17</code> process in the ERSSD.	The log file directory
<code>\$SYBASE/<SID>/log/ <SID>_REP_PR_RSSD/ backup.syb</code>	The log file for the backup operation.	The log file directory
<code>\$SYBASE/<SID>/translog/ <SID>_REP_PR_RSSD/ translog/ <SID>_REP_PR_RSSD.log</code>	The log file for the <code>dbsrv17</code> process in the ERSSD.	The translog file directory
<code>\$SYBASE/<SID>/backup/ <SID>_REP_PR_RSSD/backup/ <SID>_REP_PR_RSSD.db</code>	The backup database file for the <code>dbsrv17</code> process in the ERSSD.	The backup file directory for the database
<code>\$SYBASE/<SID>/backup/ <SID>_REP_PR_RSSD/backup/ <SID>_REP_PR_RSSD.log</code>	The backup log file for the <code>dbsrv17</code> process in the ERSSD.	The backup file directory for the database
<code>\$SYBASE/<SID>/backup/ <SID>_REP_PR_RSSD/backup/ <SID>_REP_PR_RSSD.mlg</code>	The backup mirror file for the <code>dbsrv17</code> process in the ERSSD.	The backup file directory for the database

Table 7: Replication Management Agent Instance Files

File Name	Description	Directory
\$SYBASE/<SID>/log/ AgentContainer/logs/ RMA_<YYYYMMDD>.log	Log files for the RMA instance.	The log file directory
\$SYBASE/<SID>/cfg/ AgentContainer/config/ Logger.xml	The onfiguration file for the RMA logger.	The configuration file directory
\$SYBASE/<SID>/cfg/ AgentContainer/config/ RAO.properties	The file that defines the Replication Agent for Oracle configuration properties.	The configuration file directory
\$SYBASE/<SID>/cfg/ AgentContainer/config/ RAS.properties	The file that defines the Replication Agent for SQLServer configuration properties.	The configuration file directory
\$SYBASE/<SID>/cfg/ AgentContainer/config/ RAU.properties	The file that defines the Replication Agent for UDB configuration properties.	The configuration file directory
\$SYBASE/<SID>/cfg/ AgentContainer/config/ RA_DB.properties	The file that defines RepAgent properties, used only for the <SID> database, when you set up replication.	The configuration file directory
\$SYBASE/<SID>/cfg/ AgentContainer/config/ RA_DB_master.properties	The file that defines RepAgent properties, used only for the master database, when you set up replication.	The configuration file directory
\$SYBASE/<SID>/cfg/ AgentContainer/config/ RS.properties	The file that defines replication server properties when replication is set up.	The configuration file directory
\$SYBASE/<SID>/cfg/ AgentContainer/config/ RS_DB.properties	The file that defines SAP Replication Server database connection properties when replication is set up.	The configuration file directory
\$SYBASE/<SID>/cfg/ AgentContainer/config/ RS_DB_master.properties	The file that defines SAP Replication Server database connection properties, used only for the master database, when you set up replication.	The configuration file directory
\$SYBASE/<SID>/cfg/ AgentContainer/config/ bootstrap.prop	The configuration file for starting an RMA instance.	The configuration file directory
\$SYBASE/<SID>/cfg/ AgentContainer/config/ security/csi.xml	The configuration file for the CSI.	The configuration file directory
\$SYBASE/<SID>/cfg/ AgentContainer/config/ security/roles.cfg	The configuration file for roles.	The configuration file directory
\$SYBASE/<SID>/cfg/ AgentContainer/config/ security/rsa.keys	RSA keys.	The configuration file directory

File Name	Description	Directory
<code>\$\$SYBASE/<SID>/cfg/ AgentContainer/config/ security/sa.pwd</code>	The encrypted password of user 'sa'.	The configuration file directory

Table 8: Replication Management Agent Derby Database Files

File Name	Description	Directory
<code>\$\$SYBASE/<SID>/database/ AgentContainer/configdb/ rsgerepo/seg0/*.dat</code>	Database files for the derby database.	The database file directory
<code>\$\$SYBASE/<SID>/database/ AgentContainer/configdb/ rsgerepo/ service.properties</code>	A text file with internal configuration information for derby.	The database file directory
<code>\$\$SYBASE/<SID>/database/ AgentContainer/configdb/ rsgerepo/dbex.lck</code>	A lock file for derby.	The database file directory
<code>\$\$SYBASE/<SID>/database/ AgentContainer/configdb/ rsgerepo/db.lck</code>	A lock file for derby.	The database file directory
<code>\$\$SYBASE/<SID>/translog/ AgentContainer/configdb/ rsgerepo/log/log.ctrl</code>	A translog file for derby.	The translog file directory
<code>\$\$SYBASE/<SID>/translog/ AgentContainer/configdb/ rsgerepo/log/log5.dat</code>	A translog file for derby.	The translog file directory
<code>\$\$SYBASE/<SID>/translog/ AgentContainer/configdb/ rsgerepo/log/ logmirror.ctrl</code>	A translog file for derby.	The translog file directory
<code>\$\$SYBASE/<SID>/log/ AgentContainer/configdb/ derby.log</code>	The log file for derby.	The log file directory
<code>\$\$SYBASE/<SID>/backup/ AgentContainer/backups/ RepoBackupXXXXXXXXXX.full</code>	A backup file.	The backup file directory for the database
<code>\$\$SYBASE/<SID>/backup/ AgentContainer/backups/ RepoBackupXXXXXXXXXX.diff</code>	A backup file.	The backup file directory for the database
<code>\$\$SYBASE/<SID>/backup/ AgentContainer/backups/ repository.catalog</code>	The repository catalog.	The backup file directory for the database

8.20 SQL Statement Replication

SQL statement replication replicates batched SQL statements in stored procedures, which complements log-based replication and addresses performance degradation caused by batch jobs.

In SQL statement replication, SAP Replication Server receives the SQL statement that modified the primary data, rather than the individual row changes from the transaction log.

See [SQL Statement Replication](#) in *SAP Replication Server Administration Guide Volume 2* for more details. In the HADR environment, SAP Replication Server and RepAgent commands are wrapped within RMA commands to provide a more straightforward way to manage SQL statement replication.

The SQL statement replication tasks described here all use the `sap_sql_replication` RMA command.

Related Information

[sap_sql_replication \[page 514\]](#)

8.20.1 Enabling SQL Statement Replication

Use the `on` parameter of the `sap_sql_replication` RMA command to enable SQL statement replication for either the database or table level.

Procedure

Use the `on` parameter to enable SQL statement replication:

```
sap_sql_replication { <database> | All } , on [, <option>[<option>][...] [, <table>[, <table>][,...]]]
<option> ::= { U | D | I | S }
```

- `<database> | All`
To enable SQL statement replication for a specific database, specify the `<database>` parameter. Use `All` to enable the whole HADR environment.
- `<option>[<option>][...]`
The DML operations you want to enable in SQL statement replication:
 - U – update
 - D – delete
 - I – insert select
 - S – select into
- `<table>[,<table>][,...]`
Enable SQL statement replication for the tables you specify.

i Note

See [sap_sql_replication \[page 514\]](#) for all the configuration requirements when specifying a table in this command.

This example replicates update, delete, and insert select statements as SQL statements for tables as specifies:

```
sap_sql_replication ERP_1, on, UDI, owner1.*, *x*y.a*b, *.sqldmltest_tb1,  
owner3.sqldmltest_tb2, sqldmltest_tb3
```

The command returns:

```
TASKNAME          TYPE  
VALUE  
-----  
-----  
SQL Replication Start Time      Thu Sep 13 02:28:33 UTC  
2018  
SQL Replication Elapsed Time  
00:00:00  
SQLReplication Task Name        SQL  
Replication  
SQLReplication Task State  
Completed  
SQLReplication Short Description Toggle SQL Replication in the  
system  
SQLReplication Long Description Enable the SQL Replication on Replication  
Server.  
SQLReplication Current Task Number  
2  
SQLReplication Total Number of Tasks  
2  
SQLReplication Task Start      Thu Sep 13 02:28:33 UTC  
2018  
SQLReplication Task End      Thu Sep 13 02:28:33 UTC  
2018  
SQLReplication Hostname  
rmazwang2site0.mo.sap.corp  
(11 rows affected)
```

Related Information

[sap_sql_replication \[page 514\]](#)

8.20.2 Setting SQL Statement Replication Threshold

Use the `threshold` parameter of the `sap_sql_replication` RMA command to define when SQL statement replication triggers.

Context

The threshold value is the minimum number of rows affected by a SQL statement when SQL statement replication is triggered. By default, SQL statement replication is triggered when the SQL statement affects more than 50 rows. You can adjust the value of `threshold` according to your needs. You can only set different threshold values at the database level.

Procedure

Set the threshold by specifying the `threshold` parameter in the `sap_sql_replication` command:

```
sap_sql_replication { <database> | All }, threshold, <value>
```

- `<database> | All`
To set the threshold for a specific database, specify the `<database>` parameter. Use `All` to set the threshold for the whole HADR environment.
- `<value>`
The `<value>` parameter defines the minimum number of rows affected by a SQL statement when the SQL statement replication is triggered.

This example sets the threshold to 99 for the `ERP` database, which means SQL statement replication for the `ERP` database is triggered when the SQL statement affects more than 99 rows:

```
sap_sql_replication ERP, threshold, 99
```

The command returns:

TASKNAME	TYPE	VALUE
SQL Replication	Start Time	Thu Sep 13 02:22:13 UTC 2018
SQL Replication	Elapsed Time	00:00:00
SQLReplication	Task Name	SQL Replication
SQLReplication	Task State	Completed
SQLReplication	Short Description	Toggle SQL Replication in the system
SQLReplication	Long Description	Set the SQL Replication threshold on ASE.
SQLReplication	Current Task Number	1
SQLReplication	Total Number of Tasks	1
SQLReplication	Task Start	Thu Sep 13 02:22:13 UTC 2018
SQLReplication	Task End	Thu Sep 13 02:22:13 UTC 2018
SQLReplication	Hostname	rmazwang2site0.mo.sap.corp
(11 rows affected)		

Related Information

[sap_sql_replication \[page 514\]](#)

8.20.3 Disabling SQL Statement Replication

Use the `off` parameter of the `sap_sql_replication` RMA command to disable SQL statement replication for either the database or table level.

Procedure

use the `off` parameter to disable SQL statement replication:

```
sap_sql_replication { <database> | All } , off [, <option>[<option>][...] [, <table>[, <table>][,...]]]
<option> ::= { U | D | I | S }
```

- `<database> | All`
To disable SQL statement replication for a specific database, specify the `<database>` parameter. Use `All` to disable the whole HADR environment.
- `<option>[<option>][...]`
The DML operations you want to disable in SQL statement replication:
 - U – update
 - D – delete
 - I – insert select
 - S – select into
- `<table>[,<table>][,...]`
Disable SQL statement replication for the tables you specify.

i Note

See [sap_sql_replication \[page 514\]](#) for all the configuration requirements when specifying a table in this command.

This example disables the replication of update and delete statements as SQL statements for the `ERP` database:

```
sap_sql_replication ERP, off, UD
```

The command returns:

TASKNAME	TYPE	VALUE
SQL Replication	Start Time	Thu Sep 13 02:23:47 UTC 2018
SQL Replication	Elapsed Time	00:00:00
SQLReplication	Task Name	SQL Replication


```

SQLReplication Task State Completed
SQLReplication Short Description Toggle SQL Replication in the system
SQLReplication Long Description Disable the SQL Replication on ASE.
SQLReplication Current Task Number 2
SQLReplication Total Number of Tasks 2
SQLReplication Task Start Thu Sep 13 02:23:47 UTC 2018
SQLReplication Task End Thu Sep 13 02:23:47 UTC 2018
SQLReplication Hostname rmazwang2site0.mo.sap.corp
(11 rows affected)

```

Related Information

[sap_sql_replication \[page 514\]](#)

8.20.4 Displaying SQL Statement Settings

Use the `display` parameter of the `sap_sql_replication` command to display SQL statement settings, such as the value of threshold and the tables that are enabled or disabled with SQL statement replication.

Procedure

Display SQL statement settings:

```
sap_sql_replication { <database> | All }, display
```

This example displays SQL statement settings for the database `ERP`:

```
sap_sql_replication ERP, display
```

The command returns:

DB_NAME	THRESHOLD	SQL_TYPE	LIST_TYPE	TABLE_LIST
ERP	50	D	All	[]
ERP	50	U	All	[]
ERP	50	S	None	[]
ERP	50	I	None	[]

(4 rows affected)

You can view the value of threshold and tables enabled or disabled with SQL statement replication for corresponding DML operations:

- All – SQL statement replication for all tables are enabled for the corresponding DML operation.
- None – SQL statement replication for all tables are disabled for the corresponding DML operation.

This example displays the SQL statement settings for the database `ERP_1`:

```
sap_sql_replication ERP_1, display
```

The command returns:

```

DB_NAME      THRESHOLD  SQL_TYPE  LIST_TYPE
TABLE_LIST
-----
ERP_1        50         D         In-List   [owner1.*, *x*y.a*b,
*.sqldmltest_tb1, owner3.sqldmltest_tb2, *.sqldmltest_tb3]
ERP_1        50         U         In-List   [owner1.*, *x*y.a*b,
*.sqldmltest_tb1, owner3.sqldmltest_tb2, *.sqldmltest_tb3]
ERP_1        50         S         None
[]
ERP_1        50         I         In-List   [owner1.*, *x*y.a*b,
*.sqldmltest_tb1, owner3.sqldmltest_tb2, *.sqldmltest_tb3]
(4 rows affected)

```

You can view the value of threshold and tables enabled or disabled with SQL statement replication for corresponding DML operations:

- **In-List** – SQL statement replication for corresponding DML operation is enabled for tables listed in the `TABLE_LIST` column.

This example displays the SQL statement settings for the database `ERP`:

```
sap_sql_replication ERP, display
```

The command returns:

```

DB_NAME      THRESHOLD  SQL_TYPE  LIST_TYPE
TABLE_LIST
-----
ERP          50         D         None
[]
ERP          50         U         Out-List  [*.sqldmltest_tb1, *.sqldmltest_tb2,
owner3.sqldmltest_tb3, dbo.sqldmltest_tb4]
ERP          50         S         None
[]
ERP          50         I         None
[]
(4 rows affected)

```

You can view the value of threshold and tables enabled or disabled with SQL statement replication for corresponding DML operations:

- **Out-List** – SQL statement replication for corresponding DML operation is disabled for tables listed in the `TABLE_LIST` column. Tables in the database that are not listed in the `TABLE_LIST` column still use SQL statement replication.

Related Information

[sap_sql_replication \[page 514\]](#)

8.21 Managing Configuration Files

Manage RMA configuration files at the global and the instance levels.

RMA has two level of configuration files. One is at the global level under the `RMA-16_0/config` directory, the other is at the instance level under the `RMA-16_0/instance/AgentContainer/config` directory. Changing the parameters in the global configuration file changes the configurations for all RMA instances under the same RMA installation. Changing the parameters in the instance configuration file changes the configuration for the specific RMA instance.

i Note

In the HADR environment, only one RMA instance is created under an RMA installation.

You should configure parameters using the instance configuration file because:

- The configurations in the instance file take precedence over the configurations in the global file. When RMA is executing, it checks the configurations in the instance file first and then the global file.
- In a rolling upgrade, the installer overwrites the configuration file under the global level, but not the configuration file under the instance level.

If the customized directory is enabled, the customized configuration directory is used instead of the instance configuration directory. RMA checks the configurations in the customized configuration file first and then in the global configuration file.

9 Developing Applications for an HADR System

You can develop client applications that support the HADR functionality using SDK 16.0 SP02.

SDK 16.0 SP02 supports SAP ASE high-availability disaster recovery (HADR) through OCS, SAP jConnect for JDBC 16.0 SP02 (SAP jConnect), and SAP ASE ODBC Driver 16.0 SP02.

An HADR configuration consists of two SAP ASE servers:

- Primary server: One server is the designated primary, and all transaction processing by user applications takes place on the primary.
- Warm standby: The second server acts as a warm standby to the primary server.

If the state of the primary server changes to deactivated, the standby is activated and becomes the new primary server. During the deactivation process, SAP ASE notifies client applications of its state changes. The notifications allow clients to act on state changes. For example, a client can stop initiating new transactions until it receives a message saying the new primary is activated.

To support HADR functionality, client applications can use the SAP jConnect and SAP ASE ODBC Driver features described here.

9.1 SAP jConnect Features for HADR

SAP jConnect provides special connection properties and state change messages for SAP ASE high-availability disaster recovery (HADR).

9.1.1 Connection Properties for HADR in SAP jConnect

The `HADR_MODE` property enables SAP jConnect to use the HADR functionality of SAP ASE.

The `HADR_MODE` property lets you enable or disable HADR features. By default, HADR mode is disabled. Valid settings for this property include:

- `NONE` – SAP jConnect enables no HADR features.
- `MAP` – SAP jConnect requests that SAP ASE send the HADR address list to the application upon login as well as asynchronously when the list is modified, informing the application of the available primary and standby servers, and whether any companion servers are available.
- `NOKILL` – SAP jConnect requests that SAP ASE not terminate an existing connection while the current primary is deactivating or deactivated. The connection is reset by SAP ASE when a new server is made primary.
- `NOKILL_WITH_MAP` – SAP jConnect requests that SAP ASE not terminate an existing connection while the current primary is deactivating or deactivated. In addition, when you set `NOKILL_WITH_MAP`, SAP

jConnect requests that SAP ASE send the HADR address list to the application upon login as well as asynchronously when the list is modified. The address list informs the application of the available primary and standby servers, and whether any companion servers are available. Setting `NOKILL_WITH_MAP` has the same effect as setting the `MAP` and `NOKILL` properties.

- **RECONNECT** – SAP jConnect automatically sets `NOKILL_WITH_MAP` and `REQUEST_HA_SESSION = true`. When the primary server is deactivated, SAP jConnect reconnects to the newly activated primary. After the reconnection is completed, a JZOF2 SQLState exception is displayed. To use the new connection, the application must catch the JZOF2 SQLState exception and restore the context (for example, set options if any, re-prepare statements, cursors, whether the database is in use, and so on). If any transaction is canceled due to a `force` option used in the `sp_hadr_admin deactivate` command, the application is informed, and must retry the transaction.

See the SAP ASE documentation for information on the `force` option.

i Note

For `HADR_MODE = RECONNECT`, SAP jConnect internally makes `REQUEST_HA_SESSION = true`. In this case, the client application must set the `SECONDARY_SERVER_HOSTPORT` connection property. The `SECONDARY_SERVER_HOSTPORT` connection property value specifies the companion server address. When the primary server is down, the client application connects to the companion server. If `SECONDARY_SERVER_HOSTPORT` value is not provided, the following error message is displayed:

```
JZOF1: SAP Adaptive Server Enterprise high-availability failover connection was requested but the companion server address is missing.
```

If the client application uses `HADR_MODE = RECONNECT` and explicitly sets `REQUEST_HA_SESSION = false`, SAP jConnect internally over-rides the client setting and sets `REQUEST_HA_SESSION = true`.

9.1.2 Retrieve HADR State Change Messages from SAP jConnect

SAP jConnect receives messages from SAP ASE about HADR state changes.

The different states of the HADR server that the client application can receive are:

- **NONE** – indicates that the HADR server does not support the HADR feature or the client has set connection property `HADR_MODE=NONE/null`.
- **ACTIVE** – indicates that the current connection is to the active primary server and the client can perform any operation.
- **DEACTIVATING** – indicates that the server is undergoing deactivation and the client application must not perform any new operation using the current connection. If connection property is `HADR_MODE=RECONNECT/NOKILL/NOKILL_WITH_MAP`, then the active transactions can be extended. No new transaction can be started in DEACTIVATING state. If client tries to perform any new operation a `SQLException` is thrown with error code:2377.
- **DEACTIVATED** – indicates that the server was successfully deactivated and no new operation can be performed using the current connection. If connection property is `HADR_MODE=RECONNECT/NOKILL/NOKILL_WITH_MAP`, then the connections are intact but are not usable. If the client tries to execute any query a `SQLException` is thrown with error code: 2379.
If connection property is `HADR_MODE=MAP/NONE`, and the client tries to perform operations in this state then the connection is terminated.

- `DEACTIVATION_CANCELLED` – indicates that the server deactivation was canceled and current connection object is reusable.
- `REACTIVATED` – indicates that a deactivated server is reactivated, hence the current connection object is reusable. This state is seen only when the connection property is `HADR_MODE=RECONNECT/NOKILL/NOKILL_WITH_MAP`.
- `FAILOVER` – indicates that the connection to the old active primary server is killed, and the client is now connected to the new active primary server. During failover, if the client tries to execute a query, the following message displays:

```
JZ0F2 - SAP Adaptive Server Enterprise high-availability failover has occurred.
The current transaction is aborted, but the connection is still usable.
```

The client application must re-create/reset all context objects, such as statements, prepared statements, callable statements, and so on. The `FAILOVER` state appears only when the connection property is set to `HADR_MODE=RECONNECT`.

You can retrieve the server state change messages in these two ways:

- By using the `getClientInfo()` API of the connections
- By implementing the `SybMessageHandler` interface

To retrieve the current HADR server state, client applications must pass the `HADR_CURRENT_STATE` string parameter using the `getClientInfo()` API.

Example

In this example, the driver does not make a round trip to the server to retrieve the state change messages. Instead, it reads the outbound messages sent by the server whenever the state change occurs:

```
conn = DriverManager.getConnection (url, props);
String hadrState = conn.getClientInfo("HADR_CURRENT_STATE");
```

To retrieve the server state change messages using the `getClientInfo()`, refer to the sample code from the `HADRApp.java` file:

```
public static String getCurrentHADRState(Connection conn)
throws SQLException
{
    return conn.getClientInfo("HADR_CURRENT_STATE");
}
...
public static void main(String args[])
{
    Properties props = new Properties();
    props.put("user", "USER");
    props.put("password", "PASSWORD");
    // Client application can used HADR server feature
    props.put("HADR_MODE", "RECONNECT");

    String url = "jdbc:sybase:Tds:HOST:PORT";
    Connection conn = null;
    conn = DriverManager.getConnection (url, props);
    ...
    String insertQuery = "insert into Employee (name, country) values (?, ?)";

    PreparedStatement preparedStatement = conn.prepareStatement(insertQuery);
```

```

for (int currentRow = 1; currentRow <= 50; currentRow++)
{
    try
    {
        // Set parameters to perform insert
        preparedStatement.setString(1, "Name: " + currentRow);
        preparedStatement.setString(2, "Country: " +
currentRow);

        System.out.println("Inserting ROW: " + currentRow);
        preparedStatement.executeUpdate();
        System.out.println("Insert Done");

        System.out.println("Sleeping...");
        Thread.sleep(1000);
    }
    catch (SQLException sqlEx)
    {
        // Gets current HADR server state
        String hadrState = getCurrentHADRState(conn);

        System.out.println("Current Server state: " + hadrState);
        System.out.println("Error Message: " + sqlEx.getMessage());
        if (sqlEx.getSQLState().equals("JZ0F2"))
        {
            // JZ0F2: SAP Adaptive Server Enterprise high-availability
            // failover has occurred. The current transaction is aborted,
            // but the connection is still usable.
            // Retry your transaction.

            // Creating new prepared statement object since
            // client application got connected to new server
            preparedStatement = conn.prepareStatement(insertQuery);
            System.out.println("Recreated new prepare statement object");

            // Retry current transaction
            currentRow--;
        }
        else
        if (hadrState.equals("DEACTIVATING") ||
            hadrState.equals("DEACTIVATED"))
        {
            // Following HADR error code can come:
            // Error Code: 2377 - The primary member is undergoing
            // deactivation.
            // Error Code: 2379 - The primary server has been successfully
            // deactivated.
            System.out.println("HADR Error: " + sqlEx.getMessage());
            // Since server is not in activate state,
            // current transaction is retried
            currentRow--;
        }
        else
        {
            // Something different then HADR went wrong
            throw sqlEx;
        }
    }
}

```

To retrieve the server state change messages using the `SybMessageHandler` interface, the client application implements the `SybMessageHandler`:

```

import com.sybase.jdbcx.SybMessageHandler;
public interface SybMessageHandler
{
    public SQLException messageHandler(SQLException sqe);
}

```

```
}
```

The application registers the `SybMessageHandler` class using the `SybConnection.setSybMessageHandler(SybMessageHandler hndlr)` method. SAP jConnect calls this handler whenever it receives any messages from SAP ASE, or any other error messages, including state change messages.

This example shows how to code the message handler:

```
class HADRMsgHandler implements SybMessageHandler
{
    public SQLException messageHandler(SQLException sqe)
    {
        System.out.println("=====");
        System.out.println(sqe.getMessage());
        int errorCode = sqe.getErrorCode();
        System.out.println(errorCode);
        if(errorCode == 2378)
        {
            ResultSet rs = ((SybSQLException) sqe).getEedParams();
            try
            {
                System.out.println("Deactivation time : " + rs.getTimestamp(1));
            }
            catch (SQLException s)
            {
            }
        }
        return sqe;
    }
}
//register the SybMessageHandler implementation.
_conn.setSybMessageHandler(new HADRMsgHandler());
```

9.1.3 Retrieve HADR_LIST_MAP Connection Property from SAP jConnect

Set `HADR_MODE` to `MAP` or `NOKILL_WITH_MAP` in SAP jConnect to extract `HADR_LIST_MAP` from these properties.

When you set `HADR_MODE` to `MAP` or `NOKILL_WITH_MAP`, SAP jConnect receives the `HADR_LIST_MAP` during login and whenever there is a topology change. To retrieve `HADR_LIST_MAP`, an SAP jConnect application calls the `SybConnection.getClientInfo()` method. `SybConnection.getClientInfo()` returns the property object.

The client application extracts `HADR_LIST_MAP` from these properties, which returns:

```
LinkedHashMap <string>, <object>
```

To retrieve the `HADR_LIST_MAP` components from this `LinkedHashMap`, retrieve these keys:

```
GroupName
GenerationNumber
Primary
Standby_1
```


The code in this example uses the keys above to retrieve HADR_LIST_MAP components from the LinkedHashMap:

```
public static void testHadrMap() throws Exception
{
    Class.forName("com.sybase.jdbc4.jdbc.SybDriver");
    String url = ("jdbc:sybase:Tds:SERVER:PORT?HADR_MODE=NOKILL_WITH_MAP");
    SybConnection conn = (SybConnection) DriverManager.getConnection(url, "USER",
    "PASSWORD");
    Properties props = ((SybConnection) conn).getClientInfo();
    LinkedHashMap hadrMap = (LinkedHashMap)props.get("HADR_LISTMAP");
    // Populate the ListMap
    System.out.println("Group Name : " + hadrMap.get("GroupName"));
    System.out.println("Generation Number: " + hadrMap.get("GenerationNumber"));
    printDSN(hadrMap, "Primary");
    for(int i=1; i<=hadrMap.size()-3; i++)
    {
        printDSN(hadrMap, "Standby_" + i);
    }
}
public static void printDSN(LinkedHashMap hadrMap, String dsnType)
{
    System.out.println("_____ " + dsnType + " _____");
    Properties dsnInfo = (Properties)hadrMap.get(dsnType);
    System.out.println(dsnInfo.get("DataSourceName"));
    LinkedList dataSourceAddressList = (LinkedList)dsnInfo.get("AddressList");
    LinkedList haFailoverList = (LinkedList)dsnInfo.get("HAFailoverList");
    int flags = ((Integer)dsnInfo.get("Flag")).intValue();
    for(int i=0; i< dataSourceAddressList.size(); i++)
    {
        System.out.println(dataSourceAddressList.get(i));
    }
    for(int i=0; i< haFailoverList.size(); i++)
    {
        System.out.println(haFailoverList.get(i));
    }
    System.out.println("Flag: " + flags);
}
```

9.2 SAP ASE ODBC Driver Features for HADR

The SAP ASE ODBC driver provides special connection properties, informational messages, and API support for SAP ASE high-availability disaster recovery (HADR). Use these new features to write robust database applications and take advantage of the HADR system to stay always-on.

9.2.1 Connection Properties for HADR in the SAP ASE ODBC Driver

Several properties in the SAP ASE ODBC driver support the HADR functionality of SAP ASE.

- `DRNoKillDuringDeactivation` – when this property is set to 1 (the default is 0), the SAP ASE ODBC driver requests the SAP ASE to not terminate the connection when the primary server is in a deactivated state or is undergoing deactivation.

- `HADRlist` – when this property is set to 1 (the default is 0), the SAP ASE ODBC driver requests the SAP ASE to send the HADR address list to the application upon login as well as asynchronously when the list is modified. This helps to inform the application of the available primary and standby servers, and whether any companion servers are available.
- `EnableRedirection` – when this property is set to 1 (the default), the SAP ASE ODBC driver allows the server to redirect the connection to an alternate server (cluster) or the primary server (HADR). Setting `EnableRedirection` to 0 disables redirection.
- `HADRMode` – when this property is set to 1 (the default is 0), the SAP ASE ODBC driver enables the `DRNoKillDuringDeactivation`, `HADRlist` and `HASession` connection properties to handle HADR planned and unplanned failover events in the ODBC driver. The HADR application can set only the `HADRMode` property to handle all the HADR-related events.
- `HADRPrimaryWaitTime` – this property sets the time-out value for the ODBC driver to search for a new primary server in the event of an unplanned failover. The default value is 300 seconds. The value of the `HADRPrimaryWaitTime` connection property is always the same as that of the `HADR primary wait time server configuration option`.
- `HADRLoginStallTime` – this property sets the time for the ODBC driver to stall the HADR login of the application in the event of an unplanned failover. The default value is 2 seconds. The value of the `HADRLoginStallTime` connection property is always the same as that of the `HADR login stall time server configuration option`.
- `HADRPrimaryCheckFrequency` – this property sets the frequency of the ODBC driver to check for the primary server in the event of an unplanned failover. The default value is 10 seconds. The value of the `HADRPrimaryCheckFrequency` connection property is always the same as that of the `HADR login stall time server configuration option`.

9.2.1.1 Using Connection Properties

Using the connection properties applications can control the nature of the server state change messages.

Applications that do not monitor the information messages are also notified of some of the events as errors when executing statements.

The connection property `DRNoKillDuringDeactivation` and `HADRlist` control the nature of the messages:

- `DRNoKillDuringDeactivation =0` and `HADRlist =0` – when the primary server is deactivated or failover to the new primary happens, the server kills all the connections without any HADR properties. The SAP ASE server notifies the application with the reasons for the lost connection, which could be due to:
 - Error code 2379 – this error message indicates that the primary server has been successfully deactivated and connection has been lost due to server deactivation.
Error message:
`Primary server has been successfully deactivated`
 - Error code 2376 – this error message indicates that the connection has been lost due to failover.
Error message:
`Server XYZ has been promoted to the primary role`
- `DRNoKillDuringDeactivation =1` – when the primary server starts deactivation, it does not kill the connection until a new primary server is active. Instead, the server rejects new transactions. Before deactivation, the server waits for all active transactions to complete. The server goes in the "being deactivated" state (DEACTIVATING). When the deactivation is in progress, the following error message (error code 2377) is sent to the new transaction requests:

```
Primary server is being deactivated
```

If the active transactions fail to complete until timeout, then the deactivation is forced or canceled, depending upon the `force` option of the deactivation command. All open transactions are rolled back if the deactivation is forced. When the server is deactivated, the connections are kept intact but new transactions fail with the following error message (error code 2379):

```
Primary has been successfully deactivated
```

When a server becomes the new primary server, the server informs the client, then kills the connection. If `HADRMode=1` connection property is not set, then it is the responsibility of the application to establish a new connection with the new primary server.

- `HADRList=1` – when the primary server is deactivated, inactive (idle) connections continue to survive until they perform any database activities (SQL execution) or failover happens to new primary. The connections that perform database activities are killed, with the following error message (error code 2379):

```
Primary server has been successfully deactivated
```

The ODBC driver generates the connection lost error with native message number 30046. If the same server is activated again, the idle list connections continue to operate normally.

9.2.2 Use the SAP ASE ODBC Driver to Get HADR State Change Messages from SAP ASE

SAP ASE notifies ODBC applications via the SAP ASE ODBC driver when an SAP ASE server starts to deactivate, cancels an ongoing deactivation process, transitions to a deactivated state, completes the failover, or transitions from a deactivated state to an active state.

Connections that enable the `HADRList` or `DRNoKillDuringDeactivation` connection property receive these messages when a state change occurs. See [HADR Messages from SAP ASE \[page 377\]](#).

Applications can use the `SQLGetConnectAttr()` API to get the HADR state change messages (server state). Applications can directly be in sync with the server (status) to avoid disaster and be aware of planned upgrades. Getting a server state helps database applications either identify whether the server is undergoing planned events (upgrade/maintenance) or detect disaster events and act accordingly, all without restarting the whole application. The application can always be connected and keep serving the client request as per server availability. The application is also notified about reconnections to the site that has taken over (referred as failover) in case of disaster or site upgrades. In a failover, the application reestablishes the context and continues with the client requests that could not be processed because of the failover.

Using the following connection attributes, an application can monitor the current state of the server:

- `SQL_ATTR_DR_INFO_MSG` – retrieves the current state of the connection. An application can poll the connection by calling the `SQLGetConnectAttr` function and passing in the `SQL_ATTR_DR_INFO_MSG` connection attribute. The value is set to a `SQLINTEGER` value of the most recent informational message received:
 - `SQL_DR_ACTIVATE` – indicates that the server is in an active state and can process the client requests.
 - `SQL_DR_DEACTIVATED` – indicates that the server is deactivated (or inactive) and unable to serve client requests.
 - `SQL_DR_DEACTIVATION_CANCELED` – indicates that the deactivation is canceled and the server is back to active state.

- `SQL_DR_DEACTIVATING` – indicates that the server is in deactivating state; going ahead, it might be deactivated or deactivation might get canceled.
- `SQL_DR_REACTIVATED` – indicates that the deactivated server is back to the active state.
- `SQL_DR_FAILOVER` – indicates that the planned failover is complete and the standby site has taken over to serve user requests.
- `SQL_DR_CONNECTION_LOST` – indicates that connection to the primary server has been lost due to planned or unplanned failover.

Example

The first call to the `SQLGetConnectAttr()` API returns the most recent state of the ASE in the `<hadr_status>` variable. For subsequent calls `<SQL_DR_REACTIVATED>` and `<SQL_DR_DEACTIVATION_CANCELED>` are reported as `<SQL_DR_ACTIVE>`. When the failover in the driver is complete, the `<SQL_DR_FAILOVER>` state is reported as `<SQL_DR_ACTIVE>`. For more about application failover, see [Application Failover \[page 366\]](#).

```
SQLINTEGER connectionState;
SQLGetConnectAttr(connection_handle, SQL_ATTR_DR_INFO_MSG, hadr_status,
sizeof(SQLINTEGER), SQL_NULL_HANDLE)
```

- `SQL_ATTR_DR_INFO_CALLBACK` – applications that link directly to the SAP ASE ODBC driver can avoid polling by registering a callback function using `SQLSetConnectAttr` and setting `SQL_ATTR_DR_INFO_CALLBACK` to the address of the callback function. This function is called when an HADR informational message is received. The callback function is not called on inactive connections because the connection is not proactively monitored. Messages are received when the application executes a statement or fetches rows from the result set. The syntax for the state events callback function is:

```
void HADRInfoCallback(HDBC conn, SQLINTEGER new_state);
```

Example

```
SQLSetConnectAttr(dbc, SQL_ATTR_DR_INFO_CALLBACK, (SQLPOINTER)&HADRInfoCallback,
0);
```

- `SQL_ATTR_DR_DEACTIVATION_TIMEOUT` – retrieves the time at which the deactivation occurs. An application can call the `SQLGetConnectAttr` function and pass in the `SQL_ATTR_DR_DEACTIVATION_TIMEOUT` connection attribute. The value is set to a `SQL_TIMESTAMP_STRUCT` containing the exact time (in the universal time zone) when the deactivation timeout will occur.

9.2.3 Retrieve the HADR List

Connections that enable the `HADRList` property receive an HADR data source list from the server upon login and any time the HADR data source list changes.

The data source list contains the current primary server (listed first), followed by all available standby servers. Each data source enumerates a list of addresses (which refer to the same server) and a list of high-availability companion data sources available for that data source. To retrieve these messages, SAP ASE ODBC driver applications can poll the connection by calling the `SQLGetConnectAttr` function, which uses the following connection properties:

- `SQL_ATTR_DR_LIST_GENERATION` – returns a `SQLINTEGER` identifying the generation of the data source list. The application uses this property to decide if it already has the current list before getting the `SQL_ATTR_HADR_LIST` attribute.
- `SQL_ATTR_DR_LIST` – returns a structure containing the details of the data source list. The `SQLGetConnectAttr` function copies the data into the memory provided through the `ValuePtr` parameter. If the `BufferLength` is not large enough, `SQLGetConnectAttr` sets `StringLengthPtr` to the size needed and returns `SQLERROR` with a data overflow error:

```
SQLState=HY000, MessageText=Data overflow.
```

```
Increase specified column size or buffer size, NativeError=30128
```

If the data is copied, `StringLengthPtr` is set to the number of bytes used. `ValuePtr` must be aligned on a 16-byte boundary.

- `SQL_ATTR_HADR_LIST_CALLBACK` – applications that link directly to the SAP ASE ODBC driver can avoid polling by registering a callback function using `SQLSetConnectAttr` and setting `SQL_ATTR_HADR_LIST_CALLBACK` to the address of the `callback` function. This function is called when an updated list is received from the server. The syntax of the `callback` function is:

```
void HADRListAllocCallback(HDBC <conn>, SQLINTEGER <generation_number>,
SQLLEN <size_needed>);
```

Where:

- `<conn>` – is the connection handle on which the message was received.
- `<generation_number>` – is the generation number of the new list that determines whether the application retrieves the new list or already has it from a different connection.
- `<size_needed>` – is the amount of memory needed to hold the new list.

When the `callback` function is called, the application, if it decides to update its list, may call the `SQLGetConnectAttr` function and retrieve the `SQL_ATTR_HADR_LIST` attribute to get the new list.

Following is the sample code to retrieve the HADR list:

```
DataSourceList structure
struct SQLHADRDataSourceList
{
    // The generation number of this list
    SQLINTEGER generation;

    // The group name of this list. Regardless of the setting of
    // SQL_OUTPUT_NTS, this name is null terminated.
    SQLWCHAR* group_name;
    // The length of the group_name in bytes.
    SQLLEN group_name_length;
    // The number of data sources in the data source list.
    SQLLEN number_of_data_sources;
};
```

```

// An array of size number_of_data_sources containing pointers to
// each SQLHADRDataSource in the list.
SQLPOINTER* data_source_list;
};
Struct SQLHADRDataSource
{
    // The name of the data source. Regardless of the setting of
    // SQL_OUTPUT_NTS, this name is null terminated.
    SQLWCHAR* data_source_name;
    // The length of data_source_name in bytes
    SQLLEN data_source_name_length;

    // The number of address for this data source. Each address refers
    // to the same data source (server)
    SQLLEN number_of_addresses;
    // An array of size number_of_addresses containing pointers to each
    // address in the array. The addresses are in the same format as
    // addresses in the interfaces file. Regardless of the setting of
    // SQL_OUTPUT_NTS, the addresses are null terminated.
    SQLWCHAR** address_list;
    // An array of size number_of_addresses containing the byte length of
    // each element in the address_list array.
    SQLLEN* address_list_lengths;
    // The number of HA companions available for this data source.
    SQLLEN number_of_ha_companions;
    // An array of size number_of_ha_companions containing pointers to
    // the SQLHADRDataSource for each of the HA companion servers.
    SQLHADRDataSource** ha_companion_list;
    // This SQLINTEGER is treated as a set of flags for the data source.
    // Currently, the only flag defined is SQL_DR_READONLY.
    SQLINTEGER flags;
};

```

9.2.4 Application Failover

Application failover is the reconnection of user applications to a standby server that has been promoted to the primary role upon a failure of the previous primary server, or its planned designation to a standby role for maintenance purposes.

Application failover is not triggered on inactive connections because the connection is not proactively monitored by the driver. Reconnection to a new primary does not happen unless the application executes a statement (a SQL query). When the failover is complete, the driver fails the statement execution by sending an error with native message number 30130:

```
The server is not available or has terminated your connection, you have been
successfully connected to the next available HA server. All active transactions
have been rolled back.
```

Make sure to explicitly migrate the context; this message indicates that the query execution failed because the successful failover in HADR system does not migrate the context. If query execution fails with HA failover success error, the application must reset the context. If failover is unsuccessful for any other reason, the application receives an error with native message number 30131:

```
Connection to the server has been lost, connection to the next available HA server
also failed. All active transactions have been rolled back.
```

Note

A successful failover in an HADR system does not migrate the context. The application has to reset the current database, any set options, client language, and character sets. All context information from the

connection string is applied to the new connection. For more details about migrated context see *Cluster Edition Cluster Users Guide*.

To configure the ODBC driver to handle planned and unplanned failover, set `HADRMode=1`. When the `HADRMode` property is set to 1 (the default is 0), the SAP ASE ODBC driver enables `DRNoKillDuringDeactivation`, `HADRList`, and `HASession` connection properties to handle the HADR planned and unplanned failover events in the ODBC driver.

Use the `HADR primary wait time` configuration parameter to determine the amount of time, in seconds, the standby server continues to send the redirect list to the clients in absence of primary server before failing the connection. See the *Reference Manual: Configuration Parameters*.

9.2.4.1 Handling Planned Failovers

Planned failovers in an HADR system allow the standby site to take over so that the primary site can be released for maintenance purposes.

Context

This procedure describes how to perform a planned failover:

Procedure

1. Deactivate the current primary site. When deactivation starts, the server state changes to deactivating. When the deactivation is successful, the server state changes to the deactivated state. Fetching the server state returns `SQL_DR_DEACTIVATING` or `SQL_DR_DEACTIVATED`, depending on the server state. Applications cannot start new transactions when the server is in deactivating or deactivated states; doing so results in an error. Applications have to wait and keep polling the server state until there is an active primary server. To deactivate the current primary site, after a failed query execution:
 - a. The application checks to see whether the execution failed because of a planned or unplanned HADR event.
 - b. While the server is in the deactivating or deactivated state, the application continues to fetch the server state until the state changes from deactivating or deactivated, to some other state.
2. There are three ways in which the application may connect to the active primary server. The application identifies one of the following scenarios and proceeds so that:
 - Deactivation process is canceled – server state is reported as `SQL_DR_DEACTIVATION_CANCELED`. No action is required from the client application, as all connections are intact, and running transactions are not affected. The application resubmits the transactions or query that were not started during deactivation.
 - Reactivation of the same site – server state is reported as `SQL_DR_REACTIVATED`. The application resubmits any transaction that failed when the server was transitioned to the deactivated state.

- Failover to the new site – server state is `SQL_DR_FAILOVER`. See step 4.
If the server state is reactivated or deactivation is canceled, the application continues with the failed transaction.
- 3. When the current primary server reaches the inactive state, the database administrator changes its role to warm standby. Fetching server state still returns `SQL_DR_DEACTIVATED`.
- 4. The planned failover is complete after you put the current standby server into the primary role and make it active. The server state is `SQL_DR_FAILOVER`. This indicates that failover has happened, and the new primary server is ready to serve client requests. When the client application executes a statement, the driver moves the user connection to the new primary server, and the SQL execution returns an HA failover error (30130). The application resets the context when the client application receives the HA failover success error.

9.2.4.2 Handling Unplanned Failover

An unplanned failover occurs when there is a crash or a fault in the primary server and the secondary server takes over the role of primary server, to allow normal use to continue.

If the primary server is down and the application executes a statement, the ODBC driver tries to find a server that has been promoted to the active primary role. The new active primary server may be one of these:

- The failed (previously active) server may be restarted
- The standby server is promoted to the active primary role

If there is no primary server, the driver continues to search for a new active primary server until the time-out is reached (default is 5 minutes). To change the default time-out value, use the server configuration option `HADR primary wait time`.

i Note

The timeout starts the moment the server goes down and not when the client application executes the query.

If the primary server crashes while a planned failover is in progress, the ODBC driver reports the server state as `SQL_DR_CONNECTION_LOST`. Upon receiving the state change message, the application executes a statement so the driver connects to the new active primary server. If the new active primary server is unavailable, the ODBC driver continues to search for the new primary server.

After the failover is complete, the ODBC driver fails the statement execution with an HA failover error. The application resets the context when the driver throws an HA failover success error.

i Note

If the primary server is down at the time of initial connection, the driver tries to connect to the secondary server. In such cases the application must set the `secondaryhost` and `secondaryport` connection properties. For an HADR system, the secondary server is the standby server.

To handle an unplanned failover:

1. Configure the `HADR primary wait time` option to the appropriate value.
2. Configure the application to set the `secondaryhost` and `secondaryport` connection properties.

3. Configure the application to execute a statement when the server state is `SQL_DR_CONNECTION_LOST`. If the server state is `SQL_DR_CONNECTION_LOST`, the application executes a query that triggers the driver to search for a new primary.

9.2.4.3 Determining the Cause of Query/Transaction Failure

Applications can fetch error messages to determine whether the query execution failed because of an HADR event, such as deactivate or failover, and so on.

If applications do not want to perform a search for error codes, they can rely on the callback function, which notifies applications about any changes in state.

Set the callback function for informational messages using the `SQLSetConnectAttr` (`SQL_ATTR_DR_INFO_CALLBACK`) API. Within this callback function, set the global Boolean variable to `true` to indicate that the server state has changed. If the statement execution fails and the global Boolean variable is set, that means the state has changed and application needs to handle the new state. If the query execution fails and the state has not changed, then there is some other error that the application must handle.

Example

This is an example code to determine the cause of query or transaction failure:

```
SQLSetConnectAttr(dbc, SQL_ATTR_DR_INFO_CALLBACK, SQLPOINTER) &HADRStateChanged,
0);
void HADRStateChanged(HDBC conn, SQLINTEGER new_state)
{
    server_state_changed = true;
}
```

9.2.4.4 Resetting the Application Context

After a successful planned or unplanned failover, reset the application context.

Procedure

1. Add a wrapper to the `SQLPrepare()` function and maintain a list of prepared statements. For example:

```
std::vector <SQLTCHAR*> prepared_queries;
SQLRETURN PrepareQuery(SQLHANDLE stmt, SQLTCHAR* sql)
{
    SQLRETURN sr = SQL_ERROR;
    prepared_queries.push_back(sql);
}
```

```

sr = SQLPrepare(stmt, sql, SQL_NTS);
return sr;

```

2. Set the current database. For example:

```

/* Change Database to pubs2 */
void ChangeDatabase(SQLHANDLE dbc)
{
SQLRETURN sr = SQLSetConnectAttr(dbc, SQL_ATTR_CURRENT_CATALOG,
(SQLPOINTER)_ODBCSTRTYPE("pubs2"), SQL_NTS);
PrintError(sr, SQL_HANDLE_DBC, dbc, true);
}

```

3. Prepare all the queries. For example:

```

/* This function sets the application context
1. Change the database to pubs2
2. Prepare all the queries/Re-prepare all the queries after failover */
void SetAppContext(SQLHANDLE dbc, SQLHANDLE stmt)
{
SQLRETURN sr = SQL_ERROR;
ChangeDatabase(dbc);
for (std::vector<SQLTCHAR*>::iterator it = prepared_queries.begin(); it !=
prepared_queries.end(); it++)
{
sr = SQLPrepare(stmt, (SQLTCHAR*)*it, SQL_NTS);
printError(sr, SQL_HANDLE_STMT, stmt, true);
}
}

```

9.2.4.5 Sample Code

Sample codes show how planned and unplanned failovers are handled.

The application creates an unprivileged connection and sets the application context using the `SetAppContext()` function. The application executes an update query for its entire lifecycle and handles the HADR events:

```

while (!executeQuery)
{
sr = SQLExecute(stmt);
if (sr == SQL_SUCCESS || sr == SQL_SUCCESS_WITH_INFO)
{
return sr;
}
if (sr == SQL_ERROR)
{
if (server_state_changed)
{
server_state_changed = false;
failover_completed = false;
SQLGetConnectAttr(dbc, SQL_ATTR_DR_INFO_MSG, &connection_state,
sizeof(connection_state), 0);
while (connection_state == SQL_DR_DEACTIVATED || connection_state ==
SQL_DR_DEACTIVATING)
{
cout << "wait server is deactivated" << endl;
Wait(2);
//user could wait for more time if they want to
SQLGetConnectAttr(dbc, SQL_ATTR_DR_INFO_MSG, &connection_state,
sizeof(connection_state), 0);
}
}
}
}

```

```

    }
    if (connection_state == SQL_DR_REACTIVATED || connection_state ==
SQL_DR_DEACTIVATION_CANCELED)
    {
        cout << "resuming operation as server is put into the active state"
<< endl;
        continue;
    }
    if (connection_state == SQL_DR_CONNECTION_LOST)
    {
        sr = SQLExecute(stmt);
    }
    if (connection_state == SQL_DR_FAILOVER)
    {
        failover_completed = true;
    }
}
recordnumber = 1;
diagret = SQL_ERROR;
while (diagret != SQL_NO_DATA)
{
    diagret = SQLGetDiagRec(SQL_HANDLE_STMT, stmt, recordnumber++, sqlstate,
&nativeerror, errormsg, ERR_MSG_LEN, &errormsglen);
    if (nativeerror == 30130)
    {
        SetAppContext(dbc, stmt);
        failover_completed = true;
        break;
    }
}
//Query execution failed - neither driver failover happened nor server
state changed to SQL_DR_FAILOVER, so this is some other error to be handled by
application
if (!failover_completed)
{
    return sr;
}
}
}

```

For the complete code, refer to the `hadrapp` sample in the SDK.

9.3 SAP CTLIB Features for HADR

SAP CTLIB provides special connection properties to support the SAP ASE high availability disaster recovery (HADR).

9.3.1 Connection Properties for HADR in SAP CTLIB

The following context/connection level properties in the SAP CTLIB support the HADR functionality of ASE:

- `CS_PROP_REDIRECT` – This property is enabled by default. When enabled, it allows the standby server to redirect the connection to an alternate server (cluster) or to the primary server in HADR topology.
- `CS_HAFAILOVER` – This property is disabled by default. When enabled, an HA aware client can failover to an alternate server in case a planned or unplanned failover event takes place in the HADR system.

- `CS_PROP_EXTENDEDFAILOVER` – This property is used only when the `CS_HAFAILOVER` property is enabled. When enabled, this property indicates that an HA aware client is connected to a server that supports HA and failover extensions and can receive a list of network addresses that must be used for failover instead of the information initially retrieved from directory service layer.

9.3.2 Using the CTLIB Context/Connection level properties

Using the context/connection level properties that support the HADR functionality, you can control the behavior of the server with respect to the client.

- The `CS_PROP_REDIRECT` property is enabled and set to `CS_TRUE` by default. In this case, when a client attempts to log onto a standby server, it is redirected to the primary server in the HADR system and a connection is established with the active primary server.
To disable login redirection, first disable the `CS_HAFAILOVER` property and then set the `CS_PROP_REDIRECT` property to `CS_FALSE`.
- The `CS_HAFAILOVER` property is set to `CS_FALSE` by default. To enable the `CS_HAFAILOVER` property, set it to `CS_TRUE`. When enabled, an HA aware client can failover to an alternate server in a planned or unplanned failover. In a failover event, if the `CS_HAFAILOVER` property is disabled, the client does not failover to the standby server and the connection is terminated.

Note

When you enable the `CS_HAFAILOVER` property, the `CS_PROP_REDIRECT` property is also enabled by default.

- The `CS_PROP_EXTENDEDFAILOVER` property is set to `CS_TRUE` by default, but it is used only when the `CS_HAFAILOVER` property is set to `CS_TRUE`. When enabled, the client receives a list of network addresses from server that the client must use for failover instead of relying on information initially retrieved from the directory service layer.

You can set these properties at both the connection and context levels. To set a property at a connection level, use the `ct_con_props()` function. When you set at a property at the connection level, it is applicable only for that connection. Similarly, to set a property at a context level use the `ct_config()` function. When you set a property at the context level, it is set for every connection that is created under that context.

In the following example, the `ct_config()` function enables the `CS_PROP_REDIRECT` property for a single connection and the `ct_con_props()` function disables the `CS_HAFAILOVER` property at the context level:

```
CS_INT fls = CS_FALSE;
CS_INT tru = CS_TRUE;
ret = ct_config(context, CS_SET, CS_HAFAILOVER, &fls, CS_UNUSED, NULL);
if (ret != CS_TRUE)
{
    //raise error
}
ret = ct_con_props(connection, CS_SET, CS_PROP_REDIRECT, &tru, CS_UNUSED, NULL);
if (ret != CS_TRUE)
{
    //raise error
}
```

9.3.3 Handling Failover Events

Failover is the reconnection of applications to a standby server, which has been promoted to the primary role upon a failure of the previous primary server, or its planned designation to a standby role for maintenance purposes.

Failover can be planned or unplanned. In a planned failover, the primary server is set as standby for maintenance purposes and the standby server is promoted to the primary role. Unplanned failovers usually occur when there is a crash or a fault in the primary server and the secondary server takes over the role of primary server to allow normal use to continue.

Failover events are not proactively monitored in CTLIB and the failover on the server side does not result in failover on the client side. The failover on the client side is triggered only when the client application attempts to perform any network interaction with the server. In this case the client application receives an HA failover error message, if a client error message callback handler is installed.

In case of successful failover event, the `CS_RET_HAFAILOVER` return value is returned by the attempted API operations, such as `ct_result()`, `ct_send()`, `ct_fetch()` or any routine, which performs network interaction.

The `CS_RET_HAFAILOVER` return value is returned from the API call during a synchronous connection. In an asynchronous connection the APIs return the `CS_PENDING` value to the caller and the operation is performed asynchronously. Use the `ct_poll()` function to obtain the status of a last asynchronous operation. In the event of a failover, the `ct_poll()` function returns `CS_HAFAILOVER`. Depending on the return code; perform the required processing, such as sending the next command to be executed.

The structure of a basic loop in the event of a failover is:

```
while ct_results returns CS_SUCCEED
  switch on result_type
    case CS_ROW_RESULT...
    case CS_STATUS_RESULT...
    .
    .
    case CS_CMD_DONE...
    case CS_CMD_FAIL...
    case CS_CMD_SUCCEED...
  end switch
end while
switch on ct_results' final return code
  case CS_END_RESULTS... //ct_result was successful
  case CS_CANCELED...
  case CS_FAIL...
  case CS_RET_HAFAILOVER... /*Re-execute the last query or restart transaction
as desired*/
end switch
```

9.4 Support for HA applications on HADR Clients and Servers

SAP jConnect and the SAP ASE ODBC driver provide support for high-availability applications to run with SAP ASE servers participating in an HADR system.

SAP jConnect

SAP jConnect supports high-availability (HA) applications to run with an HADR system. Existing HA client applications can use an HADR system without modifications. For more information about HA, see the *SAP jConnect for JDBC Programmers Reference > Programming Information > Database Issues > Failover Support*.

To implement failover support in SAP jConnect:

- Set `REQUEST_HA_SESSION` to `true`.
- Set `SECONDARY_SERVER_HOSTPORT` to the host name and port number where your secondary server is listening.

The legacy HADR application behaves in the following way with the HADR server:

- Server in standby inactive state (no active primary) – the client application cannot connect to a server that has no active primary server in the topology. The client application gets a login failure exception, with the following error message:

```
JZ00L: Login failed.  
9668: 01ZZZ Login failed. Adaptive Server is running in 'Standby' mode. The  
user login does not have 'allow hadr login' privilege and login redirection  
cannot occur since there is no Active Primary.  
010HA: The server denied your request to use the high-availability feature.  
Please reconfigure your database, or do not request a high-availability session.
```
- Server in primary inactive state:
 - The client application cannot connect to a server in a primary inactive state until the server is made active primary.
 - After connecting to the server, the client applications can execute any queries.
 - The client application becomes unresponsive when the server is in the primary inactive state and throws a `SQLException`.
- Server in primary active state – When the server is in the primary active state the client application can successfully connect to the server.
- Primary active server undergoes deactivation – if the primary active server undergoes deactivation and the client application tries to execute a query, the client application becomes unresponsive until the primary server is reactivated. If the standby server is promoted to the role of primary active, then the client application gets a failover `SQLException` with error code `JZF02`. In this case, the client application has to re-create all context object such as statement, prepared statement, callable statement, and so on, and re-execute the failed transaction/query.

SAP ASE ODBC Drivers

An HA application is one that sets the `HASession` connection property to 1 and handles the failover error returned after an HA failover has successfully completed. For more on HA and `HASession`, see the *Adaptive Server Enterprise ODBC Driver by Sybase Users Guide for Microsoft Windows and UNIX > Failover in High Availability Systems* in the Software Developer's Kit documentation set.

This is a sample connection string:

```
DSN=HADRPrimaryServer;UID=UnPrivilegedUser;PWD=HADRPWD123;SecondaryServer=localhost;
SecondaryPort=1600;HASession=1;
```

When a query is executed on the primary deactivated server, the query is blocked by the server until there is an active primary server. The query execution proceeds normally when the same server is reactivated. If there is a failover, the server fails the command with error codes 2379 and 2376. As a part of processing response of the query, the ODBC driver fails over to the new primary server. After the successful failover, the ODBC driver generates the `HAFailover` error with the native message number 30130. The ODBC application resets the application context and re-executes the failed query/transaction. To reset the application context and re-execute the failed query/transaction, set the `HASession` connection property to 1 so that the ODBC driver fails over to the new primary server.

i Note

The default value of the `commandtimeout` connection property is 30 seconds. The ODBC driver cancels the blocked commands after it reaches this time-out value. To delay the cancellation of the blocked commands, adjust the value of the `commandtimeout` connection property to a higher value.

SAP CTLIB

SAP CTLIB provides compatibility support for existing HA and cluster applications. To use the high availability features in your applications, enable the `CS_HAFAILOVER` property. These applications require minimal or no modifications to run against the HADR servers.

9.5 Support for Cluster Edition Applications on HADR Clients and Servers

The SAP jConnect and the SAP ASE ODBC driver provide support for Cluster Edition applications to run with SAP ASE servers participating in an HADR system.

For details on the SAP ASE Cluster Edition, see the *Cluster Edition Cluster Users Guide*.

SAP jConnect

SAP jConnect supports the SAP Adaptive Server Enterprise Cluster Edition to run with HADR system, where multiple SAP ASE servers connect to a shared set of disks and a high-speed private interconnection. This allows SAP ASE server to scale using multiple physical and logical hosts.

For more information about HA see the *SAP jConnect for JDBC 16.0 Programmers Reference* *SAP jConnect for JDBC 16.0 > Programming Information > Advanced Features*.

Use the connection string to enable connection failover by setting `REQUEST_HA_SESSION` to `true`, where `server1:port1, server2:port2, ... , serverN:portN` is the ordered failover list:

Example

```
URL="jdbc:sybase:Tds:server1:port1,server2:port2,...,serverN:portN/mydb?REQUEST_HA_SESSION=true"
```

SAP jConnect tries to connect to the first host and port specified in the failover list. If unsuccessful, SAP jConnect goes through the list until a connection is established or until it reaches the end of the list.

SAP ASE ODBC Driver

This is a sample connection string:

```
DSN=HADRPrimaryServer;UID=UnPrivilegedUser;PWD=HADRPWD123;AlternateServers=localhost:1600;HASession=1;
```

When a query is executed on the primary deactivated server, the query is blocked by the server until there is an active primary server. The query execution proceeds normally when the same server is reactivated. If there is a failover, the server fails the command with error codes 2379 and 2376. As a part of a processing response of the query, the ODBC driver fails over to the new primary server.

After the successful failover, the ODBC driver generates the `HAFailover` error with the native message number 30130. The ODBC application resets the application context and re-executes the failed query/transaction. To reset the application context and re-execute the failed query/transaction, set the `HASession` connection property to 1 so that the ODBC driver does not fail over to the new primary server.

Note

The default value of the `commandtimeout` connection property is 30 seconds. The ODBC driver cancels the blocked commands after it reaches this time-out values. To delay the cancellation of the blocked commands, adjust the value of the `commandtimeout` connection property to a higher value.

9.6 HADR Messages from SAP ASE

SAP jConnect and the SAP ASE ODBC driver receive messages about HADR from SAP ASE.

Error Code	Message	Description	Action
2376	New primary has been activated.	Available through: <ul style="list-style-type: none"> • ODBC – SQL_ATTR_DR_INFO_MS G or SQL_ATTR_DR_INFO_CA LLBACK • jConnect – SybConnection.getClientInfo() 	Client application must establish a new connection to the newly activated primary.
2377	Primary server is being deactivated.	Available through: <ul style="list-style-type: none"> • ODBC – SQL_ATTR_DR_INFO_MS G or SQL_ATTR_DR_INFO_CA LLBACK • jConnect – SybConnection.getClientInfo() 	Application should complete the current transaction and avoid any new transactions.
2378	Deactivation did not complete. Primary is once again active.	Available through: <ul style="list-style-type: none"> • ODBC – SQL_ATTR_DR_INFO_MS G or SQL_ATTR_DR_INFO_CA LLBACK • jConnect – SybConnection.getClientInfo() 	Client application may resume normal processing.
2379	Primary has been deactivated.	Available through: <ul style="list-style-type: none"> • ODBC – SQL_ATTR_DR_INFO_MS G or SQL_ATTR_DR_INFO_CA LLBACK. • jConnect – SybConnection.getClientInfo() 	Client application should roll back any open transactions and avoid any new transactions.
2380	Primary has been reactivated.	Available through: <ul style="list-style-type: none"> • ODBC – SQL_ATTR_DR_INFO_MS G or 	Client application may resume normal processing.

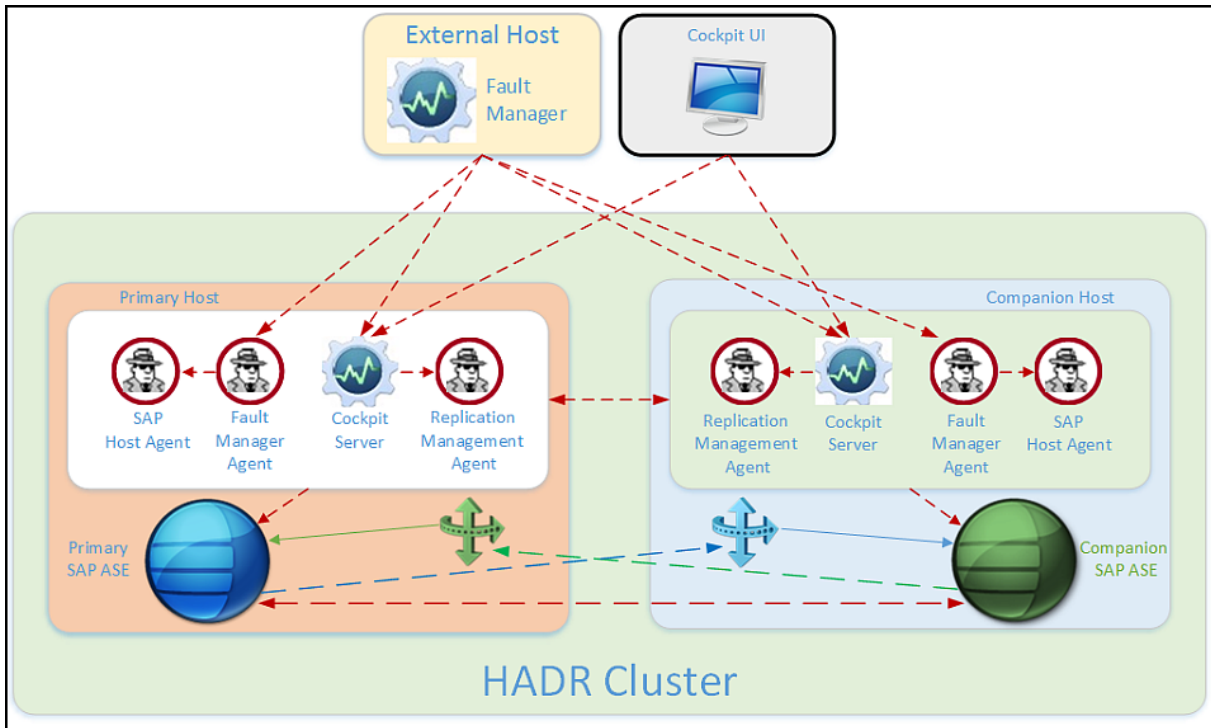
Error Code	Message	Description	Action
		SQL_ATTR_DR_INFO_CALLBACK • jConnect – SybConnection.getClientInfo()	
9662	Unable to issue a checkpoint on database '%.*s'.	Available through jConnect SybConnection.getClientInfo()	See <i>Troubleshooting: Error Messages Advanced Resolutions</i> for general information about resolving SAP ASE errors.
9663	The RepAgent Thread for database '%.*s' is not running.	Available through jConnect SybConnection.getClientInfo()	See <i>Troubleshooting: Error Messages Advanced Resolutions</i> .
9664	At least one RepAgent Thread did not complete its processing of the transaction log.	Available through jConnect SybConnection.getClientInfo()	See <i>Troubleshooting: Error Messages Advanced Resolutions</i> .
9665	Adaptive Server is running in 'Standby' mode. The user does not have 'allow hadr login' privilege.	Available through jConnect SybConnection.getClientInfo()	See <i>Troubleshooting: Error Messages Advanced Resolutions</i> .
9666	Adaptive Server is running in 'Primary Inactive' mode. The user does not have 'allow hadr login' privilege nor the 'hadrgentlecap' capability is set on the connection.	Available through jConnect SybConnection.getClientInfo()	See <i>Troubleshooting: Error Messages Advanced Resolutions</i>
9667	Adaptive Server is running in 'Primary Deactivating' mode. The user does not have 'allow hadr login' privilege nor the 'hadrgentlecap' capability is set on the connection.	Available through jConnect SybConnection.getClientInfo()	See <i>Troubleshooting: Error Messages Advanced Resolutions</i>

Error Code	Message	Description	Action
9668	Login redirection to Primary server failed.	Available through jConnect SybConnection.getClientInfo()	See <i>Troubleshooting: Error Messages Advanced Resolutions</i>
9670	The truncation point of database '%.*s' has not been established with DBCC SETTRUNC.	Available through jConnect SybConnection.getClientInfo()	See <i>Troubleshooting: Error Messages Advanced Resolutions</i>
9671	Deactivation failed due to %s. Resetting state to %s.	Available through jConnect SybConnection.getClientInfo()	See <i>Troubleshooting: Error Messages Advanced Resolutions</i> .
9672	Server reached INACTIVE state.	Available through jConnect SybConnection.getClientInfo()	See <i>Troubleshooting: Error Messages Advanced Resolutions</i>
9673	Initiating log drain mechanism.	Available through jConnect SybConnection.getClientInfo()	See <i>Troubleshooting: Error Messages Advanced Resolutions</i>
9674	User connections statistics:: %d in xact, %d in chained mode, %d in unchained mode, %d holding server side cursors.	Available through jConnect SybConnection.getClientInfo()	See <i>Troubleshooting: Error Messages Advanced Resolutions</i> .
3957	New transaction cannot be started due to an ongoing HADR deactivate operation. The command could not be completed.	Available through jConnect SybConnection.getClientInfo()	See <i>Troubleshooting: Error Messages Advanced Resolutions</i> .
3958	Ongoing transactions have been rolled back due to HADR deactivate.	Available through jConnect SybConnection.getClientInfo()	See <i>Troubleshooting: Error Messages Advanced Resolutions</i> .
15968	%s: parameter %d is not specified.	Available through jConnect SybConnection.getClientInfo()	See <i>Troubleshooting: Error Messages Advanced Resolutions</i> .

10 HADR Internals

The SAP ASE HADR system is built on top of SAP Replication Server technology.

Replication Servers used in an HADR system are embedded in SAP ASE and use synchronous replication mode when configured for high availability (HA). SAP ASE uses asynchronous replication mode when the cluster is configured for disaster recovery (DR).



SAP ASE topology includes one primary SAP ASE and one standby SAP ASE (called the "companion" in HA configuration). Applications access data on the primary SAP ASE. Administrator users may connect to either SAP ASE server, but ordinary users are either rejected or redirected to the primary SAP ASE when they attempt to connect to the standby SAP ASE.

SAP ASE HADR configuration and administration is performed using the `sp_hadr_admin` system procedure. Most of its commands are designed for SAP ASE to use internally. Use only the documented `sp_hadr_admin` parameters, and use them carefully.

Replication topology consists of one source and one target SAP ASE. The RMA module manages SAP ASE and Replication Server state transitions using procedures that start with `sap_` (see [RMA Commands \[page 444\]](#)). Many of these procedures are designed for internal use; use only the documented `sap_` interfaces for administration and monitoring.

Use the SAP installer or the `setuphadr` utility to configure an HADR cluster: Do not use `sap_` commands or the `sp_hadr_admin` system procedure to perform this task.

SAP ASE Cockpit provides cluster-wide state monitoring and the ability to administer some HADR functionality.

The Fault Manager module runs on a separate host, preferably on the host running the application server. It monitors the state of the different modules in the HADR cluster, takes action based on events, and sends alerts through SAP ASE Cockpit. The Fault Manager includes agents on SAP ASE hosts that run as heartbeat clients and work in tandem with the Fault Manager.

The SAP Host Control module runs under `sudo` (root) privilege (on UNIX platforms) on SAP ASE hosts, and provides services to the Fault Manager for running database and operating system commands.

When SAP ASE is in synchronous replication mode, committed transactions are sent to Replication Server on remote hosts, apart from writing to the database log device. The replication state may temporarily switch to an asynchronous state in synchronous mode if there is no response from the remote Replication Server within the amount of time specified by the `max_commit_wait` Replication Agent configuration parameter to prevent application performance due to temporary glitches in network connectivity between the primary SAP ASE and the remote Replication Server. However, automatic failover by the Fault Manager is disabled until SAP ASE resumes replication back to a synchronous state after catch-up. Set `max_commit_wait` to a high value to ensure zero data loss.

When the replication mode and state are synchronous, transactions are committed by SAP ASE after receiving notification from the remote Replication Server that the data modifications have been written to a persistent storage device that provides protection against data loss in the event of host or site failure.

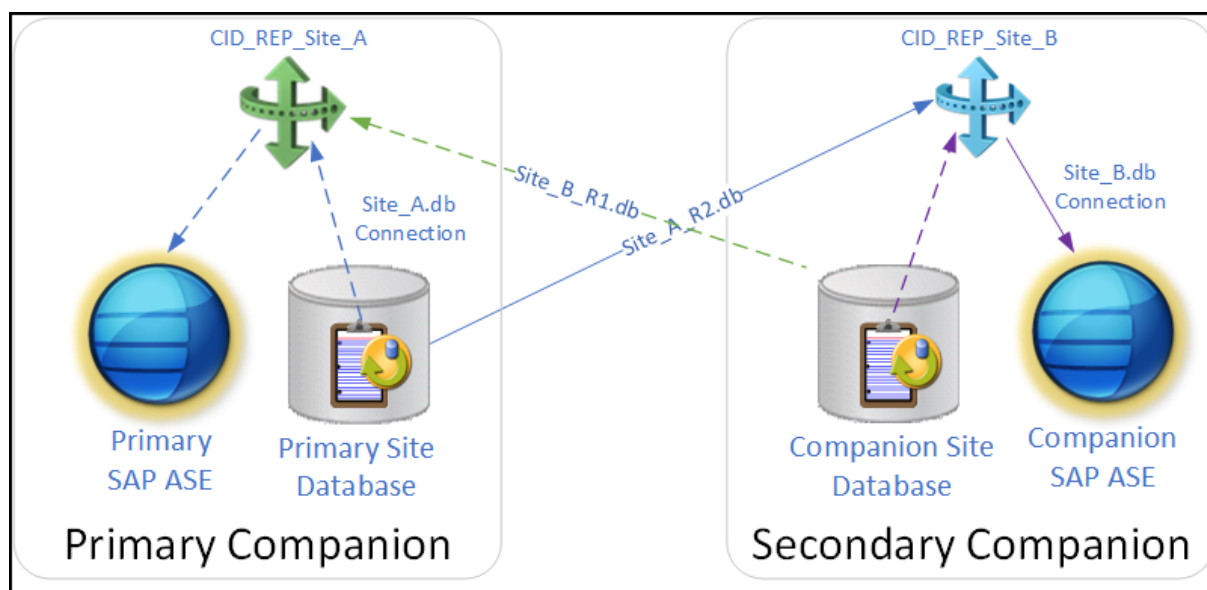
Architecture for Current Connections

10.1 Connections

Two connections are created for each database in an HADR system:

- A local connection to the local Replication Server on the host on which it is running (connection `Site_A.db` in the image below).
- A remote connection to the remote Replication Server on the opposite host (connection `Site_A_R2.db` in the image below).

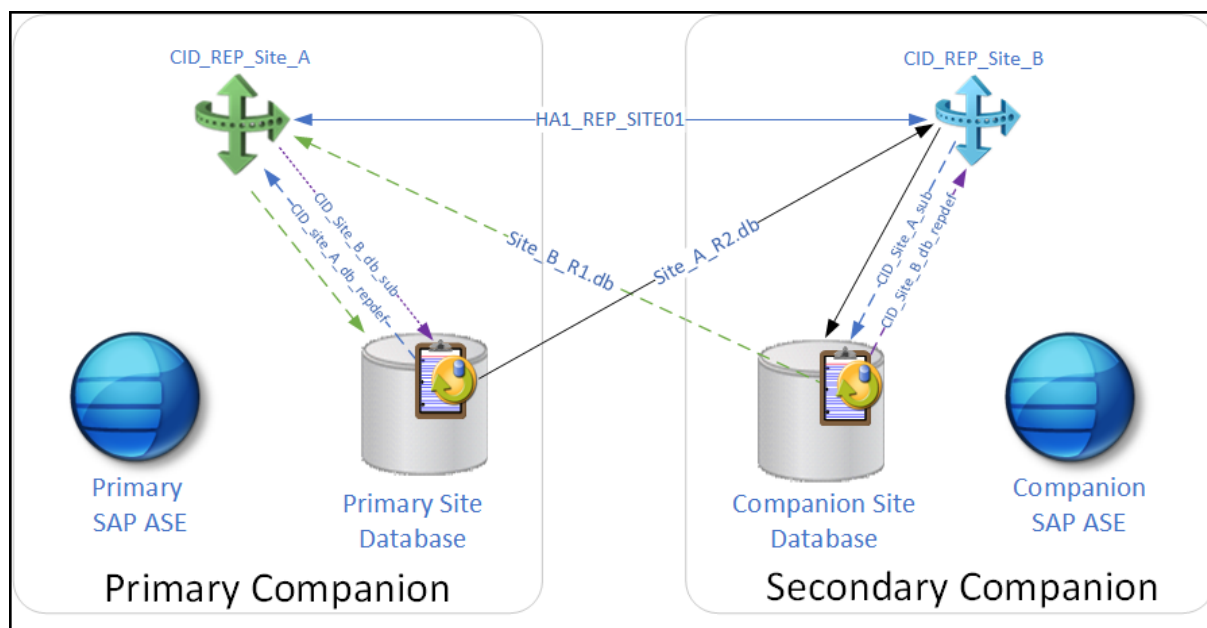
However, only one connect is active at a time, and for an isolated cluster, only the remote connections (those suffixed with the _R1 and R2) are used. In the example below, `site_A_R2.db` connection is active:



During certain failure conditions, local connections may be used with an external Replication Server when the HADR cluster replicates to a third site (for example, such as a reporting database in SAP IQ).

Local multi-site availability replication definitions (repdefs) and subscriptions are similar. Both use:

- The same source names: `<CID>_<site_name>_repdef`. Neither name includes the _R1 or R2 suffixes because they are local.
- The route created between the primary and companion Replication Server for communication. The naming convention is `<CID>_REP_<site>` (for example, `HA1_REP_SITE01`).



The SAP HADR system does not currently use the connection between the two Replication Servers (`HA1_REP_SITE01` in the image above).


The HADR system includes a number of remote, active topology data paths.

- A database repdef is created on each remote source database.
- A database subscription is created on the standby database on this remote database repdef.

The name scheme for these is:

- Remote database repdef – <CID>_<Site>_R2_<database_name>_repdef. For example, HA1_site01_R2_pubs2_repdef.
- Remote database subscription – <CID>_<Site>_R2_<database_name>_sub. For example, HA1_site01_R2_pubs2_pubs2.

All names reference the source connection and site.



Unused connections are suspended. For example, in this output is from the HA1_SITE02 companion server with an active Replication Server (the primary site is named HA1_SITE01), the REP AGENT, NRM, DSI EXEC, and DSI connections (in red) are unused and, consequently, suspended:


```
1> admin who
2> go
```

Spid	Name	State	Info	
88	DSI EXEC	Awaiting Command	107 (1) HA1_SITE02.demo_db	Parallel DSIs are enabled by default
89	DSI EXEC	Awaiting Command	107 (2) HA1_SITE02.demo_db	
90	DSI EXEC	Awaiting Command	107 (3) HA1_SITE02.demo_db	
37	DSI	Awaiting Message	107 HA1_SITE02.demo_db	Local standby IBQ, SQT, and DIST connections are running but not used
45	DSIT	Awaiting Wakeup	107 HA1_SITE02.demo_db	
83	SQT	Awaiting Wakeup	107:1 DIST HA1_SITE02.demo_db	
23	SQM	Awaiting Message	107:1 HA1_SITE02.demo_db	
22	SQM	Awaiting Message	107:0 HA1_SITE02.demo_db	Standby repagent local connection is suspended
	REP AGENT	Down	HA1_SITE02.demo_db	
	NRM	Down	HA1_SITE02.demo_db	Primary DSI remote connections are suspend, and will not be used
	DSI EXEC	Suspended	108 (1) HA1_SITE01_R2.demo_db	
	DSI	Suspended	108 HA1_SITE01_R2.demo_db	Primary IBQ, SQT, and DIST connections are active
46	DSIT	Awaiting Wakeup	108 HA1_SITE01_R2.demo_db	
81	SQT	Awaiting Wakeup	108:1 DIST HA1_SITE01_R2.demo_db	
25	SQM	Awaiting Message	108:1 HA1_SITE01_R2.demo_db	Primary repagent remote connection is active
24	SQM	Awaiting Message	108:0 HA1_SITE01_R2.demo_db	
62	CAP	Awaiting Command	108 HA1_SITE01_R2.demo_db	
	REP AGENT CI	Active	HA1_SITE01_R2.demo_db	
	SPQ WRITER	Dup	108 HA1_SITE01_R2.demo_db	
	SPQ READER	QWait	108 HA1_SITE01_R2.demo_db	

Each SAP ASE server has a proxy connection with the other SAP ASE server in the HADR system. This connection is used primarily for DR_admin, but is also used when local RMA and Host Agents need to verify connections (for example, when the RMA sap_status procedure needs to access the remote node for status information, or if the local Fault Manager heartbeat cannot see either the remote SAP ASE or the Fault Manager, it assumes the SAP ASE is isolated from the network and initiates a change to standby mode). SAP ASE uses the proxy connections to fill out the information in these proxy tables in the master database:

- hadrGetLog
- hadrGetTicketHistory
- hadrStatusActivePath
- hadrStatusResource
- hadrStatusRoute
- hadrStatusSynchronization

The proxy connections in the HADR system look similar to:



11 Troubleshooting

You can troubleshoot many of the Replication Server, SAP ASE Cockpit, and HADR system issues.

11.1 Troubleshooting the HADR System

Troubleshooting the HADR system often includes rectifying permission and space issues.

Diagnosing Problems in the HADR System

View error logs for diagnostic information.

- Check the `dev_sybdbfm` Fault Manager error log for errors – To get diagnostic information, set the trace level to 3 by adding the line `'ha/syb/trace = 3'` to `SYBHA.PFL`. Restart the Fault Manager for the change to take effect.
- Increase the trace level of SAP Host Agent services – Add the line `service/trace = 3` to `/usr/sap/hostctrl/exe/host_profile.SAP Host Agent` by issuing: `/usr/sap/hostctrl/exe/saphostexec -restart`.

These logs will subsequently display additional information:

```
/usr/sap/hostctrl/exe/dev_saphostctrl  
/usr/sap/hostctrl/exe/dev_sapdbctrl
```

- Check the RMA logs, located in `$SYBASE/DM/RMA-16_0/instances/AgentContainer/logs/RMA_*.log`
- Check SAP ASE logs, located in `$SYBASE/ASE-16_0/install/*.log` for any issue related to replication.
- Check the Replication Server log, located in `$SYBASE/DM/<CID>_REP_<logical_site>` and is named `<CID>_REP_<logical_site>.log`.

Common Issues

Permissions and space issues are common problems in the HADR system.

- Verify that all HADR directories have the appropriate permissions, specifically the SAP ASE installation directory, the Fault Manager installation and execution directories, and `/tmp`. The Fault Manager creates temporary directories under `/tmp` and adds temporary files there. If it is unable to do so, calls to the SAP Host Agent fail, but the Fault Manager cannot know that the call failed because it could not add the temporary files. For this reason, verify that the user executing the Fault Manager has permissions on all the directories.

- Verify that `/tmp` has adequate space available. If `/tmp` is full, the Fault Manager cannot create its temporary files. To verify that `/tmp` has space, execute:

```
df -k /tmp
```

If this command shows 100 percent usage, you may have to make room in `/tmp`.

- Verify that the version for GLIBC is 2.7 or later by executing:

```
ldd --version
```

- Make sure you enter the correct passwords for `sa`, `DR_admin`, and `sapadm`. It can be very difficult to find the root cause of the errors when password mismatches are the culprit. By default, `sapadm` may not have a password when you create it with the SAP Host Agent, but requires one for the Fault Manager. Add or update the `sapadm` password using the `passwd` command.
- Verify that the user limits value for `open files` is set to an adequate number (4096 or larger) before configuring the HADR system for large databases. Use this command to view the value for `open files`:

- On the C-shell:

```
limit descriptors
```

- On the Bourne shell:

```
ulimit -a
```

- Verify that there is sufficient amount of memory and swap space before configuring the HADR system for large databases. In particular, materialization requires an adequate amount of memory and swap space for large databases.

Failure to Add an Instance

The HADR system displays an error message when it cannot add an instance. For example:

```
2014 09/12 04:06:42.996 executing: asehostctrl -host star4 -user sapadm
***** -function LiveDatabaseUpdate -dbname JD3 -dbtype syb -dbinstance
STAR4 -updatemethod Execute -updateoption TASK=ADD_ASE_INSTANCE -updateoption
SYBASE_PATH=/star/big_system/ase160 -updateoption SYBASE_USER=important_user -
updateoption
DRA_HOST=star4 -updateoption DRA_PORT=7001 .
2014 09/12 04:06:42.996 starting control call.
NiInit3: NI already initializes (init=6;cur=2048)
NiHLGetNodeAddr: found hostname 'star4' in cache
NiGetNodeAddr: hostname 'star' = addr 10.172.162.61
NiGetServNo: servicename '1128' = port 1128
NiCreateHandle: hdl 6 state NI_INITIAL_CON
NiInitSocket: set default settings for new hdl 6/sock 6 (UD; ST)
NiBlockMode: set blockmode for hdl 6 FALSE
NiConnectSocket: hdl 6 is connecting to /tmp/.sapstream1128 (timeout=-1)
NiHLGetHostName: found address 10.172.162.61 in cache
NiGetHostName: addr 10.172.162.61 = hostname 'star4'
NiConnect: SiConnect failed for hdl 6/sock 6
(SI_ECONN_REFUSE/13; UD; ST; /tmp/.sapstream1128)
NiInitSocket: set default settings for new hdl 6/sock 6 (I4; ST)
NiBlockMode: set blockmode for hdl 6 FALSE
NiConnectSocket: hdl 6 is connecting to 10.172.162.61:1128 (timeout=-1)
SiPeekPendConn: connection of sock 6 broken (111-Connection refused)
NiHLGetHostName: found address 10.172.162.61 in cache
NiGetHostName: addr 10.172.162.61 = hostname 'star4'
```

```

***LOG Q0I=> NiPConnect2: 10.172.162.61:1128: connect (111: Connection refused)
[/bas/CGK_MAKE/src/base/ni/nixxi.cpp 3324]
*** ERROR => NiPConnect2: SiPeekPendConn failed for hdl 6/sock 6
(SI_ECONN_REFUSE/111; I4; ST; 10.172.162.61:1128) [nixxi.cpp 3324]
NiICloseHandle: closing initial hdl 6
2014 09/12 04:06:42.998 dbctrl call cnt reset: 0 .
2014 09/12 04:06:42.998 control call ended.
2014 09/12 04:06:42.999 call_saphostctrl completed ok.
2014 09/12 04:06:42.999 saphostctrl executed.
2014 09/12 04:06:42.999 dbctrl call cnt reset 2: 0 .
2014 09/12 04:06:42.999 ase instance status is: .
2014 09/12 04:06:42.999 add instance status: failure

```

This section of the error message indicates that the system could not create a connection to the hostagent:

```

***LOG Q0I=> NiPConnect2: 10.172.162.61:1128: connect (111: Connection refused)
[/bas/CGK_MAKE/src/base/ni/nixxi.cpp 3324]
*** ERROR => NiPConnect2: SiPeekPendConn failed for hdl 6/sock 6
(SI_ECONN_REFUSE/111; I4; ST; 10.172.162.61:1128) [nixxi.cpp 3324]

```

To resolve the issue, execute this command to check if the `sapstartsrv` process is running (it should be started with the SAP Host Agent):

```
ps -aef | grep sapstartsrv
```

Restarting the HADR Option on Windows

On Windows, start SAP ASE as a service, not with a startup script.

For example, if Host1 (the primary companion) goes down and fails over to Host2 (the secondary companion), restart SAP ASE Host1 as a Windows service instead of using startup scripts such as the `<RUN_server>` file.

11.2 Recovering from Errors in an HADR System

Rectifying errors in the HADR system often involves reviewing the error logs.

The default location for the errors logs are:

- SAP ASE – `$SYBASE/$SYBASE_ASE/install/server_name.log`
- Backup Server – `$SYBASE/$SYBASE_ASE/install/server_name.log`
- Replication Server – `$SYBASE/DM/log/rep_server.log`
- RMA – `$SYBASE/DM/RMA-16_0/instances/AgentContainer/logs/RMA_*.log`

The following examples assume the configuration described in this RMA `sap_set` command:

```

sap_set
go
PROPERTY                                VALUE
-----                                -
maintenance_user                        NW7_maint
sap_sid                                  NW7
installation_mode                        BS

```

```

participating_databases [master,NW7]
NW7SEC, sybase_dir /sybase/NW7_REP
NW7SEC, hadrflag ha
NW7SEC, ase_port 4901
NW7SEC, ase_user DR_admin
NW7SEC, ase_backup_server_port 4902
NW7SEC, ase_hostname mo-699d5a593
NW7SEC, db_dump_dir /sybase/NW7_REP
NW7SEC, delay 0
NW7SEC, ase_instance NW7_NW7SEC
NW7SEC, synchronization_mode sync
NW7SEC, distribution_mode remote
NW7SEC, distribution_target NW7PRI
NW7SEC, replication_port 4905
NW7SEC, rssid_port 4906
NW7SEC, rssid_ra_port 4907
NW7SEC, replication_user DR_admin
NW7SEC, device_buffer_dir /sybase/NW7_REP/repdata_1
NW7SEC, rs_hostname mo-699d5a593
NW7SEC, device_buffer_size 30720
NW7SEC, simple_persistent_queue_dir /sybase/NW7_REP
NW7SEC, simple_persistent_queue_size 2000
NW7SEC, simple_persistent_queue_max_size 2000
NW7SEC, rs_instance NW7_REP_NW7SEC
NW7SEC, internal_name NW7SEC
NW7SEC, dr_plugin_port 4909
NW7PRI, sybase_dir /sybase/NW7_REP
NW7PRI, hadrflag ha
NW7PRI, ase_port 4901
NW7PRI, ase_user DR_admin
NW7PRI, ase_backup_server_port 4902
NW7PRI, ase_hostname mo-4f28f910f
NW7PRI, db_dump_dir /sybase/NW7_REP
NW7PRI, delay 0
NW7PRI, ase_instance NW7_NW7PRI
NW7PRI, synchronization_mode sync
NW7PRI, distribution_mode remote
NW7PRI, distribution_target NW7SEC
NW7PRI, replication_port 4905
NW7PRI, rssid_port 4906
NW7PRI, rssid_ra_port 4907
NW7PRI, replication_user DR_admin
NW7PRI, device_buffer_dir /sybase/NW7_REP/repdata_1
NW7PRI, rs_hostname mo-4f28f910f
NW7PRI, device_buffer_size 15360
NW7PRI, simple_persistent_queue_dir /sybase/NW7_REP
NW7PRI, simple_persistent_queue_size 2000
NW7PRI, simple_persistent_queue_max_size 2000
NW7PRI, rs_instance NW7_REP_NW7PRI
NW7PRI, internal_name NW7PRI
NW7PRI, dr_plugin_port 4909

```

Path is Suspended from Primary Server to the Companion Server

Symptom: The connection from the primary to the companion server is suspended.

In this situation, issuing the `sap_status path` RMA command results in a message similar to the following from the RMA command line, where the line in bold indicates that a server is down:

```

./OCS-16_0/bin/isql -UDR_admin -P<pass_word> -S<host_name>:<RMA_port_number>
> sap_status path
> go

```

```
Path Name Value Info
```

```
-----  
-----  
Start Time 2015-04-15 03:10:54.379 Time command started executing.  
Elapsed Time 00:00:00 Command execution time.  
NW7PRI Hostname mo-2897e9422.mo.sap.corp Logical host name.  
NW7PRI HADR Status Primary : Active Identify the primary and standby sites.  
NW7SEC Hostname mo-338995c0a.mo.sap.corp Logical host name.  
NW7SEC HADR Status Standby : Inactive Identify the primary and standby sites.  
NW7PRI.NW7SEC.NW7 State Suspended Path is suspended. Transactions are not being  
replicated.  
NW7PRI.NW7SEC.NW7 Failing Command Error Message Failing Command Error Message  
The connection in the Replication Server on the NW7SEC host to  
'NW7_NW7SEC.NW7' is not configured.  
NW7PRI.NW7SEC.NW7 Latency Unknown No latency information for database 'NW7'.  
NW7PRI.NW7SEC.NW7 Commit Time Unknown No last commit time for the database  
'NW7'.  
NW7PRI.NW7SEC.master State Active Path is active and replication can occur.  
NW7PRI.NW7SEC.master Latency Time 2015-04-15 02:03:33.980 Time latency last  
calculated  
NW7PRI.NW7SEC.master Latency 540 Latency (ms)  
NW7PRI.NW7SEC.master Commit Time 2015-04-15 02:03:33.020 Time last commit  
replicated  
NW7SEC.NW7PRI.NW7 State Suspended Path is suspended. Transactions are not being  
replicated.  
NW7SEC.NW7PRI.NW7 Latency Unknown No latency information for database 'NW7'.  
NW7SEC.NW7PRI.NW7 Commit Time Unknown No last commit time for the database  
'NW7'.  
NW7SEC.NW7PRI.master State Suspended Path is suspended. Transactions are not  
being replicated.  
NW7SEC.NW7PRI.master Latency Unknown No latency information for database  
'master'.  
NW7SEC.NW7PRI.master Commit Time Unknown No last commit time for the database  
'master'.
```

The output describes four paths: two should be active and two should be suspended. The two that should be active are the paths that include the primary companion. In this example, the paths that start with "NW7PRI".

Action: To reactivate a suspended path:

1. Use the Replication Server `admin who_is_down` command to determine which server is down (see the *Replication Server Reference Manual > Replication Server Commands*). In this example, `admin who_is_down` indicates that path `NW7_NW7SEC.NW7` is down, which is associated with thread number 106:

```
admin who_is_down  
Spid Name State Info  
-----  
-----  
REP AGENT Suspended NW7_NW7SEC.master  
NRM Suspended NW7_NW7SEC.master  
DSI EXEC Suspended 106(1) NW7_NW7SEC.NW7  
DSI EXEC Suspended 106(2) NW7_NW7SEC.NW7  
DSI EXEC Suspended 106(3) NW7_NW7SEC.NW7  
DSI Suspended 106 NW7_NW7SEC.NW7  
REP AGENT Suspended NW7_NW7SEC.NW7  
NRM Suspended NW7_NW7SEC.NW7
```

2. Retry the connection to see if a timeout caused the failure:

```
resume connection to NW7_NW7SEC.NW7
```

3. If the `resume connection` command fails, issue either of these commands to purge all data from the queue that is being removed:

⚠ Caution

Using the `sqlm_purge_queue` and `resume connection ... skip tran` parameters can cause a data mismatch between the standby and primary companions. They are very risky operations, and you must perform a rematerialization once they are finished.

→ Tip

You can choose to run the `sysadmin sqlm_purge_queue` command to purge queues, without necessarily hibernating on the Replication Server. Instead, you can suspend the appropriate modules in the Replication Server, and then purge queues as usual. Running `sysadmin sqlm_purge_queue` with the `[, check_only]` parameter facilitates this scenario, as it checks and reports if the appropriate modules were suspended successfully (it **does not** purge queues), thus enabling you to make an informed decision before purging queues. Note that you can continue to purge queues like you did before – by hibernating on the Replication Server. For more information, see the *Usage* section under *SAP Replication Server Reference Manual > SAP Replication Server Commands > sysadmin sqlm_purge_queue*.

- `resume connection ... skip tran` to resume the connection but skip the indicated transactions.
- `sysadmin sqlm_purge_queue` to purge all messages from a stable queue.

4. Execute the Replication Server `sysadmin hibernate_on` command to enable the server hibernation mode:

```
sysadmin hibernate_on
```

5. Execute the Replication Server `sysadmin sqlm_purge_queue` command (the combination '106, 0', below, tells `sqlm_purge_queue` to operate on queue 106:0):

```
sysadmin sqlm_purge_queue,106, 0
```

6. Disable hibernation:

```
sysadmin hibernate_off
```

7. Execute the Replication Server `resume connection` command to resume the suspended connection:

```
resume connection to NW7_NW7SEC.NW7
```

8. Execute the RMA `sap_status` command to verify that the suspended path is now active:

```
sap_status path
Path Name Value Info
-----
-----
Start Time 2015-04-15 03:37:34.342 Time command started executing.
Elapsed Time 00:00:01 Command execution time.
NW7PRI Hostname mo-2897e9422.mo.sap.corp Logical host name.
NW7PRI HADR Status Primary : Active Identify the primary and standby sites.
NW7SEC Hostname mo-338995c0a.mo.sap.corp Logical host name.
NW7SEC HADR Status Standby : Inactive Identify the primary and standby sites.
NW7PRI.NW7SEC.NW7 State Active Path is active and replication can occur.
NW7PRI.NW7SEC.NW7 Latency Time 2015-04-15 02:03:33.180 Time latency last
calculated
NW7PRI.NW7SEC.NW7 Latency 340 Latency (ms)
NW7PRI.NW7SEC.NW7 Commit Time 2015-04-15 03:08:13.340 Time last commit
replicated
NW7PRI.NW7SEC.master State Active Path is active and replication can occur.
```

```

NW7PRI.NW7SEC.master Latency Time 2015-04-15 02:03:33.000 Time latency last
calculated
NW7PRI.NW7SEC.master Latency 550 Latency (ms)
NW7PRI.NW7SEC.master Commit Time 2015-04-15 02:03:33.040 Time last commit
replicated
NW7SEC.NW7PRI.NW7 State Suspended Path is suspended. Transactions are not
being replicated.
NW7SEC.NW7PRI.NW7 Latency Unknown No latency information for database 'NW7'.
NW7SEC.NW7PRI.NW7 Commit Time Unknown No last commit time for the database
'NW7'.
NW7SEC.NW7PRI.master State Suspended Path is suspended. Transactions are not
being replicated.
NW7SEC.NW7PRI.master Latency Unknown No latency information for database
'master'.
NW7SEC.NW7PRI.master Commit Time Unknown No last commit time for the
database 'master'.

```

RMA Command Failure

Symptom: An RMA command fails in the HADR system.

An RMA command failure results in an error message similar to the following, with the error being indicated with an Error Task, **Task State Error**, indicated in bold:

```

sap_failover NW7PRI, NW7SEC, 600
go
TASKNAME      TYPE                VALUE
-----
Failover      Start Time          Tue Apr 07 08:32:59 UTC 2015
Failover      Elapsed Time        00:00:00
Failover      Task Name           Failover
Failover      Task State          Error
Failover      Short Description   Failover moves primary responsibility
from current logical source to logical target.
Failover      Long Description    Validating properties required for
failover
have been
set.
Failover      Failing Command     sap_failover NW7PRI, NW7SEC, 600, 600
Failover      Failing Command Error Message Invalid optional failover argument
'600'.
Valid arguments are 'force' and
'unplanned'.
Failover      Corrective Action   Execute sap_failover with 'force' to
force
deactivation to occur when timeout period expires or with 'unplanned' to force
deactivation
(if the primary ASE is available) and force failover (if the primary ASE is
NOT available).
Failover      Hostname            mo-f6a45b44b.mo.sap.corp

```

Action: Search the error log for a reason why the RMA command failed (all RMA commands start with "TDS LANGUAGE"), including:

- Find the timestamp of the execution indicated by "Start Time" in the sap_failover output. In the example, this is "Tue Apr 07 08:32:59 UTC 2015".
- Search the RMA error log for the "TDS LANGUAGE: sap_failover" keywords near the time indicated by the timestamp (in this case, "04-07-15 07:53:47").

- Determine whether failover completed all required steps. Search the RMA error log for "Task Failover - Is performing step - " keywords. Failover should include 14 steps.
- Determine if failover completed. Search the error log for "Task Failover - State - Completed" keywords. If failover did not complete, the error log includes the "Task Failover - State - Error" keywords.
- Determine if the `host_available` command failed (indicating the host was not available). If `host_available` failed, search for the "TDS LANGUAGE: sap_host_available" and "Task HostAvailable - Is performing step - " keywords.
- Analyze the error messages in RMA
- Analyze SAP ASE and Replication Server error logs on the primary and companion servers.

Inconsistent Data

Symptom: Data is inconsistent between the primary and companion servers.

Action: Rematerialize the servers using these RMA commands (in these examples, the primary server is named `PRI`, and the companion server is named `STA`):

1. Disable the replication on the primary server:

```
sap_disable_replication PRI,NW7
```

2. Re-enable replication:

```
sap_enable_replication PRI,NW7
```

3. Disable replication on the primary server's master database:

```
sap_disable_replication PRI,master
```

4. Re-enable replication on the primary server's master database:

```
sap_enable_replication PRI,master
```

5. Rematerialize the master database on the primary and companion servers:

```
sap_materialize auto,PRI,STA,master
```

i Note

The maintenance user password is changed and managed by Replication Server after you run `sap_materialize`, preventing the database administrator from accessing the data of primary and standby databases.

6. Rematerialize the primary and companion servers on host NW7:

```
sap_materialize auto,PRI,STA,NW7
```

SAP ASE Log Is Full

Symptom: SAP ASE reports that its log is full.

Action: Dump the database without making a backup copy:

```
dump transaction <database_name> with truncate_only
```

⚠ Caution

Use `dump transaction with no_log` as a last resort, and use it only once after `dump transaction with truncate_only` fails.

The `with truncate_only` and `with no_log` parameters allow you to truncate a log that has become dangerously short of free space. Neither parameter provides a means to recover transactions that have committed since the last routine dump.

See the *SAP ASE Reference Manual: Commands*.

Multiple Standby Servers after a Split-Brain Check

Symptom: The primary server starts as the standby server if the primary server is started before the companion server.

Action:

1. Log on the actual primary server.
2. Move to the master database:

```
use master
```

3. Promote the server to the primary role:

```
sp_hadr_admin primary
```

4. Resume user application transaction activity on the primary server:

```
sp_hadr_admin activate
```

Replication Agent Is Down

Symptom: Issuing the `sap_status path` RMA command results in a suspended path because the Replication Agent on SAP ASE is down:

```
1> sap_status path
2> go
PATH          NAME          VALUE
INFO
-----
-----
```

```

-----
-----
command started executing.
Start Time                2015-04-17 03:35:57.526 Time
Elapsed Time              00:00:01                Command
execution time.
NW7SEC                   Hostname                mo-f9bb75e82          Logical
host name.
NW7SEC                   HADR Status            Standby : Inactive    Identify
the primary and standby sites.
NW7SEC                   Synchronization Mode   Synchronous          The
configured Synchronization Mode value.
NW7SEC                   Synchronization State   Inactive
Synchronization Mode in which replication is
currently operating.
NW7SEC                   Distribution Mode       Remote
Configured value for the distribution_mode
replication model property.
NW7SEC                   Replication Server Status Active                The
status of Replication Server.
NW7PRI                   Hostname                mo-f6a45b44b          Logical
host name.
NW7PRI                   HADR Status            Primary : Active      Identify
the primary and standby sites.
NW7PRI                   Synchronization Mode   Synchronous          The
configured Synchronization Mode value.
NW7PRI                   Synchronization State   Inconsistent
Replication Synchronization Mode is currently
operating as 'Inconsistent'
instead of the configured value 'Synchronous'. Please refer to Replication
Server logs for more details.
NW7PRI                   Distribution Mode       Remote
Configured value for the distribution_mode
replication model property.
NW7PRI                   Replication Server Status Active                The
status of Replication Server.
NW7PRI.NW7SEC.NW7       State                   Suspended             Path is
suspended (Replication Agent Thread).
Transactions are not being
replicated.
NW7PRI.NW7SEC.NW7       Additional Info         Additional Info        The
REPLICATION AGENT connection in the
Replication Server on the mo-f9bb75e82 host to 'NW7_NW7PRI_R2.NW7' is
suspended.
NW7PRI.NW7SEC.NW7       Latency Time           2015-04-17 03:30:04.238 Time
latency last calculated
NW7PRI.NW7SEC.NW7       Latency                436                  Latency
(ms)
NW7PRI.NW7SEC.NW7       Commit Time            2015-04-17 03:35:12.566 Time last
commit replicated
NW7PRI.NW7SEC.NW7       Distribution Path       NW7SEC                The path
of Replication Server through which
transactions travel.
NW7PRI.NW7SEC.master State Active                Path is
active and replication can occur.
NW7PRI.NW7SEC.master Latency Time 2015-04-17 03:30:03.020 Time
latency last calculated
NW7PRI.NW7SEC.master Latency      0                    Latency
(ms)
NW7PRI.NW7SEC.master Commit Time 2015-04-17 03:30:04.412 Time last
commit replicated
NW7PRI.NW7SEC.master Distribution Path NW7SEC                The path
of Replication Server through which
transactions travel.
NW7SEC.NW7PRI.NW7 State Suspended Path is
suspended (Replication Agent Thread).
Transactions are not being replicated.

```

NW7SEC.NW7PRI.NW7 Latency latency information for database 'NW7'.	Unknown	No
NW7SEC.NW7PRI.NW7 Commit Time commit time for the database 'NW7'.	Unknown	No last
NW7SEC.NW7PRI.NW7 Distribution Path of Replication Server through which transactions travel.	NW7PRI	The path
NW7SEC.NW7PRI.master State suspended (Replication Agent Thread). Transactions are not being replicated.	Suspended	Path is
NW7SEC.NW7PRI.master Latency latency information for database 'master'.	Unknown	No
NW7SEC.NW7PRI.master Commit Time commit time for the database 'master'.	Unknown	No last
NW7SEC.NW7PRI.master Distribution Path of Replication Server through which transactions travel.	NW7PRI	The path

Action: Restart the Replication Agent.

To determine which Replication Agent to restart, find the name of the primary and companion servers, the path marked "suspended," the location where the error originates, and on which host the error originates.

From the output above:

- Primary and companion servers – These lines indicate that NW7PRI is the primary server and that NW7SEC is the companion server:

```
NW7PRI HADR Status Primary : Active Identify the primary and standby sites
NW7SEC HADR Status Standby : Inactive Identify the primary and standby sites
```

- Suspended path – Paths that start with NW7PRI should be active, but sap_status path reports that NW7PRI.NW7SEC.NW7 is "suspended."
- Error originates – This line NW7PRI.NW7SEC.NW7 indicates that the suspension error results from the Replication Agent thread: "State Suspended Path is suspended (Replication Agent Thread). Transactions are not being replicated."
- Host on which the error originates – This line NW7PRI.NW7SEC.NW7 indicates that the Replication Agent thread on host mo-f9bb75e82 is stopped: Additional Info Additional Info The REPLICATION AGENT connection in the Replication Server on the mo-f9bb75e82 host to 'NW7_NW7PRI_R2.NW7' is suspended.

To resolve this issue, restart the Replication Agent running on host mo-f9bb75e82:

1. Log in SAP ASE on host mo-f9bb75e82:

```
$$SYBASE/$SYBASE_OCS/bin/isql -Usa -Ppassword -Smo-f9bb75e82
```

2. Issue sp_start_rep_agent against database NW7 to start the Replication Agent:

```
sp_start_rep_agent NW7
go
Replication Agent thread is started for database 'NW7'.
(return status = 0)
```

3. Issue sp_status path to verify that the path is now active.

Primary SAP ASE Cannot be Activated

Symptom: after using `sap_disable_replication` to disable the replication and then restarting the primary SAP ASE server, the primary SAP ASE server cannot be activated any more and the user applications fail to connect to the primary databases.

`sap_disable_replication` only stops the RepAgent, but does not disable the RepAgent and set the second truncation point (STP) to end. However, SAP ASE cannot be activated if one RepAgent is enabled and the STP is not valid.

Action: either set the STP to end manually or disable the RepAgent manually on the primary SAP ASE.

- To set the STP to end manually and then activate the primary SAP ASE:

```
use master
go
dbcc settrunc('ltm','end')
go
use ERP
go
dbcc settrunc('ltm','end')
go
use master
go
sp_hadr_admin activate
go
use master
go
dbcc settrunc('ltm','ignore')
go
use ERP
go
dbcc settrunc('ltm','ignore')
go
```

- To disable the RepAgent and then activate the primary SAP ASE:

```
sp_config_rep_agent master,disable
go
sp_config_rep_agent 'ERP',disable
go
sp_hadr_admin activate
go
```

11.3 Recovering the Primary Data Server If SAP Replication Server is Unavailable

If SAP Replication Server is unavailable during an SAP ASE startup after an unplanned failover, use SAP ASE commands to recover a database that is enabled for synchronous replication, and make it accessible online.

Context

If the replication mode is synchronous for the primary data server and SAP Replication Server is unavailable during SAP ASE startup after an unplanned failover, SAP ASE cannot recover the original primary data server and make it assume the role of a standby data server, since SAP ASE cannot connect to SAP Replication Server to obtain information about the last transaction that arrived at SAP Replication Server. For example, you see the following if the database name is D01 and <dbid> represents the database ID, in the SAP ASE error log:

```
Error: 9696, Severity: 17, State: 1
Recovery failed to connect to the SAP Replication Server to get the last oqid for
database 'D01'.
Database 'D01' (dbid <dbid>): Recovery failed.
Check the ASE errorlog for further information as to the cause.
```

Procedure

1. Check the SAP ASE error log to see if the latest attempt to connect to the SAP Replication Server failed.
2. Verify that the original primary database has not been recovered.
For example, if the database name is D01, log in to `isql` and enter:

```
use D01
go
```

If the database has not been recovered, you see:

```
Msg 921, Level 14, State 1:
Server 'paris', Line 1:
Database 'D01' has not been recovered yet - please wait and try again.
```

3. Use `isql` to connect to SAP Replication Server and execute `admin who` to verify that SAP Replication Server has restarted and is receiving connections.
4. In SAP ASE, enable trace flag 3604 to log all events and any errors that occur during database recovery:

```
dbcc traceon(3604)
go
```

5. Recover D01 and make it available for online access:

```
dbcc dbrecover(D01)
go
```

The recovery is successful and the database is accessible online if you see the events logged by the trace flag ending with:

```
...
Recovery complete.
Database 'D01' is now online.
DBCC execution completed. If DBCC printed error messages, contact a user with
System Administrator (SA) role.
```

6. Verify that the database is recovered and can be accessed.

```
use D01
go
```

11.4 Restarting the Primary Data Server Without Synchronization

You can use the `dataserver --recover-syncrep-no-connect` parameter to restart the primary SAP ASE data server without synchronization to SAP Replication Server if you cannot restart SAP Replication Server during an unplanned failover.

Context

During failover of SAP ASE from primary to standby, SAP ASE needs to connect to SAP Replication Server to query the last transaction it received from the primary ASE `dataserver`. When this connection is not possible, the following error is logged in the ASE error log, and the replicated database(s) are not recovered:

```
Error 9696: "Recovery failed to get the last oqid for database '<name>' from SAP
Replication Server because it was either unable to connect or it received an
error".
```

The `--recover-syncrep-no-connect` parameter starts SAP ASE and tries to connect to the SAP Replication Server during recovery. If the connection attempts to SAP Replication Server fail, error 9696 is not invoked. SAP ASE recovers the databases, but the primary and standby databases may not be synchronized. Without synchronized replication between the databases and SAP Replication Server, you cannot recover from an unplanned failover with the assurance of no data loss that synchronized replication provides.

Procedure

1. Add the `dataserver --recover-syncrep-no-connect` parameter to the `run_server` file to restart the primary data server without synchronization:

```
...
dataserver --recover-syncrep-no-connect
...
```

2. Execute the `run_server` file.

11.5 Installation

Troubleshooting a failed installation involves rectifying the SAP installer, recovering from a failed setup, and performing a teardown.

11.5.1 SAP Installer Issues

Troubleshooting the SAP installer often includes rectifying the Replication Server configuration and materialization.

Troubleshooting Replication Configuration

Perform these steps if the SAP installer encounters an error when configuring the Disaster Recovery environment:

1. Check the SAP installation log for errors; it shows errors from the SAP installer perspective.
2. If the error cannot be resolved, check the RMA log, located at `$SYBASE/DM/RMA-16_0/instances/AgentContainer/logs/RMA_*.log` on the primary Replication Server machine.
3. Examine SAP ASE and Replication Server log files for installation errors.
After the error is resolved, retry the configuration.

Materialization Fails While the Database Is in Use

If the SAP installer encounters an error during materialization, and you determine after investigating the logs that the database being materialized is in use:

1. Log in to the standby SAP ASE and issue `sp_who` to determine which processes are using the database.
2. For any existing processes, have its associated user log off the server to remove the process. As a last resort, use the `kill` command to remove the process.

After the error is resolved, retry the configuration.

Materialization Fails

If the SAP installer encounters an error during materialization, perform the following if you determine after investigating the logs that you must perform materialization again (you must first reset replication).

1. Log in to the RMA:

```
isql -w999 -U<user> -P<password> -S<host:port>
```


2. Issue these commands:

```
sap_disable <source_logical_host_name>, <database>  
sap_enable <source_logical_host_name>, <database>
```

When you resolve the error, retry the configuration by clicking the [Retry](#) button in the SAP installer.

For more information about RMA commands, see the *RMA Configuration and User Guide* or issue `sap_help` at the command line.

Net Password Encryption Required

Perform these steps if the SAP installer encounters an error during materialization and after investigating the logs you determine the net password encryption required is out of sync.

1. Log into the SAP ASE that is out of sync.
2. Reset the `net_password_encryption_reqd` configuration parameter. Setting it to:
 - 0 – (the default) allows the client to determine whether to use encryption
 - 1 – requires an encryption algorithms
 - 2 – requires strong password encryption

11.5.2 Recovering from a Failed Setup

You can recover from failed setups from the SAP installer and from the `setuphadr` utility.

11.5.2.1 Recovering from a Failed HADR Setup from the Installer

Perform tasks on the first and second site to recover from a failed setup.

On the First Site

1. If the HADR setup failed, click [Next](#) to complete the installation. The installer starts RMA.
2. Check the `setuphadr` utility log file, located in `$SYBASE/ASE-16_0/init/logs` for the cause of failure, and correct it.
3. Enter the passwords in `setuphadr` utility response file, located in `$SYBASE/ASE-16_0/setuphadr.rs`.
4. Execute this command to finish the setup using the `setuphadr` utility:

```
setuphadr $SYBASE/ASE-16_0/setuphadr.rs
```

i Note

If RMA failed to start, check the RMA log in `$SYBASE/DM/RMA-16_0/instances/AgentContainer/logs`. Make any required corrections, then start RMA.

- (UNIX) – `$$SYBASE/$SYBASE_ASE/bin/rma`
- (Windows) – start the RMA Windows service by:
 - Starting *Sybase DR Agent* - `<cluster_ID>` from the Services panel, or
 - Issuing this command, where `<cluster_ID>` is the ID of the cluster:

```
net start SybaseDRAgent_<cluster_ID>
```

On the Second Site

1. If the HADR setup failed, click *Next* to complete the installation. The installer starts RMA.
2. Check the `setuphadr` utility log file in `$$SYBASE/ASE-16_0/init/logs` and the RMA log in `$$SYBASE/DM/RMA-16_0/instances/AgentContainer/logs` directories to find out why the setup failed, and make any required corrections.
3. Perform the following depending on when the HADR setup failed:

- If HADR setup failed before the `Setup Replication` task in `setuphadr` utility log file:
 1. Enter the passwords in `setuphadr` utility response file, located in `$$SYBASE/ASE-16_0/setuphadr.rs`.
 2. Issue this command to finish the setup using the `setuphadr` utility:

```
setuphadr $$SYBASE/ASE-16_0/setuphadr.rs
```

- If HADR setup failed during, or after, the `Setup Replication` task in `setuphadr` utility log file:
 1. Perform the teardown with the instructions in [Troubleshooting the Replication System \[page 410\]](#) > *Recovering Replication Server*.
 2. If `$$SYBASE/ASE-16_0/setuphadr.rs` does not exist on first site, copy it from second site, then:
 - Enter the passwords
 - Set `is_secondary_site_setup` property to `false`
 - Set the value of the `setup_site` property to the first site
 3. On the first site, run `setuphadr` utility with the edited `setuphadr.rs` responses file:

```
setuphadr $$SYBASE/ASE-16_0/setuphadr.rs
```

4. Edit the `$$SYBASE/ASE-16_0/setuphadr.rs` response file on second site to include the passwords.
5. On the second site, run `setuphadr` utility with the edited `setuphadr.rs` responses file:

```
setuphadr $$SYBASE/ASE-16_0/setuphadr.rs
```

11.5.2.2 Recovering from a Failed setuphadr Setup

How you recover from failed `setuphadr` setup depends on when the failure occurred.

1. Check the `setuphadr` utility log file in `$$SYBASE/ASE-16_0/init/logs` and the RMA log in `$$SYBASE/DM/RMA-16_0/instances/AgentContainer/logs` directories for the reason the setup failed, and make any required corrections.
2. If the HADR setup failed before the `Setup Replication` task in `setuphadr` utility log file, rerun `setuphadr` on the current site.

i Note

Remove any existing RMA service before you rerun `setuphadr` on Windows 64-bit.

3. If the HADR setup failed during, or after, the `Setup Replication` task in `setuphadr` utility log file:
 1. Perform a teardown according to the instructions in [Performing a Teardown \[page 403\]](#).
 2. Rerun the `setuphadr` utility on first site.
 3. Rerun the `setuphadr` utility on the second site.

11.5.3 Performing a Teardown

Performing a teardown requires tasks on the primary and standby sites.

The steps described in this section require you to issue the `sap_teardown` command, which automatically performs these tasks:

- Stops the Replication Server and deletes its instance directory, partition files, and simple persistent queue directories, and kills all Replication Server related processes.
- Deactivates the primary SAP ASE, then changes its mode to standby, if the source host (the machine on which SAP ASE runs) is available.
- Drops all servers from the HADR server list on both SAP ASE servers.
- Drops the HADR group from both servers.
- Disables HADR on both servers.
- Disables CIS RPC Handling.

i Note

- The `sap_teardown` command does not drop the logins for the administrator or maintenance user. Drop and re-create these logins after running `sap_teardown`.
- Clean up the SPQ directories on each host after running `sap_teardown`, otherwise you may encounter errors when re-creating the HADR system.

11.5.3.1 Performing a Teardown using `sap_teardown`

Tearing down a replication environment includes disabling replication in the SAP ASE servers, stopping the SAP Replication Servers, and deleting all directories and files created during setup, including the SAP Replication Server instances.

After the teardown is complete, the system is no longer an HADR system. The SAP ASE is left running after the teardown and should be treated like a regular, SMP server.

Use the `sap_teardown` command to tear down the replication environment. The command does not modify any data that has been replicated to the standby databases. Additionally, the databases on both the primary and standby hosts are not unmarked for replication. The command does not remove any software, but it does remove the SAP Replication Servers and configurations that support replication. Executing `sap_teardown`:

- Disables Replication Agents and secondary truncation points.
- Shuts down and deletes SAP Replication Server instances, including their stable queue files.
- Does nothing to the data in the standby database (the data remains current as of when replication was last active).
- Deletes these directories:
 - Instance directories and their contents on the primary and standby Replication Servers
 - On the primary server – `$$SYBASE/DM/<SID_REP_primary_logical_host_name>`
 - On the standby server – `$$SYBASE/DM/<SID_REP_standby_logical_host_name>`
 - Any device files created by Replication Server

The primary and standby dump directories are not deleted during teardown. The dump directories are defined using `sap_set` and setting the `db_dump_dir` property. These directories can get very large depending on the amount of data materialized. It is the responsibility of the user to maintain these directories.

The primary and standby device directories are not deleted during teardown. These dump directories are defined using `sap_set` and setting property, `device_buffer_dir`.

To perform a teardown using `sap_teardown`:

1. Log in to the DR Agent:

```
isql -UDR_admin -PSybase123 -Shost1:8899
```

2. Execute:

```
sap_teardown
```

11.5.3.2 Manually Removing HADR Replication

You should use the `sap_teardown` command to perform a teardown.

If you need to manually remove HADR replication:

1. Log into the primary and standby SAP ASE servers and remove the HADR proxy tables:

```
use master
go
drop table hadrGetTicketHistory
go
drop table hadrGetLog
go
drop table hadrStatusPath
go
drop table hadrStatusResource
go
drop table hadrStatusRoute
go
```

2. Log into the primary and standby SAP ASE servers and remove these Replication Server system objects from the `master` and participating databases:

```
drop procedure rs_get_lastcommit
go
drop procedure rs_syncup_lastcommit
go
drop procedure rs_update_threads
```

```

go
drop procedure rs_initialize_threads
go
drop procedure rs_marker
go
drop procedure rs_check_repl_stat
go
drop procedure rs_ticket
go
drop procedure rs_ticket_v1
go
drop procedure rs_ticket_report
go
drop procedure rs_send_repserver_cmd
go
drop table rs_lastcommit
go
drop table rs_threads
go
drop table rs_ticket_history
go
drop table rs_dbversion
go
drop table rs_mat_status
go

```

3. Log into the primary server to remove and disable HADR member information:

```

sp_hadr_admin deactivate, '300', 'teardown', 'force', 'nodrain'
go
sp_hadr_admin standby, force
go
sp_hadr_admin dropserver, '<CID>_<standby_logical_host_name>'
go
sp_hadr_admin dropserver, '<CID>_<primary_logical_host_name>'
go
sp_hadr_admin dropserver, '<CID>_<primary_logical_host_name>_DRA'
go
sp_hadr_admin dropgroup, '<group_name>'
go
sp_configure 'HADR mode',-1
go

```

4. Log into the standby server to remove and disable HADR member information:

```

sp_hadr_admin dropserver, '<CID>_<primary_logical_host_name>'
go
sp_hadr_admin dropserver, '<CID>_<standby_logical_host_name>'
go
sp_hadr_admin dropserver, '<CID>_<standby_logical_host_name>_DRA'
go
sp_hadr_admin dropgroup, '<group_name>'
go
sp_configure 'HADR mode',-1
go

```

5. Shut down Replication Server on the primary and standby sites.
6. Remove the Replication Server instance directory from the primary site at `$$SYBASE/DM/CID_REP_primary_logical_host_name`.
7. Remove the Replication Server instance directory from the standby site at `$$SYBASE/DM/CID_REP_standby_logical_host_name`.
8. Shut down RMA on the primary and standby sites.

9. Remove the RMA configuration database by removing all files and subdirectories from:

```
$SYBASE/DM/RMA-16_0/instances/AgentContainer/configdb/*
$SYBASE/DM/RMA-16_0/instances/AgentContainer/backups/*
```

10. Restart RMA.

11.5.4 Removing an HADR Environment Using the removehadr Utility

Use the `removehadr` utility to clean up an HADR environment.

The `removehadr` utility performs the following:

- Performs a shutdown of the current running RMA instance.
- Deletes the RMA instance configurations from the `config` and `configdb` directories.
- Deletes the RMA logs from the `logs` directory.
- Drops `DR_admin`, the default ASE administration and maintenance user.

If `HADR mode` is enabled and there is no running RMA instance, `removehadr` performs the following:

- Removes the SAP Replication Server instance.
- Performs a shutdowns of the running SAP Replication Server processes, such as `repserver`, `dbltm`, and `dbsrv17` in Linux and `repsrvr.exe`, `dbltm.exe`, and `dbsrv17.exe` in Windows.
- Deletes the replication-related tables in SAP ASE, which all use the `rs_` prefix.
- Deletes the proxy tables in SAP ASE.
- Drops the HADR servers and groups and resets `HADR mode` to `disabled`.
- Stops the RepAgent processes and disables the RepAgent.

Execute the `sap_teardown` command before running the `removehadr` utility or else the utility logs on to the RMA server and finishes the teardown process.

The `removehadr.sh` (for Linux) or `removehadr.cmd` (for Windows) files as well as `removehadr.jar` are present in the `$Sybase/RMA-16_0/bin` directory.

The syntax to use the `removehadr` utility is:

```
removehadr.sh [-R <res_file>] [-U <sa_username>] [-P <sa_password>] [-D
<DR_admin_password>] [-I <interface_file>] [-S <server_name>]
```

Where:

- `<res_file>` - is the path of the resource file that stores the HADR installation information.
- `<sa_username>` - is the `sa` username used to connect to SAP ASE and RMA.
- `<sa_password>` - is the `sa` login password used to connect to SAP ASE and RMA.
- `<DR_admin_password>` - is the DR admin password used to connect to DR and RMA.
- `<server_name>` - is the SAP ASE server name that is used to connect.
- `<interface_file>` - is the interface file used for connection.

Examples:

1. Example 1

This example displays all the parameter options for `removehadr`:

```
removehadr.sh -h, -help, --help
```

```
-h,--help print this help
-R <res_file> resources file for HADR to be removed
-U <sa_username> ASE sa login name
-P <sa_password> ASE sa login password
-D <DR_admin_password> DR_admin login password
-S <server name> ASE server name
-I <interface file> interfaces or sql.ini file
```

2. Example 2

This example removes the HADR environment with the given *.rs file. The resource file here refers to the same .res file `setuphadr` used to set up the HADR environment:

```
removehadr.sh -R setup.rs
```

```
Executing ASE Command: 'sp_configure 'HADR mode''
Executing ASE Command: 'use master'
Executing ASE Command: 'sp_role 'grant',replication_role,sa'
Executing RMA Command: 'sap_teardown'
Shutting down RMA instance.
Dropping DR_admin user....
....
```

3. Example 3

This example removes the HADR environment by using the SAP ASE `sa` login:

```
removehadr.sh -Usa -Dsybase -Ssite0:5000
```

```
SA user password:
sybase
Executing ASE Command: 'use master'
Executing ASE Command: 'sp_configure 'HADR mode''
Executing ASE Command: 'use master'
Executing ASE Command: 'sp_role 'grant',replication_role,sa'
Executing RMA Command: 'sap_teardown'
Shutting down RMA instance.
Dropping DR_admin user....
```

4. Example 4

This example removes the HADR environment by using the `sa` login, with an interfaces file and a server logical name:

```
removehadr.sh -Usa -Psybase -Dsybase -I/opt/sap/DM/interfaces -SPI2_PR
```

11.6 Monitoring

You can monitor the health of the Replication Server in the HADR system.

Latency is measured as the time the commit was received at the primary server until the time the commit was seen on the standby server.

- Monitor the health of Replication Server in the HADR environment
- Issue an `rs_ticket` command at regular intervals, approximately once every 10 to 30 minutes, and check the `rs_ticket_history` table on the standby server to verify that the ticket was received, and whether there is any latency.
- Check the Replication Server log to make sure no errors are reported. If there are warnings of memory hitting its threshold, increase the amount of memory available to Replication Server (or verify that Replication Server rectifies the memory shortage).
- From Replication Server:
 - Issue `admin who_is_down` to determine if any components are down.
 - Issue `admin who, sqm` to check if the queues are moving.
 - Check the size of the `SPQ`, `InBound`, and `OutBound` queues.

11.6.1 Utilities for Monitoring the HADR System

SAP includes a number of utilities to monitor the HADR system.

Check the health of replication using the Replication Server `admin` command:

- Check if any components or connections in Replication Server are down:

```
admin who_is_down
```

- Check the flow of the `sqm` parameter:

```
admin who, sqm
```

- Check the disk space for the Replication Server partitions:

```
admin disk_space
```

- Check the current memory used:

```
admin stats, mem_in_use
```

- Check the maximum amount of memory used so far:

```
admin stats, max_mem_use
```

- View details of the memory used and maximum amount of memory used:

```
admin stats, mem_detail_stats
```

11.6.2 Monitoring the Replication Agent

There are a number of monitoring tables for monitoring the Replication Agent

Including:

- `monRepCoordinator`
- `monRepLogActivity`

- `monRepMemoryStatistics`
- `monRepScanners`
- `monRepScannersTotalTime`
- `monRepSchemaCache`
- `monRepStreamStatistics`
- `monRepSyncTaskStatistics`
- `monRepTruncationPoint`

See the *SAP ASE Reference Manual: Tables > Monitoring Tables* and *Performance and Tuning Series: Monitoring Tables* for information about using these monitoring tables.

11.6.2.1 Monitoring the Scanners

RegAgent has a `syslogs` scanner and a `sysimrlogs` scanner if in-memory row storage (IMRS) is enabled on a database.

Use the following monitoring tables to monitor the scanners:

- `monRepLogActivity`
- `monRepMemoryStatistics`
- `monRepScanners`
- `monRepScannersTotalTime`
- `monRepSchemaCache`
- `monRepStreamStatistics`

For an IMRS database, the `sysimrlogs` scanner is enabled by default. You can turn on the trace flag 9126 to disable it if necessary.

See the *Reference Manual: Tables > Monitoring Tables* for more information about these tables.

11.7 Replication

Troubleshoot Replication Server issues to rectify replication, RMA, and unreplicated data issues.

11.7.1 Troubleshooting the Replication System

There are a number of steps you can perform to troubleshoot the replication system.

Determining if Replication Is Enabled for Primary Connection

After you enable replication for a primary connection, executing the `admin who` command from the `isql` prompt on Replication Server displays the status of the Synchronous Replication components. For example:

```
CAP           Awaiting Command      110 QG30.tpcc12
REP AGENT CI  Active                 QG30.tpcc12
SPQ WRITER   Dup                   110 QG30.tpcc12
SPQ READER   QWait                 110 QG30.tpcc12
```

Where:

- `CAP` – the Capture component reads stream replication packages from the simple persistent queue (SPQ). It translates stream replication commands into Replication Server commands and writes them into inbound queue. `CAP` is one of:
 - `Awaiting Command` – waiting for message from the SPQ reader.
 - `Active` – processing a package from the SPQ.
 - `Down` – Capture has failed and shut down, or Replication Server is in hibernation mode.
 - `Suspended` – suspended.
- `REP AGENT CI` - the Replication Agent or log transfer components. `REP AGENT CI` is one of:
 - `Active` – Replication Agent is connected.
 - `Down` – Replication Agent is not connected.
 - `Suspended` – Log transfer is suspended.
- `SPQ WRITER` – the inbound stream replication connection. It receives messages from Replication Agent and writes them into the SPQ. `SPQ WRITER` is one of:
 - `QWait` – waiting due to a full writer queue.
 - `Dup` – detects a duplicate message.
 - `Writing` – currently writing a message to the SPQ file.
 - `Ready` – waiting for incoming message from Replication Agent.
 - `Down` – inactive because Replication Agent is down.
- `SPQ READER` - the outbound stream replication connection. It reads stream replication packages from SPQ and sends them to Capture. `SPQ READER` is one of:
 - `QWait` – waiting due to a full reader queue.
 - `NCWait` – detects NC commands in SPQ.

- Reading – reading a message from the SPQ file.
- Ready – no additional messages in the SPQ file.
- Down – the SPQ reader is not active because Capture is down.

Verifying that Replication Agent is Connected to Replication Server

Use the `admin who` from the `isql` prompt on Replication Server to check the status of the Replication Agent (displayed as `REP_AGENT CI` in the output). The status is Active if the Replication Agent is connected.

When Replication Agent connects to Replication Server, it includes messages similar the following in the Replication Server log file:

```
T. 2015/02/11 23:27:57. (212): RATCI(QG30.tpcc123) connected to repserver
```

If Replication Server indicates that Replication Agent is running on the primary database but is not connected, use `sp_help_rep_agent <database_name>, process` to check the status of the Replication Agent process. Connect to the SAP ASE acting as the primary node and execute:

```
sp_help_rep_agent pdb,process
go
Replication Agent Coordinator Process Status
 dbname spid sleep_status state
-----
  pdb      58 sleeping      sleeping
(1 row affected)

Replication Agent Scanner Process Status
 dbname spid start_marker end_marker current_marker sleep_status state
-----
  pdb      60 (13399,39)  (13399,39) (13399,39)  sleeping      sleeping

Replication Agent Truncation Point Manager Process Status
 dbname spid trunc_pts_confirmed trunc_pts_processed total_processing_time
longest_process_time idle_time sleep_status state
-----
  pdb      59             0             99             95
0
```

See the Replication Server *Reference Manual* > *SAP ASE Commands and System Procedures* for information about `sp_help_rep_agent`.

Issue `sp_who` on SAP ASE to determine if the Replication Agent is running on the primary SAP ASE. Connect to the SAP ASE acting as the primary node and execute:

```
sp_who
go
fid  spid  status      loginame  origname  hostname  blk_spid  dbname  tempdbname
cmd  block_xloid threadpool
-----
:.....:
0    58    background  NULL     NULL     NULL     0    pdb    tempdb    REP
AGENT
0    59    background  NULL     NULL     NULL     0    pdb    tempdb    REP
AGENT CI STPMGR
0    60    background  NULL     NULL     NULL     0    pdb    tempdb    REP
AGENT CI SCANNER
```

Collecting More Information About the Stream Replication Libraries

Use trace information to collect information about the stream replication libraries.

Enable the traces with this syntax:

```
trace 'on', ci, <trace_flag>
```

Where `<trace_flag>` is one of:

- `CI_TRACE_DUPLICATE` – logs information about duplicate stream replication packages.
- `CI_TRACE_GRP1` – displays trace connection related information in stream replication libraries.
- `CI_TRACE_GRP2` – displays trace data path related information in stream replication libraries.
- `CI_TRACE_GRP3` – displays trace buffer management related information in stream replication libraries.

Collecting Information About the SPQ Layer

Use trace information to collect information about the SPQ layer.

Enable the traces with this syntax:

```
trace 'on', spq, <trace_flag>
```

Where `<trace_flag>` is one of:

- `SPQ_TRACE_PACKAGE` – dumps information about every stream replication package processed (read or write) in the SPQ.
- `SPQ_TRACE_DISPATCHER` – logs package dispatcher activities in the SPQ.
- `SPQ_IGNORE_TRUNCATION` – discards truncation point movement requests.

Collecting Information About the Capture Layer

Use trace information to collect information about the capture module. Enable the traces with this syntax:

```
trace 'on', cap, <trace_flag>
```

Where `<trace_flag>` is one of:

- `CAP_TRACE_CI_SCHEMAS` – dumps every schema in stream replication packages received by Capture.
- `CAP_TRACE_CI_PACKAGES` – dumps information about all stream replication package received by Capture (for example, last OQID, number of commands, and number of schemas).
- `CAP_TRACE_CI_COMMANDS` – dumps all commands in the stream replication packages received by Capture.
- `CAP_TRACE_COMMANDS` – dumps all commands from stream replication packages in a SQL format.

Error Messages: "Capture fails with a bad row buffer..." or "Capture receives a corrupted row buffer"

A message is written to the Replication Server error log file when Capture fails to parse a command and shuts down.

The full message is one of the following:

```
Capture fails with a bad row buffer. Row buffer length=%d, Column name=%s, Column type=%d, Max column length=%d, Column length=%d, Column offset=%d.
```

```
Capture receives a corrupted row buffer. Row buffer dump(%d bytes): <hex dump>
```

Action – Enable Capture trace to collect more information and contact an SAP Replication Server administrator (see the following section).

Capture Fails and Shuts Down Due to Schema Issue

Capture issues an error message to the Replication Server error log file when it fails and shuts down.

The full text is one of the following:

```
Schema <schema_number> with version <version_number> is not found.
```

```
Inconsistent schema for table <table_name>.
```

```
Cannot decode a row for table <table_name>. There are <number_of_columns> columns in the schema, but there are %d columns in the data.
```

Action – Use the `sysadmin` Replication Server command from the `isql` prompt to collect more information and contact an SAP Replication Server administrator. The syntax is:

```
sysadmin dump_schemas [, cache | rssid [, <dbid> [, <sid> [, ver]]]]
```

The `sysadmin` command dumps table schemas to the schema cache or the Replication Server System Database (RSSD). The Replication Agent sends table schema when a row of the table is processed for the first time. See the *Replication Server Reference Manual*.

Capture adds these schema to a cache and persists them in the RSSD to parse the SAP ASE raw row buffer. When parsing a command, Capture gets the schema it requires from the cache. However, if it is not in the cache, it is loaded from the RSSD.

Purging SPQ Prior to Rematerializing the Database

Use the `sysadmin` Replication Server command to purge the SPQ.

The syntax is:

```
sysadmin spq_purge_queue, <dsname>, <dbname>
```

Before you purge the SPQ, suspend the log transfer and capture, or make sure Replication Server is in hibernation mode.

To purge the SPQ, from the Replication Server's `isql` command line:

1. Suspend the log transfer:

```
suspend log transfer from <server_name.database_name>
```

2. Suspend the capture:

```
suspend capture <server_name.database_name>
```

3. Purge the queue:

```
sysadmin spq_purge_queue, <server_name> , <database_name>
```

4. Resume the log transfer:

```
resume log transfer from <server_name.database_name>
```

5. Resume the capture:

```
resume capture <server_name.database_name>
```

sap_set_host Fails Because Agent Cannot Contact Another Agent

RMA displays an error similar to the following when an issued `sap_set_host` command cannot connect to another agent (see bold text):

```
# isql -UDR_admin -Psybase -Slocalhost:7001 -w300
sap_set_host HA site1 5000 site1 5005 7001
  TASKNAME    TYPE
  VALUE

-----
-----

Set Host      Start Time                Thu Jun 23 02:13:06 EDT
2016

Set Host      Elapsed Time
00:00:00

SetHost       Task Name                  Set
Host

SetHost       Task State
Error
```

```

SetHost      Short Description      Define a logical host with hostname
and port for both ASE and Replication
Server.

SetHost      Long Description          Obtaining the Agent that represents
local host 'HA' and Replication Server
'site1'.

SetHost      Failing Command Return Code
0

SetHost      Failing Command Error Message SetHost failed due to an unexpected
runtime
exception.

SetHost      Failing Exception          SQL ERROR -- Message: 'JZ006: Caught
IOException: java.net.ConnectException: Connection refused (port 7001 to
address ::ffff:192.168.253.129 (site1))' SQLState: 'JZ006' Remote Code: '0'
Cause: ''.
SetHost      Task Start                Thu Jun 23 02:13:06 EDT
2016

SetHost      Task End                  Thu Jun 23 02:13:06 EDT
2016

SetHost      Hostname                  site0

(12 rows affected)

```

To resolve the issue:

1. Log in to the remote RMA and run `sap_AgentInfo` to review agent connection information to use in a subsequent step to alter the agent connection:

```

./OCS-16_0/bin/isql -UDR_admin -P<pass_word> -S<host_name>:<RMA_port_number>
sap_AgentInfo

```

`sap_AgentInfo` displays agent information, including:

- o Host_Name
- o HOST_IP
- o HOST_FULL_NAME
- o HOST_NAME
- o RMI_PORT
- o TDS_PORT
- o JMX URL

2. Add this line to the `$$SYBASE/DM/RMA-16_0/instances/AgentContainer/config/bootstrap.prop` file:

```

rsge.bootstrap.debug=true

```

3. Shut down RMA.
4. Restart RMA:
 - o (UNIX) – `$$SYBASE/$SYBASE_ASE/bin/rma`
 - o (Windows) – Start the RMA Windows service by:

- o Starting *Sybase DR Agent* - <cluster_ID> from the Services panel, or
- o Issuing this command, where <cluster_ID> is the ID of the cluster:

```
net start SybaseDRAgent_<cluster_ID>
```

The RMA displays information about its starting address, hostname resolution, and binding information.

5. Add this information specifying the IP address to the \$SYBASE/DM/RMA-16_0/instances/AgentContainer/config/bootstrap.prop file:

```
java.rmi.server.hostname=<IP_Address>
```

6. Shut down, then restart RMA.

Determining if Rematerialization Is Necessary After Unplanned Failover

Unplanned failovers may result in lost data (even in synchronous mode), and you may need to rematerialize the databases to resynchronize the data.

When you execute `sap_status` after executing `sap_failover`, the RMA produces messages that contain the phrases: "Additional Info 2" or "Corrective Action". Messages that do not include the "Corrective Action" phrase means that you need not rematerialize the database. The steps provided by the "Corrective Action" describe the steps you need to perform to rematerialize the databases.

- If Replication Server is configured for, and is running in, synchronous mode – RMA produces this message:

```
Additional Info 2: The primary Replication Server '{0}:{1}' is configured for synchronization mode and was found running in synchronization mode.
```

In this case, rematerialization of the database is not necessary. However, for the two cases below, you may need to rematerialize the databases.

- If Replication Server is configured for, but is not running in, synchronous mode – RMA produces this message:

```
Additional Info 2: The primary Replication Server '{0}:{1}' is configured for synchronization mode however it is NOT running in synchronization mode.
```

The corrective action information is:

```
Corrective Action: Rematerialize site '{0}' from site '{1}' if necessary.
```

- If the version of Replication Server does not support synchronous mode – RMA produces this message:

```
Corrective Action: Manually verify data is in-sync between sites.
Rematerialize site '{0}' from
site '{1}' if necessary
```

Recovering Replication Server

There are a number of situations for which you need to recover and rebuild the HADR system.

If the primary server reports that the database error log is full, replication will be slow, and you can rematerialize the databases.

If the standby server reports that the database error log is full, the replication path is likely broken. In this case, check the log for errors. Fix any replication issues you find in the log. If there are none, you may need to increase the log on the standby server.

To recover the HADR system when SAP ASE reports that the database log is full:

- Disable the Replication Server from the HADR system by logging in to the RMA and running:

```
sap_disable_replication Primary_Logical_Site_Name
```

After you disable replication, the HADR system preserves the mode and state so that all HADR-aware applications continue to log into the primary companion to execute their work. Synchronize the standby and primary companions by rematerializing the master and CID databases (the second TP is removed from the primary companion for the master and CID databases):

1. Enable replication:

```
sap_enable_replication <Primary_Logical_Site_Name>
```

2. Materialize the master database:

```
sap_materialize auto, <Primary_Logical_Site_Name>,  
<Standby_Logical_Site_Name>, master
```

i Note

The maintenance user password is changed and managed by Replication Server after you run `sap_materialize`, preventing the database administrator from accessing the data of primary and standby databases.

3. To manually materialize the CID database, disable the automatic database dump and backup process, and verify a dump is not currently running:

```
sap_materialize start, <Primary_Logical_Site_Name>,  
<Standby_Logical_Site_Name>, $CID_<DBNAME>
```

To automatically materialize the database, issue:

```
sap_materialize auto, <Primary_Logical_Site_Name>,  
<Standby_Logical_Site_Name>, <User_DB_Name>
```

See [sap_materialize \[page 479\]](#).

- Initiate the teardown to run the primary companion in standalone mode and shut down the HADR system:

```
sap_teardown
```

`sap_teardown` disables the HADR system on the primary and companion servers. After the teardown is finished, the application server can connect to any SAP ASE server because it is running in standalone mode. The Replication Server is shut down and is available to reconfigure.

i Note

The `sap_teardown` command does not drop the logins for the administrator or maintenance user. Drop and recreate these logins after running `sap_teardown`.

11.7.2 Troubleshooting the RMA

Execute these commands to perform general administration for RMA:

- `sap_help` – displays a list of commands and their usage.
- `sap_version` – displays the versions of all the servers (SAP ASE, Replication Server, RMA).
- `sap_set` – displays all the settings in the environment, including ports, host names, and user names.
- `sap_status <path>` – displays the path state, or status of replication along a path.

Examine the log file (located in `$SYBASE/DM/RMA-16_0/instances/AgentContainer/logs/RMA_*.log`) to view commands sent to RMA and the servers in the replication environment.

Hidden Maintenance User

The RMA supports the Replication Server hidden maintenance user password.

Replication Server periodically changes the maintenance user password. After executing `sap_teardown`, you may need to reset the SAP ASE maintenance user password in both servers before configuring a new replication environment.

Use `sp_password` to reset the SAP ASE maintenance user password. For example:

```
sp_password Sybase123, Sybase123, D01_maint
```

Executing Commands from Local and Remote RMAs

Do not execute RMA administrative commands (for example, `sap_setup_replication` or `sap_failover`) simultaneously on the local and remote RMAs. However, you can execute the `sap_status` commands concurrently from either RMA.

SAP ASE and Replication Server Domains

The RMA assumes that the SAP ASE and Replication servers running on a logical host use the same network domain. If they do not, the HADR setup fails during the ERP database materialization and issues an error stating that SAP ASE could not communicate with a remote SAP ASE.

Use `sp_hadr_admin addserver` to manually change the network name in SAP ASE. Use `sap_set` to verify the name change after you have added all the hosts using `sap_set_host`. Check the `rs_hostname` and `ase_hostname` for each logical hostname to confirm the fully qualified domain names use the same suffixes.

SAP ASE Dump Directory

The RMA cannot validate the dump directory's location or permissions when SAP ASE and Replication Server are located on separate host computers.

.sqlanywhere16 Directory

Replication Server uses an SQL Anywhere database to host the embedded RSSD.

A directory named `.sqlanywhere16` is created in the operating system user's home directory when you create this database. The SQL Anywhere-embedded RSSD continues to function correctly if the `.sqlanywhere16` directory is accidentally deleted, and SQL Anywhere writes the directory in another location if the home directory does not exist.

However, SQL Anywhere also creates another directory to store temporary files. If you set the `<SATMP>` environment variable, SQL Anywhere uses this location to store its temporary file. To set `<SATMP>` in the:

- Bourne shell – issue `export SATMP=/tmp/.SQLAnywhere`
- C shell – issue `setenv SATMP /tmp/.SQLAnywhere`

If `<SATMP>` not set, SQL Anywhere uses the value specified by the `<TMP>` environment variable for the location of the temporary files. If `<TMP>` is not set, it uses the value specified by `<TMPDIR>`. If `<TMPDIR>` is not set, it uses the value specified by `<TEMP>`. If none of these environment variables are set, the temporary files are created in `/tmp`.

11.7.3 Checking RMA Version from the Executable

You can quickly determine the version of RMA from the executable using the `-v` or `-version` arguments.

- On a UNIX system, issue:
 - `$$SYBASE/$SYBASE_ASE/bin/rma -v`
 - `$$SYBASE/$SYBASE_ASE/bin/rma -version`
- On a Windows system, issue:
 - `%SYBASE%\%SYBASE_ASE%\bin\rma.bat -v`
 - `%SYBASE%\%SYBASE_ASE%\bin\rma.bat -version`

11.7.4 Troubleshooting Data That is Not Replicating

Data may not replicate because the inbound or the outbound queues are full.

Inbound Queue is Full

Symptom: The inbound queue (IBQ) is reported as full because the downstream components (DIST or DSI) are suspended, or they are shut down due to another issue. Replication Server displays a message similar to this in the log, indicating the partition is full:

```
'WARNING: Stable Storage Use is Above 90 percent.'  
'SQM_ADD_SEGMENT ('%d:%d'): Going to wait for a segment to be freed.  
Administrative action is needed'
```

Recovery Procedure: Check the data accumulated in the IBQ. After fixing the issues, resume the suspended component or restart the failed component.

Symptom: The IBQ is reported as full because the downstream components (Distributor or DSI) cannot keep up with the upstream components (Capture or Replication Agent), and as a result, data accumulates in the IBQ. The Replication Server error log includes one or more message similar to this, indicating the current status of the SRS partition:

```
'WARNING: Stable Storage Use is Above 90 percent.'  
'SQM_ADD_SEGMENT ('%d:%d'): Going to wait for a segment to be freed.  
Administrative action is needed'
```

In this situation:

- The `admin disk_space` command indicates that the Replication Server partition is full.
- The `admin who` command indicates that all components on the path are running.
- The `admin who, sqm` command indicates there is backlog in the IBQ.

Recovery Procedure:

- If the Replication Server's partition is too small (for example, the partition size is 100 MB or 1 GB), issue a command similar to this to add more space:

```
create partition <partition_name> on '<path_to_partition>' with size '10000'
```

- If the IBQ is reported as full due to poor DIST performance, consider tuning the Distributor so that it can keep up with upstream Capture and Replication Agent components. In this situation, the `admin who, sqm` command indicates there is little or no backlog in the outbound queue (OBQ).
- If the IBQ is reported as full due to poor DSI performance, tune the DSI to keep up with Distributor, Capture, and Replication Agent. In this situation, `admin who, sqm` indicates a backlog in the IBQ and OBQ.

Symptom: The IBQ is reported as full due to open transactions. In this situation, when the open transaction is discovered, the IBQ cannot be truncated and eventually fills up. The Replication Server error log includes a message similar to this, indicating the partition is full:

```
' WARNING: Stable Storage Use is Above 90 percent.'  
'SQM_ADD_SEGMENT ('%d:%d'): Going to wait for a segment to be freed.  
Administrative action is needed'
```

In this situation:

- The `admin who` command indicates that all the components on the path are running.
- The `admin disk_space` command indicates that the Replication Server partition is full.
- The `admin who, sqm` command indicates backlog in the IBQ.
- The `admin who, sqt` command indicates open transaction in the IBQ.

Recovery Procedure: Check the reason for the open transaction:

- If the open transaction occurs because Replication Agent disconnects and then reconnects without sending a `purge open` command, purge the open transactions by issuing:

```
sysadmin purge_first_open, q_number, q_type
```

- If the open transaction is a long-running transaction, add additional partitions:

```
create partition <partition_name> on '<path_to_partition>' with size '10000'
```

Outbound Queue is Full

Symptom: The outbound queue (OBQ) is reported as full because DSI is suspended, or is down and issues a message similar to:

```
'WARNING: Stable Storage Use is Above 90 percent.'  
'SQM_ADD_SEGMENT ('%d:%d'): Going to wait for a segment to be freed.  
Administrative action is needed'
```

Recovery Procedure: After fixing the issues, resume the suspended DSI or restart the failed DSI.

Symptom: The OBQ is reported as full because the DSI cannot keep up with the Distributor, Capture or Replication Agent upstream components, and as a result, data accumulates in the OBQ. The Replication Server's log message indicates that the partition is exhausted:

```
'WARNING: Stable Storage Use is Above 90 percent.'  
'SQM_ADD_SEGMENT ('%d:%d'): Going to wait for a segment to be freed.  
Administrative action is needed.'
```

In this situation:

- The `admin who` command indicates that all the components on the path are running
- The `admin disk_space` command indicates that the Replication Server partition is full
- The `admin who, sqm` command indicates backlog in the OBQ

Recovery Procedure:

- If the Replication Server's partition is too small (for example, the partition size is 100 MB or 1 GB), issue a command similar to this to add more space:

```
create partition <partition_name> on '<path_to_partition>' with size '10000'
```

- If the OBQ is reported as full due to poor DSI performance, consider tuning the DSI component so that it can keep up with upstream Distributor, Capture, and Replication Agent components.

- If the OBQ is reported as full due to poor RDB performance, tune the RDB to keep up with the DSI component.

SPQ is Full

Symptom: The SPQ is reported as full because the downstream Capture, DIST, or DSI components are suspended, or they are shut down due to another issue. In this situation, the Replication Agent stops sending messages to the Replication Server, and replication eventually stops. The Replication Server error log indicates the current SPQ status:

```
'The SPQ instance is full. Truncation needs to occur to free space.'
```

Recovery Procedure: Either disable replication or enlarge the SPQ (see the next recovery procedure). Use this syntax to disable replication:

```
sap_disable_replication <primary_logical_hostname> [, <database>]
```

Symptom: The SPQ is reported as full because the downstream component cannot keep up with the Replication Agent, resulting in data accumulating in the SPQ. The Replication Server error log includes one or more messages similar to this, indicating the current status of the SPQ:

```
'The SPQ instance is full. Truncation needs to occur to free space.'
```

Recovery Procedure:

- If the SPQ is configured too small (that is, the value of `spq_max_size` is significantly less), you may need to increase the size of SPQ. For example, if the maximum size is 100 MB or 1 GB, increase the SPQ size by issuing:

```
suspend log transfer from dsname.dbname
go
suspend capture dsname.dbname
go
alter connection to dsname.dbname set spq_max_size to '20480'
go
resume capture dsname.dbname
go
resume log transfer from dsname.dbname
go
```

- If the SPQ is reported as full because of any issue related to Capture, you can tune the Capture so that it can keep up with Replication Agent. In this situation, the `admin who, sqm` command indicates that there is little or no backlog in the IBQ and OBQ:

```
Suspend capture dsname.dbname
go
Alter connection to dsname.dbname set cap_prs_num_threads to '4'
go
Resume capture dsname.dbname
go
```

- If the SPQ is reported as full due to poor DSI performance, you can tune the DSI to keep up with the upstream components. In this situation,
 - The `admin who, sqm` command indicates that the IBQ or OBQ contains a lot of backlog.

- The `admin disk_space` command indicates that the Replication Server partition is full.

11.7.5 Troubleshooting a Secondary Truncation Point That is Not Moving

A situation in which the transaction log of the primary SAP ASE continues to grow but issuing a dump transaction does not free up space may indicate that the transaction log needs more space or that the secondary truncation point does not move.

SAP ASE uses truncation points to ensure that only transactions processed by the Replication Agent are truncated. A secondary truncation point marks the place in the primary database log up to which the Replication Agent has processed transactions. The Replication Agent periodically updates the secondary truncation point to reflect transactions successfully passed to the Replication Server. SAP ASE does not truncate the log past the secondary truncation point. See the *Troubleshooting Guide* for more information about truncation points.

To determine if the secondary truncation point does not move, connect to the SAP ASE acting as primary node , and execute a `select` statement from `syslogshold` at regular intervals. For example:

```
select * from master..syslogshold
2> go
  dbid  reserved    spid  page        xactid          masterxactid
starttime
name
   xloid
-----
-----
-----
-----
-----
 4      0        0    13399 0x000000000000 0x000000000000
Aug 26 2015 11:20PM
$replication_truncation_point
0
```

In this output:

- The row that includes the name equal to `$replication_truncation_point` displays data related to the secondary truncation point.
- The `page` column contains the page number to which the secondary truncation point is pointing.

If the `page` value does not change during multiple executions of `select * from master..syslogshold`, check if you have a long-running transaction exhausting the transaction log. If you do not, check if Replication Server is running and that Replication Agent is connected to it. If the system is replicating, connect to Replication Server and issue `admin who`.

If the status of the `downSPQ READER` has a value of `NCwait`, this may indicate that the secondary truncation point cannot be moved because there are `NC` (non-confirmed) transactions. To verify this, create a dummy table, mark it for replication, insert data into the table, and issue `select * from master..syslogshold` to see whether the secondary truncation point moves.

11.8 Performance

Use the `rs_ticket` Replication Server stored procedure to troubleshoot performance issues.

To start troubleshooting performance issues, execute an `rs_ticket 'begin'` command from the primary server at any stage of your workload, letting the workload continue normally. Allow the `rs_ticket` command to run for about 60 minutes, then issue `rs_ticket 'end'`. The `rs_ticket` command flows through the Replication Server.

Check the `rs_ticket_history` table on the standby server:

- If both the tickets have reached the `rs_ticket_history` table, you can calculate the time taken to reach the end component by subtracting the `begin` value from the `end` value. In this example, the `begin` value is `08/29/15 15:08:53.243`, and the `end` value is `08/29/15 15:18:53.363`, so the time required was approximately 10 minutes:

```
V=2;H1=begin;PDB(HA1)=08/29/15 15:08:53.243;RA(HA1)=08/29/15
15:08:53.243;EXEC(51)=08/29/15
15:18:53.374;B(51)=34770028920;DIST(11)=08/29/15
15:18:53.623;DSI(75)=08/29/15
15:18:53.907;DSI_T=8143963;DSI_C=426597987;RRS=HA1_REP_hasite2
V=2;H1=end;PDB(HA1)=08/29/15 15:18:53.363;RA(HA1)=08/29/15
15:18:53.363;EXEC(52)=08/30/15
00:00:01.619;B(52)=34770028920;DIST(11)=08/30/15
00:00:01.872;DSI(75)=08/30/15
00:00:02.165;DSI_T=8143964;DSI_C=426597990;RRS=HA1_REP_hasite2
```

- If a component is taking a lot of time and not catching up with other components, identify the component causing the bottleneck by subtracting the `end` value from the `begin` value for each component, and then return the component.

Additionally, you can frequently check the Replication Server error log for a "memory limit exceed" message, indicating that Replication Server has reached its memory limit. If you see this error message, you may need to increase the value for the Replication Server `memory_limit` configuration parameter.

i Note

By default, Replication Server attempts to manage the memory as much as possible without any manual intervention.

11.9 Failover

`sap_status` shows the reason for failure and the corrective actions. View the log files of RMA, SAP ASE, and Replication Server to see additional details of the failures.

i Note

The Fault Manager triggers an unplanned failover by connecting to the RMA on the companion node and issuing `sap_failover`. Begin troubleshooting by first looking into the Fault Manager log and issuing `sap_status` on the RMA on the companion node. The steps to rectify an unplanned failover are identical to those of a planned failover.

RMA log file is located in `$$SYBASE/DM/RMA-16_0/instances/AgentContainer/logs` directory. The log files have the format `<RMA_MMDDYYY.log>`.

The RMA rotates the log file each day. Log files from earlier days are in the same directory.

The Replication Server log file is located in `$$SYBASE/DM/<CID>_REP_<logical_site>` and is named `<CID>_REP_<logical_site>.log`.

By default, the SAP ASE log file for the primary and standby server are located in `$$SYBASE/$$SYBASE_ASE/install/<server_name.log>`.

A planned failover executes these steps to switch the activity to the new primary site:

- Deactivating the primary SAP ASE and waiting for the backlog to be drained.
- Ensuring that all data has been applied to the standby site.
- Reconfiguring the system, changing the replication direction.
- Activating the new SAP ASE as the primary server.

`sap_failover` is an asynchronous operation. Use `sap_status 'task'` to check the status of the `sap_failover` command.

If `sap_failover` fails, use the `sap_status 'task'` parameter to see the reason for the failure and any verification you can perform for your system.

Example 1

In this example, a planned failover failed during the deactivation step because the transaction log of the source SAP ASE was not drained during the time provided by the `sap_failover` command:

```
[Status, Start Time, Fri Aug 28 05:41:43 EDT 2015]
[Status, Elapsed Time, 00:01:01]
[Failover, Task Name, Failover]
[Failover, Task State, Error]
[Failover, Short Description, Failover moves primary responsibility from current
logical source to logical target.]
[Failover, Long Description, Reverting ASE HADR back to active for the source
ASE 'PI2_PR:PR' because timeout occurred waiting for the source to be
deactivated.]
[Failover, Current Task Number, 4]
[Failover, Total Number of Tasks, 18]
[Failover, Task Start, Fri Aug 28 05:41:43 EDT 2015]
[Failover, Task End, Fri Aug 28 05:42:44 EDT 2015]
[Failover, Hostname, site0]
```

In this situation, check if all Replication Agents are running on the databases participating in the HADR system by executing `sp_help_rep_agent scan` to check the progress in the transaction log, or using the `sp_help_rep_agent scan_verbose` parameter to determine the number of pages pending for Replication Agent to process before reaching the end of the log. For example:

```
sp_help_rep_agent pdb, scan_verbose
go
Replication Agent Scan Status
dbname pathname scanner_spid start_marker end_marker current_marker
log_pages_left log_recs_scanned oldest_transaction
-----
pdb      n/a                60 (13399,39)  (13399,39)
(13399,39) 0                2837525 (0,0)
```

Example 2

In this example, the `sap_status 'task'` parameter shows that the validation that checks if all data was applied to the standby site has failed:

```
[Status, Start Time, Fri Aug 28 05:49:43 EDT 2015]
  [Status, Elapsed Time, 00:01:02]
  [Failover, Task Name, Failover]
  [Failover, Task State, Error]
  [Failover, Short Description, Failover moves primary responsibility
from current logical source to logical target.]
  [Failover, Long Description, Reverting ASE HADR back to active for the
source ASE 'PI2_PR:PR' because timeout occurred waiting for the end of data
marker.]
  [Failover, Failing Command Error Message, Timeout occured waiting for
the end of data marker for database 'PI2' to be received.]
  [Failover, Corrective Action, Log into the HA Replication Server and
verify DSI connection 'PI2_HA.PI2' has a state of 'Awaiting Message'.]
  [Failover, Current Task Number, 6]
  [Failover, Total Number of Tasks, 18]
  [Failover, Task Start, Fri Aug 28 05:49:43 EDT 2015]
  [Failover, Task End, Fri Aug 28 05:50:45 EDT 2015]
  [Failover, Hostname, site0]
```

In this situation, `sap_status 'task'` suggests checking whether all connections to Replication Server are running.

In case of failure, check the SAP ASE and Replication Server logs for more information.

Example 3

In this example, the `sap_status` command indicates that there is a Replication Agent instance enabled on the database in the standby host, which is not in a normal state:

```
TASKNAME    TYPE
VALUE
-----
-----
Status      Start Time          Thu Apr 04 22:58:59 PDT
2019
Status      Elapsed Time
00:00:03
Failover    Task Name
Failover    Task State
Error
Failover    Short Description   Failover makes the current standby ASE
as the primary server.
Failover    Long Description    Verifying logical host 'xxxxx' has been
made available.
Failover    Failing Command Error Message Logical host 'xxxxx' has not been made
available.Rep Agent is still enabled.
Failover    Additional Info 2   The primary Replication Server
'xxxxx:xxxx' is configured for synchronization mode and was found running in
synchronization mode.
Failover    Corrective Action   Run command 'sap_host_available
xxxxxxx' to make logical host available. Afterwards run command 'sap_failover
xxxxx, xxxxxx, 300, unplanned' again.
Failover    Task Start          Thu Apr 04 22:58:59 PDT
2019
Failover    Task End            Thu Apr 04 22:59:02 PDT
2019
Failover    Hostname            xxxxxxxxxxxxxxxx
```

In this situation, check the RMA log to see if the execution of the previous `sap_host_available` command succeeded. If the command failed, execute the command again or disable all the Replication Agent instances on the standby host, and then re-execute the failover command.

11.10 Access, and Login Redirection


By default, SAP ASE running in standby mode redirects the client login to the primary SAP ASE if the client login does not have the `allow_hadr_privilege` permission granted.

To check if the login has this privilege, connect to the primary SAP ASE and run:

```
sp_helpprotect @permission_name='allow_hadr_login'
```

The SAP ASE log contains information if it cannot redirect the login to the primary SAP ASE. In this situation, the login fails because it is an unprivileged client connection.

11.10.1 Troubleshooting the SAP ASE Cockpit

- Symptom – Password decrypt error in the agent log.
Action – Run the `passencrypt` script to set up the password. `passencrypt` is located in `com.sybase.ase/agent-plugin.xml`.
- Symptom – The phrase "no data" is displayed in the statistics chart of the [Overview](#) screen
Action – Create the technical user:
 1. Remove `tustore.xml` from the `SYBASE/plugins/ASEMAP` directory and restart the SAP ASE Cockpit server.
 2. Delete the `cockpit_user` or `tech_user`:
 1. Select the [Explore](#) tab.
 2. Select [Security](#) > [Logins](#) > [cockpit_user](#) > [Delete](#) .
 - If you see a message similar to "The login 'cockpit_user' was not dropped because it is currently active in process ID(s) 'xx'":
 1. Select the [Monitor](#) tab.
 2. Select [Processes](#).
 3. Select the process with the ID specified in the message.
 4. Right click, and select [Terminate Process](#).
 5. Click [OK](#).
 3. Create the Cockpit technical user:
 1. Select the [Explore](#) tab.
 2. Select [ASE Servers](#).
 3. Select the server.
 4. Click to select [Create Cockpit Technical User](#).You may safely ignore the error message posted to `ASEMAP.log` file, "Cockpit Technical User is not created."

i Note

As a last resort, you can stop the SAP ASE Cockpit and remove the repository's db directory. Follow the steps described above to create a new technical user.

- Symptom – The *HADR* menu is disabled in the *Monitoring* tab.
Action – Authenticate the RMA. However, you must create the technical user on the server that runs RMA and is configured for the HADR system.
- Symptom – *HADR* menu does not appear in Monitoring view.
Action – Enable the always-on option by setting:

```
sp_configure 'hadr mode', 1
```

Changing the Password for sccadmin After Installation

To change the sccadmin password after it has been installed:

1. Execute:

```
$$SYBASE/COCKPIT-4/bin/passencrypt -csi
```

2. Enter a new password at the prompt. passencrypt prints a ciphertext value.
3. Copy the entire ciphertext string, and paste the value into csi_config.xml for sccadmin:

```
<authenticationProvider controlFlag="sufficient"
name="com.sybase.security.core.PreConfiguredUserLoginModule">
<options name="username" value="sccadmin"/>
<options name="password" value="{SHA-256:QZT6pRIx6l8=}sENxwVBfp65aaTjd/
vAZFQAQOmkrk4Zl7iXLOFyJqCM="/>
<options name="roles" value="Cockpit Administrator"/>
</authenticationProvider>
```

4. Restart SAP ASE Cockpit to refresh the new password.

11.10.2 Troubleshooting Replication Agent

Replication Agent running on stream replication consists of three processes: the coordinator, the scanner, and the truncation point manager.

The Replication Agent is described in sp_who output:

0	16	background	0	NULL pdb	NULL tempdb	NULL REP AGENT
	0	NULL 17	background	0	NULL pdb	NULL tempdb
STPMGR						REP AGENT CI
	0	NULL 18	background	0	NULL	NULL

SCANNER	0	pdb	tempdb	REP AGENT CI
NULL		0		

This output displays the following:

- Coordinator – REP AGENT
- Truncation Point Manger – REP AGENT CI STPMGR
- Scanner – REP AGENT CI SCANNER

Sample Replication Agent Error Log

Replication Agent messages are indicated by RAT-CI, and client interface messages are indicated with CI-Info.

For example:

```
S(1) REPAGENT(4): [2014/03/27 04:10:39.05] RAT-CI (Coordinator): task is started.
00:0002:00000:00020:2014/03/27 04:10:39.05 kernel RAT-CI Coordinator-task is
spawned and running.
S(1) REPAGENT(4): [2014/03/27 04:10:39.05] RAT-CI (Coordinator): DBINFO rep
stream mode: 1
S(1) REPAGENT(4): [2014/03/27 04:10:39.05] RAT-CI (Coordinator): DBINFO rep
stream mode: 1
S(1) REPAGENT(4): [2014/03/27 04:10:39.05] RAT-CI (Coordinator): Stream mode
async was not changed to async.
S(1) REPAGENT(4): [2014/03/27 04:10:39.05] RAT-CI (Coordinator): RAT-CI is
started in ASYNC mode.
S(1) REPAGENT(4): [2014/03/27 04:10:39.05] RAT-CI (Coordinator): CI-library is
loaded, using CI-lib version : 0.0
S(1) REPAGENT(4): [2014/03/27 04:10:39.05] RAT-CI (STPMGR): STPMGR did not set
affinity to any engine
S(1) REPAGENT(4): [2014/03/27 04:10:39.05] RAT-CI (STPMGR): Request truncation
point and process results.
00:0002:00000:00021:2014/03/27 04:10:39.05 kernel RAT-CI STPMGR-task is spawned
and running.
S(1) REPAGENT(4): [2014/03/27 04:10:39.05] RAT-CI (Coordinator): 1 STPMGR
task(s) successfully bootstrapped.
S(1) REPAGENT(4): [2014/03/27 04:10:39.05] RAT-CI (Scanner): Scanner did not set
affinity to any engine.
00:0002:00000:00022:2014/03/27 04:10:39.05 kernel RAT-CI Scanner-task is
spawned and running.
S(1) REPAGENT(4): [2014/03/27 04:10:39.05] RAT-CI (Scanner): Interfaces file
used:
/marslinux4_eng1/elora/cepheus/interfaces
00:0002:00000:00022:2014/03/27 04:10:39.06 kernel (CI-Info) Message: 1001,
Severity: 0 (pdb)
Failed to allocate pkg buffer size 262144
00:0002:00000:00022:2014/03/27 04:10:39.06 kernel (CI-Info) Message: 1001,
Severity: 0 (pdb)
Failed to allocate package, ask 8, allocated 8, maxDObj 20 szPkgBuf 262144,
szXmitBuf 786432
No producer override set, using OCS configuration [ocs:RSHADR]
No consume override set, using OCS configuration [ocs:RSHADR]
S(1) REPAGENT(4): [2014/03/27 04:10:39.06] RAT-CI (Coordinator): 1 Scanner
task(s) successfully bootstrapped.
00:0002:00000:00020:2014/03/27 04:10:39.06 server Started Rep Agent on
database, 'pdb' (dbid = 4)
using streaming replication (sync mode) with library version 1.0.1
S(1) REPAGENT(4): [2014/03/27 04:10:41.10] RAT-CI (Scanner): RAT-CI connected to
RSHADR
```

```
S(1) REPAGENT(4): [2014/03/27 04:10:41.10] RAT-CI (Scanner): Maint user from
Replication Server
suid:3 uid:0 name:pdb_maint
S(1) REPAGENT(4): [2014/03/27 04:10:41.10] RAT-CI (Scanner): RAT-CI async
catchup: Number of pages to end of log: 7
S(1) REPAGENT(4): [2014/03/27 04:10:41.10] RAT-CI (Scanner): RAT-CI async
catchup: Number of pages to end of log: 0
S(1) REPAGENT(4): [2014/03/27 04:10:41.12] RAT-CI (Coordinator): DBINFO rep
stream mode: 4
S(1) REPAGENT(4): [2014/03/27 04:10:41.12] RAT-CI (Coordinator): Stream mode
changed from async to sync.
00:0002:00000:00020:2014/03/27 04:10:41.12 server Rep Agent on database 'pdb'
switched from mode 'async' to mode 'sync' because scanner reached end of log.
```

In the error log output, note that:

- The client interface version is indicated by:

```
using streaming replication (sync mode) with library version 1.0.1
```

- The maintenance user is indicated by:

```
RAT-CI (Scanner): Maint user from Replication Server suid:3 uid:0
name:pdb_maint
```

- Streaming mode is indicated by:

```
RAT-CI (Coordinator): RAT-CI is started in ASYNC mode.
```

RAT-CI however, changes from async to sync mode:

```
Rep Agent on database 'pdb' switched from mode 'async' to mode 'sync' because
scanner reached end of log.
```

- The memory is not properly configured:

```
00:0002:00000:00022:2014/03/27 04:10:39.06 kernel (CI-Info) Message: 1001,
Severity: 0 (pdb)
Failed to allocate pkg buffer size 262144
00:0002:00000:00022:2014/03/27 04:10:39.06 kernel (CI-Info) Message: 1001,
Severity: 0 (pdb)
Failed to allocate package, ask 8, allocated 8, maxDObj 20 szPkgBuf 262144,
szXmitBuf 786432
```

- Replication Agent has opened the stream to replicate and the client interface native thread has opened the CT-lib connection to Replication Server (that is, the channel is ready to replicate):

```
S(1) REPAGENT(4): [2014/03/27 04:10:41.10] RAT-CI (Scanner): RAT-CI connected
to RSHADR
```

- These messages show the number of pages before the client interface reaches the end of the log, at which point Replication Agent will switch to the configured mode:

```
S(1) REPAGENT(4): [2014/03/27 04:10:41.10] RAT-CI (Scanner): RAT-CI async
catchup: Number of pages to end of log: 7
S(1) REPAGENT(4): [2014/03/27 04:10:41.10] RAT-CI (Scanner): RAT-CI async
catchup: Number of pages to end of log: 0
```

Shutdown Messages

The Replication Agent issues messages similar to this during shutdown:

```
00:0002:00000:00034:2014/03/27 07:12:14.03 server Shutting down Rep Agent for
database, 'pdb' (dbid = 4).
The Replication Agent thread for database 'pdb' is being stopped.
(return status = 0)
1> 00:0002:00000:00034:2014/03/27 07:12:14.04 server Rep Agent on database 'pdb'
switched from mode 'sync' to mode 'async' because of a shutdown.
00:0000:00000:00000:2014/03/27 07:12:14.04 kernel (CI-Info) Message: 1001,
Severity: 0
(pdb)Producer sender shutdown is requested, exiting...
00:0002:00000:00034:2014/03/27 07:12:14.03 server Shutting down Rep Agent for
database, 'pdb' (dbid = 4).
S(1) REPAGENT(4): [2014/03/27 07:12:14.04] RAT-CI (Coordinator): DBINFO rep
stream mode: 1
S(1) REPAGENT(4): [2014/03/27 07:12:14.04] RAT-CI (Coordinator): Stream mode
changed from sync to async .
00:0002:00000:00034:2014/03/27 07:12:14.04 server Rep Agent on database 'pdb'
switched from mode 'sync '
to mode 'async' because of a shutdown.
00:0000:00000:00000:2014/03/27 07:12:14.04 kernel (CI-Info) Message: 1001,
Severity: 0 (pdb)Producer
sender shutdown is requested, exiting...
S(1) REPAGENT(4): [2014/03/27 07:12:16.04] RAT-CI (Scanner): Closed CI Stream...
S(1) REPAGENT(4): [2014/03/27 07:12:16.04] RAT-CI (Scanner): Closed CI Provided
Interface...
S(1) REPAGENT(4): [2014/03/27 07:12:16.04] RAT-CI (Coordinator): 1 Scanner
task(s) successfully stopped.
S(1) REPAGENT(4): [2014/03/27 07:12:16.04] RAT-CI (Coordinator): 1 STPMGR
task(s) successfully stopped.
S(1) REPAGENT(4): [2014/03/27 07:12:16.04] RAT-CI (Coordinator): task is stopped.
```

The Coordinator drives the shutdown process: the scanner closes the stream, stops the STPMGR, then stops the Coordinator.

During shutdown, Replication Agent switches to asynchronous mode to communicate to Replication Server.

Messages for Stream Replication

During stream replication, packages are composed of several commands. Every package may contain metadata (schema) and commands. Enable trace flag 9229 to see what Replication Agent is sending. For example:

```
S(1) REPAGENT(4): [2014/03/27 09:17:18.92] RAT-CI (Scanner): CMD (Header) :
BEGINXACT (CMD type 0).
S(1) REPAGENT(4): [2014/03/27 09:17:18.92] RAT-CI (Scanner): CMD (Header) : XACT
TranId:
0x00000000000001f5e00074841445231706462.
S(1) REPAGENT(4): [2014/03/27 09:17:18.92] RAT-CI (Scanner): CMD (Header) : XACT
OQID:
0x00000000000001f5e0000653600070000653600070000a2fb0099115700000000.
S(1) REPAGENT(4): [2014/03/27 09:17:18.92] RAT-CI (Scanner): CMD (Content):
BeginTime:
179198930522455, TranName: _ins UserName: sa Spid:23, Attributes:2.
S(1) REPAGENT(4): [2014/03/27 09:17:18.92] RAT-CI (Scanner): Metadata is added
for (1008003591) t1.
S(1) REPAGENT(4): [2014/03/27 09:17:18.92] RAT-CI (Scanner): Metadata
```

```

(TblSchema) name:t1 owner:dbo objid:
1008003591 version:1 attrib:0 numcols:2 numpkcols:0.
S(1) REPAGENT(4): [2014/03/27 09:17:18.92] RAT-CI (Scanner): Metadata (ColDesc)
Catalog colid:
1 Cached colid:1 citype:8 attrib:0 maxlen:4 name:i prec:0 scale:0
compresslevel:0.
S(1) REPAGENT(4): [2014/03/27 09:17:18.92] RAT-CI (Scanner): Metadata (ColDesc)
Catalog colid:
2 Cached colid:2 citype:0 attrib:0 maxlen:3 name:c prec:0 scale:0
compresslevel:0.
S(1) REPAGENT(4): [2014/03/27 09:17:18.92] RAT-CI (Scanner): CMD (Header) :
INSERT (CMD type 2).
S(1) REPAGENT(4): [2014/03/27 09:17:18.92] RAT-CI (Scanner): CMD (Header) : XACT
TranId:
0x000000000000001f5e00074841445231706462.
S(1) REPAGENT(4): [2014/03/27 09:17:18.92] RAT-CI (Scanner): CMD (Header) : XACT
OQID:
0x000000000000001f5e0000653600080000653600070000a2fb0099115700000000.
S(1) REPAGENT(4): [2014/03/27 09:17:18.92] RAT-CI (Scanner): CMD (Content): Row:
objid:1008003591 length:9 numcols:2.
S(1) REPAGENT(4): [2014/03/27 09:17:18.92] RAT-CI (Scanner): CMD (Content): Col
(after):
Catalog id:1 Cached id:1 citype:8 colval:0x156312ad2 len:4 prec:0 scale:0
pkey:0 irl:0.
S(1) REPAGENT(4): [2014/03/27 09:17:18.92] RAT-CI (Scanner): CMD (Content): Col
(after):
Catalog id:2 Cached id:2 citype:0 colval:0x156312ad6 len:3 prec:0 scale:0
pkey:0 irl:0.
S(1) REPAGENT(4): [2014/03/27 09:17:18.92] RAT-CI (Scanner): CMD (Header) :
ENDXACT (CMD type 1).
S(1) REPAGENT(4): [2014/03/27 09:17:18.92] RAT-CI (Scanner): CMD (Header) :
XACT TranId:0x000000000000001f5e00074841445231706462.
S(1) REPAGENT(4): [2014/03/27 09:17:18.92] RAT-CI (Scanner): CMD (Header) :
XACT OQID:0x000000000000001f5e0000653600090000653600070000a2fb0099115700000000.
S(1) REPAGENT(4): [2014/03/27 09:17:18.92] RAT-CI (Scanner): CMD (Content):
EndTime: 179198930522455.

```

The rows that contain the string "Metadata" (in bold) indicate that metadata is being added to the package. For example, execute this:

```

1> begin tran
2> insert t1 values (3, 'ccc')
3> insert t2 values (1, 'hola')
4> delete t1 where i = 1
5> go

```

The error log then includes messages similar to:

```

S(1) REPAGENT(4): [2014/03/27 09:25:56.10] RAT-CI (Scanner): Metadata is added
for (1040003705) t2.
S(1) REPAGENT(4): [2014/03/27 09:25:56.10] RAT-CI (Scanner): Metadata
(TblSchema) name:t2 owner:dbo
objid:1040003705 version:1 attrib:0 numcols:2 numpkcols:0.

```

There is no need to include the schema for object t1 in this package because this schema was already sent in a previous package.

Secondary Truncation Point Management

Replication Agent generates a unique identifier for every log record to be replicated. This identifier is internally named OQID (origin queue ID). This value is used internally to indicate whether a message was already sent. OQIDs look similar to this in the log files:

```
S(1) REPAGENT(4): [2014/03/27 09:58:23.88] RAT-CI (Scanner): CMD (Header) : XACT
OQID:0x000000000000020120000654500020000654500020000a2fb00a45a010000000.
```

11.11 Troubleshooting the Fault Manager

Problems in the Fault Manager typically occur when it takes actions on participating nodes.

Since the Host Agent is responsible for executing all local actions, it is useful to understand how to troubleshoot it to resolve issues for the Fault Manager.

Root Partition Is Full

On hosts running the primary or standby servers, the Fault Manager heartbeat log file, named `dev_hbeat`, may grow very large (may be gigabytes in size), causing the host's root (`/`) partition to fill, and the `asehostctrl` to fail. Check the size of the file with this command to determine if a failure is caused by `dev_hbeat` growing too large:

```
sudo du -sh /usr/sap/hostctrl/work/dev_hbeat
16G /usr/sap/hostctrl/work/dev_hbeat
```

To resolve this, delete the `dev_hbeat` file.

Increase the Trace Level for Troubleshooting

You should set the trace to its highest level on the SAP Host Agent and the Fault Manager so the error log output is as verbose as possible.

- For the SAP Host Agent – Increasing the trace level on the SAP Host Agent requires you to set the trace level in the profile file and restart the SAP Host Agent for the change to take effect. The profile file is located in:
 - (UNIX) `/usr/sap/hostctrl/exe/host_profile`
 - (Windows) `%ProgramFiles%\SAP\hostctrl\exe\host_profile`
- 1. Add this line to the profile file:

```
service/trace = 3
```

- Restart the SAP Host Agent:

```
/usr/sap/hostctrl/exe/saphostexec -restart
```

- For the Fault Manager – Increasing the trace level on the Fault Manager requires you to set the trace level in the profile file and restart the Fault Manager for the change to take effect. Increasing the trace level increases the number of log entries, and may increase the file size. The profile file is named `SYBHA.PFL`, and is located in the installation directory of the Fault Manager on all platforms:

- Add this line to the profile file:

```
ha/syb/trace = 3
```

- Restart the Fault Manager:

```
<Fault_Manager_install_dir>/FaultManager/sydbfm_<CID>
```

The Fault Manager makes these calls through the SAP Host Agent:

- `LiveDataBaseUpdate` – used to take an action (for example, restarting components, initiating failovers, and so on). `LiveDataBaseUpdate` takes a `<TASK>` argument with it, which defines the action to perform. `<TASK>` is one of:
 - `HEARTBEAT_STARTUP`
 - `HEARTBEAT_CHECK`
 - `HEARTBEAT_STOP`
 - `REPLICATION_STATUS`
 - `RS_STANDBY_AVAILABLE`
 - `RESUME_REPLICATION`
 - `HA_VERSION HB_VERSION`
 - `ADD_ASE_INSTANCE`
 - `SET_USER_PASSWORD`
 - `FAILOVER_UNPLANNED`
- `GetDatabaseStatus` – Used to view the status of components. However, no action is taken on any component.
- `StartDatabase`
- `StopDatabase`

Hostagent Timeout

If the following messages are displayed frequently in the SAP ASE Cockpit, they indicate that `sapdbctrl` calls from the Fault Manager are timing out; you may need to increase the configured timeout value for `SAP dbctrl`:

```
Primary SAP Host ({Primary_Site}) Agent cannot be contacted Primary SAP Host
({Primary_Site}) Agent
contact is restored
Secondary SAP Host ({Secondary_Site}) Agent cannot be contacted Secondary SAP
Host ({Secondary_Site}) Agent
contact is restored
```

To resolve this, increase the timeout period for `sapdbctrl` by increasing the values for the `ha/syb/dbctrl_timeout` parameter in the Fault Manager profile file. The default value is 30.

Restart the Fault Manager after updating profile file.

SAP Host Agent debugging

Check the SAP Host Agent log file on the respective host if Fault Manager calls to the SAP Host Control fail. The SAP Host Control log file is located in the `/usr/sap/hostctrl/work/dev_sapdbctrl` file. See the *SAP Host Agent Troubleshooting Guide* at: <http://scn.sap.com/docs/DOC-34217> for more information.

Stopping the Fault Manager

- Symptom – While using `stop` command to shut down the Fault Manager, you see this message:

```
fault manager did not change to mode UNKNOWN within 60 seconds.  
fault manager running, pid = 15922, fault manager overall status = OK,  
currently executing in mode DIAGNOSE
```

This message occurs because the Fault Manager exits and displays the currently running mode if it is unable to stop.

- Action – Re-execute the `stop` command.

⚠ Caution

Never stop the Fault Manager using the `kill -9` operating system command.

Troubleshooting 'replication status' Messages

- Symptom – The status of the primary and companion HADR nodes is healthy, but the `sanity report` displays the 'replication status' as one of following:
 - Suspended
 - Dead
 - Unusable
 - Indoubt
 - Unknown
- Action – Consult the Replication Server error logs for information.

sybdbfm "No Fault Manager Found" Message

- Symptom – You see a message similar to this when running any Fault Manager `sybdbfm` command:

```
no fault manager found for current working directory
```

```
error: stop failed.
```

You are not running the `sybdbfm` command from the directory where the `profile` file and other Fault Manager-generated files (such as `sp_sybdbfm`, `stat_sybdbfm`, and so on) are located.

- Action – Rerun the `sybdbfm` command from the directory where these files are located.

sybdbfm status Message "Fault Manager Is not Responding"

- Symptom – You see a message similar to this when issuing `sybdbfm status`:

```
fault manager is not responding, currently busy. Please retry.
```

The Fault Manager is currently busy.

- Action – Re-run the `sybdbfm status` command.

Troubleshooting dev_sybdbfm Messages

- Symptom – `dev_sybdbfm` displays a message similar to:

```
2015 01/30 06:43:41.488 connected using using DRIVER={COCKPIT
4};server=mo-4a63cdeba;port=4998;uid=sccadmin;pwd=xxxxxxxx;logintimeout=3;Com
mandTimeOut=3;BackEndType=ReplicationServer.
2015 01/30 06:43:49.495 Sending notification to SAP ASE Cockpit failed with
error: -1.
2015 01/30 06:43:49.495 SQLGetDiagRec 0
2015 01/30 06:43:49.495 ERROR in function NotifySCC (1166) (SQLExecute
failed): (30046) [08S01] [SAP][ASE ODBC Driver]Connection to the server has
been lost. Unresponsive Connection was disconnected during command timeout.
Check the server to determine the status of any open transactions.
2015 01/30 06:43:49.495 ERROR in function NotifySCC (1166) (SQLExecute
failed): (30149) [HYT00] [SAP][ASE ODBC Driver]The command has timed out.
2015 01/30 06:43:49.495 notification of standby SAP ASE Cockpit failed.
2015 01/30 06:43:49.495 disconnect connection
```

Action – This is not a fatal problem. This message is displayed when the Fault Manager cannot communicate with the SAP ASE Cockpit on its configured TDS ports (the default port number is 4998). This message is typically displayed when the wrong TDS port number is entered, or when the SAP ASE Cockpit is not running. The Fault Manager continues to run, but does not send out notifications. If the SAP ASE Cockpit is not running, restart it with:

```
$$SYBASE/COCKPIT-4/bin/cockpit.sh -stop
$$SYBASE/COCKPIT-4/bin/cockpit.sh -start
```

- Symptom – `dev_sybdbfm` displays a message similar to:

```
2014 11/18 21:18:01.116 executing: asehostctrl -host <ASE host> -user sapadm
*****
-function LiveDatabaseUpdate -dbname <ASE Server Name> -dbtype syb -
dbinstance <Site
Name for ASE in RMA> -updatemethod Check -updateoption TASK=HB_VERSION .
```

This message displays when:

- The Fault Manager fails to start during the bootstrap cycle because of mismatched heartbeat versions
- There are older Fault Manager heartbeat clients that are already running on any of the participating HADR nodes

Action – Kill the heartbeat clients that are running with the `kill -9` command. Check `/usr/sap/hostctrl/work/dev_sapdbctrl` and `/usr/sap/hostctrl/work/dev_saphostexec` for more information (these files require sudo access on the host).

- Symptom – If this command fails in `dev_sybdbfm`:

```
2015 03/04 10:26:55.089 executing: asehostctrl -host <ASE Host> -user sapadm
*****
-function LiveDatabaseUpdate -dbname <DB Name> -dbtype syb -dbinstance <ASE
Sitename>
-updatemethod Execute -updateoption TASK=HEARTBEAT_STARTUP .
```

This message displays when:

- The Fault Manager fails to start heartbeat client on one of the HADR nodes during the bootstrap cycle
- There is already a process listening on the set or default port for heartbeat clients

Action – Check `/usr/sap/hostctrl/work/dev_sapdbctrl`, `/usr/sap/hostctrl/work/dev_saphostexec`, and the heartbeat log file on the participating node (`/usr/sap/hostctrl/work/dev_hbeat`), and other log files, for more information.

- Symptom – `dev_sybdbfm` displays a message similar to:

```
2015 03/04 10:26:41.952 executing: asehostctrl -host <ASE Host> -user sapadm
*****
-function LiveDatabaseUpdate -dbname <DB Name> -dbtype syb -dbinstance <ASE
Sitename>
-updatemethod Execute -updateoption TASK=HEARTBEAT_STOP .
```

This message typically occurs when the Fault Manager fails during the bootstrap cycle.

Action – Check `/usr/sap/hostctrl/work/dev_sapdbctrl` and `/usr/sap/hostctrl/work/dev_saphostexec` for more information (these files require sudo access on the host).

- Symptom – `dev_sybdbfm` displays a message similar to:

```
2015 03/04 10:31:32.465 executing: asehostctrl -host <ASE Host> -user sapadm
*****
-function LiveDatabaseUpdate -dbname <DB Name> -dbtype syb -dbinstance <ASE
Sitename>
-updatemethod Check -updateoption TASK=REPLICATION_STATUS .
```

This message typically occurs when the Fault Manager fails during the bootstrap cycle.

Action – Check `/usr/sap/hostctrl/work/dev_sapdbctrl` and `/usr/sap/hostctrl/work/dev_saphostexec` for more information (these files require sudo access on the host).

- Symptom – `dev_sybdbfm` displays a message similar to:

```
2015 01/08 03:44:13.423 executing: asehostctrl -host <ASE Host> -user sapadm
*****
-function LiveDatabaseUpdate -dbname <DB Name> -dbtype syb -dbinstance <ASE
Sitename>
-updatemethod Execute -updateoption TASK=RS_STANDBY_AVAILABLE.
```

This message typically occurs when the Fault Manager fails during the bootstrap cycle.

Action – Check `/usr/sap/hostctrl/work/dev_sapdbctrl` and `/usr/sap/hostctrl/work/dev_saphostexec` for more information (these files require sudo access on the host).

- Symptom – dev_sybdbfm displays a message similar to:

```
2015 03/04 10:35:31.598 executing: asehostctrl -host <ASE Host> -user sapadm
*****
-function LiveDatabaseUpdate -dbname <DB Name> -dbtype syb -dbinstance <ASE
Sitename>
-updatemethod Execute -updateoption TASK=RESUME_REPLICATION.
```

This message typically occurs when the Fault Manager fails during the bootstrap cycle.

Action – Check /usr/sap/hostctrl/work/dev_sapdbctrl and /usr/sap/hostctrl/work/dev_saphostexec for more information (these files require sudo access on the host).

- Symptom – dev_sybdbfm displays a message similar to:

```
2015 03/04 10:25:33.983 executing: asehostctrl -host <ASE Host> -user sapadm
*****
-function LiveDatabaseUpdate -dbname <DB Name> -dbtype syb -dbinstance <ASE
Sitename>
-updatemethod Execute -updateoption TASK=ADD_ASE_INSTANCE -updateoption
SYBASE_PATH=<ASE Installation Path>/DM -updateoption SYBASE_USER=<OS User
designated for use of ASE>
```

This message typically occurs when:

- The Fault Manager fails to add the SAP ASE instance
- One of the SAP Host Agent services is not running

Action:

- Verify that these services are running:

```
/usr/sap/hostctrl/exe/saphostexec
/usr/sap/hostctrl/exe/sapstartsrv
/usr/sap/hostctrl/exe/saposcol
```

- Check and restart all services:

```
/usr/sap/hostctrl/exe/saphostexec -restart
```

- Check /usr/sap/hostctrl/work/dev_sapdbctrl and /usr/sap/hostctrl/work/dev_saphostexec for more information (these files require sudo access on the host).

- Symptom – dev_sybdbfm displays a message similar to:

```
2015 03/04 10:25:38.700 executing: asehostctrl -host <ASE Host> -user sapadm
*****
-function LiveDatabaseUpdate -dbname <DB Name> -dbtype syb -dbinstance <ASE
Sitename> -updatemethod Check
-updateoption TASK=HA_VERSION
```

This message typically occurs when

- The Fault Manager fails to match the HA version during the bootstrap cycle
- You are using an older SAP Host Agent version with a newer Fault Manager

Action – Kill the heartbeat clients that are running with the kill -9 command. Check /usr/sap/hostctrl/work/dev_sapdbctrl and /usr/sap/hostctrl/work/dev_saphostexec for more information (these files require sudo access on the host).

- Symptom – dev_sybdbfm displays a message similar to:

```
2015 03/04 10:25:32.814 executing: asehostctrl -host <ASE Host> -user sapadm
*****
-function LiveDatabaseUpdate -dbname <DB Name> -dbtype syb -dbinstance <ASE
Sitename>
```

```
-dbuser DR_admin -dbpass ***** -updateoption Execute -updateoption  
TASK=SET_USER_PASSWORD  
-updateoption USER=DR_ADMIN
```

This message occurs when the Fault Manager fails to set a username and password for a particular component during the bootstrap cycle, typically when there is a password mismatch from the ones entered into SecureStore when installing Fault Manager and the current passwords for the same user name. Action – Reconfigure the Fault Manager to automatically add the latest user and password combinations, or update the individual passwords using the `rsecssfx` binary (located in `<installation_directory>/FaultManager/bin/`). Check `/usr/sap/hostctrl/work/dev_sapdbctrl` and `/usr/sap/hostctrl/work/dev_saphostexec` for more information (these files require sudo access on the host).

- Symptom – `dev_sybdbfm` displays a message similar to:

```
2015 03/04 10:33:12.892 executing: asehostctrl -host <ASE Host> -user sapadm  
*****  
-function LiveDatabaseUpdate -dbname <DB Name> -dbtype syb -dbinstance <ASE  
Sitename>  
-updateoption Execute -updateoption TASK=FAILOVER_UNPLANNED -updateoption  
PRIMARY_SITE=meereen  
-updateoption STANDBY_SITE=<ASE Sitename>
```

This message occurs during execution when the Fault Manager fails to fail over, which may be caused by more than one component in the HADR setup failing, or by a network outage during failover. Action – Check the health of all components listed in the last Fault Manager health report, or query RMA with the `sap_status path` command. A report that describes more than a single component as unavailable probably explains the failure of the failover. Restart any required components. If failover fails due to an intermittent network outage, it will attempt the failover again when it is back online. Check `/usr/sap/hostctrl/work/dev_sapdbctrl` and `/usr/sap/hostctrl/work/dev_saphostexec` for more information (these files require sudo access on the host).

Troubleshooting Commands Outside the Fault Manager

For errors that occur in the Fault Manager, you should reproduce the errors in a smaller system to determine the root cause by scanning the error log for the failed command and replacing the `asehostctrl` command with `/usr/sap/hostctrl/exe/saphostctrl`.

Executing these commands individually at the command line gives you greater control over reproducing the error and accessing the root cause. Check `/usr/sap/hostctrl/work/dev_sapdbctrl` and `/usr/sap/hostctrl/work/dev_saphostexec` for more details on what happens when you execute these commands (these files require sudo access on the host).

Viewing Diagnostic Information in Log Files

To view diagnostic information from these log files, set the trace level to 3:

- `/usr/sap/hostctrl/exe/dev_saphostctrl`
- `/usr/sap/hostctrl/exe/dev_sapdbctrl`

Check the RMA log files in `$$SYBASE/DM/RMA-16_0/instances/AgentContainer/logs/RMA_*.log` for replication issues. Check the SAP ASE log files in `$$SYBASE/ASE-16_0/install/*.log` for server issues.

Avoiding Common Errors

There are a number of issues that commonly cause errors in the HADR system.

- Directories do not have appropriate permissions – The installation directories for SAP ASE, Fault Manager, execution directory, and `/tmp` require the appropriate permissions. The Fault Manager creates temporary directories under `/tmp`, to which it adds temporary files. If permissions prevent it from doing this, the SAP Host Agent call will fail but it will not know the reason. Verify that the user executing the Fault Manager has permissions on all the required directories.
- Verify that `/tmp` is not full – If `/tmp` is full, Fault Manager cannot create temporary files. Check the status of `/tmp` by executing `df -k /tmp`. If this command returns a value of 100% for `Use%`, you may need to make additional room in `/tmp`.
- Verify that GLIBC version is 2.7 or later – The Fault Manager is built with `glibc` version 2.7, so the hosts running it must use `glibc` version 2.7 or later. Use this command to check the `glibc` version:

```
ldd --version
```

- Enter the correct passwords for `sa`, `DR_admin` and `sapadm` – It is very difficult to find the root cause of errors when password mismatches are the culprit, and it is best to verify the passwords are correct before the errors occur. The SAP Host Agent installation may not include a default password for `sapadm`, but is required by Fault Manager. Add or update a password with the `passwd` command.

Errors While Installing the Fault Manager

An error that reads "Error: Invalid Credentials" during the Fault Manager installation indicates that one of the components for which you entered information is offline, or that the username and password combination for that component is incorrect.

Consequently, the Fault Manager cannot connect to that component to verify that it is running. This connection fails only when the username and password combination are incorrect or the component is down.

The Fault Manager Cannot Connect to SAP ASE After Restarting the Host

When you restart the host, the SAP Host Agent and its services are also restarted according to how their entries are added to `init.d` during installation.

However, installation changes the environment and starts the SAP Host Agent with the `LANG` environment variable set to `en_US.UTF-8`.

The host may have been changed prior to a restart, and the SAP Host Agent may be started in an environment that is different from what is specified in `init.d`. After you restart the host, you should source `SYBASE/SYBASE.sh` and restart the SAP Host Agent with this command:

```
/usr/sap/hostctrl/exe/saphostexec -restart
```

Sample Error Log and Corrective Action

This Fault Manager error log indicates that the Fault Manager could not create a connection to the Host Agent.

```
2014 09/12 04:06:42.996 executing: asehostctrl -host star4 -user sapadm
***** -function LiveDatabaseUpdate -dbname JD3 -dbtype syb -dbinstance
STAR4 -updatemethod Execute -updateoption TASK=ADD_ASE_INSTANCE -updateoption
SYBASE_PATH=/star4_eng/jdabhole/asel60
-updateoption SYBASE_USER=jdabhole -
updateoption DRA_HOST=star4 -updateoption DRA_PORT=7001 .
2014 09/12 04:06:42.996 starting control call.
NiInit3: NI already initializes (init=6;cur=2048)
NiHLGetNodeAddr: found hostname 'star4' in cache
NiIGetNodeAddr: hostname 'star4' = addr 10.172.162.61
NiIGetServNo: servicename '1128' = port 1128
NiICreateHandle: hdl 6 state NI_INITIAL_CON
NiIInitSocket: set default settings for new hdl 6/sock 6 (UD; ST)
NiIBlockMode: set blockmode for hdl 6 FALSE
NiIConnectSocket: hdl 6 is connecting to /tmp/.sapstream1128 (timeout=-1)
NiHLGetHostName: found address 10.172.162.61 in cache
NiIGetHostName: addr 10.172.162.61 = hostname 'star4'
NiPConnect: SiConnect failed for hdl 6/sock 6
(SI_ECONN_REFUSE/13; UD; ST; /tmp/.sapstream1128)
NiIInitSocket: set default settings for new hdl 6/sock 6 (I4; ST)
NiIBlockMode: set blockmode for hdl 6 FALSE
NiIConnectSocket: hdl 6 is connecting to 10.172.162.61:1128 (timeout=-1)
SiPeekPendConn: connection of sock 6 broken (111-Connection refused)
NiHLGetHostName: found address 10.172.162.61 in cache
NiIGetHostName: addr 10.172.162.61 = hostname 'star4'
***LOG Q0I=> NiPConnect2: 10.172.162.61:1128: connect (111: Connection refused)
[/bas/CGK_MAKE/src/base/ni/nixxi.cpp 3324]
*** ERROR => NiPConnect2: SiPeekPendConn failed for hdl 6/sock 6
(SI_ECONN_REFUSE/111; I4; ST; 10.172.162.61:1128) [nixxi.cpp 3324]
NiICloseHandle: closing initial hdl 6
2014 09/12 04:06:42.998 dbctrl call cnt reset: 0 .
2014 09/12 04:06:42.998 control call ended.
2014 09/12 04:06:42.999 call_saphostctrl completed ok.
2014 09/12 04:06:42.999 saphostctrl executed.
2014 09/12 04:06:42.999 dbctrl call cnt reset 2: 0 .
2014 09/12 04:06:42.999 ase instance status is: .
2014 09/12 04:06:42.999 add instance status: failure
```

The section of the error log that describes the error is:

```
***LOG Q0I=> NiPConnect2: 10.172.162.61:1128: connect (111: Connection refused)
[/bas/CGK_MAKE/src/base/ni/nixxi.cpp 3324]
*** ERROR => NiPConnect2: SiPeekPendConn failed for hdl 6/sock 6
(SI_ECONN_REFUSE/111; I4; ST; 10.172.162.61:1128) [nixxi.cpp 3324]
```

This error occurs when `sapstartsrv` is not running. To solve the problem, issue this command to check if the `sapstartsrv` process is running:

```
ps -aef | grep sapstartsrv
```

11.11.1 Fault Manager and SAP Host Agent Commands

You can start and run the Fault Manager and the SAP Host Agent from the command line.

See [The SAP Host Agent \[page 28\]](#) for information about the SAP Host Agent commands. See [Configuring the Fault Manager from the Command Line \[page 147\]](#) for information about the Fault Manager commands.

11.12 Configuring the RMI Port

RMA Remote Method Invocation (RMI) is an API that RMA instances use in the HADR environment to invoke methods from other RMA instances remotely. To support remote invocation, RMA RMI needs five ports: the configured port it occupies, as well as the previous four consecutive ports.

Configuring a port number for RMA RMI in the `setuphadr` response file automatically reserves the four consecutive integers before that configured number. For example, when the RMA RMI port number is configured with 7000, 6999, 6998, 6997, and 6996 are also occupied by RMA RMI.

To avoid unexpected errors, make sure the configured port and its four predecessors are neither blocked by firewalls nor used by any other applications.

If you have run into errors with your environment setup when using the `setuphadr` utility, troubleshoot the error on each site using the following steps:

1. Obtain the RMA process ID:

```
ps -ef|grep rma
```

2. Kill this RMA process ID:

```
kill <process id>
```

3. Delete the configuration information for the RMA instance:

```
rm -rf $SYBASE/DM/RMA-15_5/instances/AgentContainer/configdb
```

4. Find five consecutive port numbers that are all available, then modify the `setuphadr` response file for all servers by changing `rma_rmi_port` to the highest-numbered port in your set of five consecutive ports. For example:

```
From:
PRIM.rma_rmi_port=7000
To:
PRIM.rma_rmi_port=8000
From:
COMP.rma_rmi_port=7000
To:
COMP.rma_rmi_port=8000
```

5. Modify `$SYBASE/DM/RMA-15_5/instances/AgentContainer/config/bootstrap.prop` by changing `rmiport` to the port number as configured in last step. For example:

```
From
rsge.bootstrap.rmiport=7000
To
```

```
rsgc.bootstrap.rmiport=8000
```

6. Rerun the `setuphadr` utility on each site.

12 HADR Reference

Use commands, system procedures, and proxy tables to administer the HADR system.

12.1 RMA Commands

Use the Replication Management Agent (RMA) commands to administer, monitor, or modify the properties of your replication environment.

12.1.1 sap_add_device

Use the `sap_add_device` command to add SAP Replication Server device spaces. The `sap_add_device` command issues an add partition command to the underlying SAP Replication Server that is defined for the requested logical host.

For more details, see *Replication Server Reference Manual > Replication Server Commands > create partition*.

Syntax

```
sap_add_device <logical_host_name>, <logical_device_name>, <device_file>, <size>
```

Parameters

<logical_host_name>

Specifies the logical host name of the SAP Replication Server on which the device space is to be added.

<logical_device_name>

Specifies the logical name of the device to be added.

<device_file>

Specifies the path to the directory in which you are creating the device and the name of the device file.

<size>

Specifies the size of the device file (MB).

Examples

Example 1

Add 20 MB device space named `part9` on `site0` SAP Replication Server under `/testenv7/partition9.dat`:

```
sap_add_device site0, part9, /testenv7/partition9.dat, 20
go
```

The returned result is:

TASKNAME	TYPE	VALUE
Add Replication Server Device	Start Time	Thu May 16 05:11:42 UTC 2019
Add Replication Server Device	Elapsed Time	00:00:01
AddDevice	Task Name	Add Replication Server Device
AddDevice	Task State	Completed
AddDevice	Short Description	Add Replication Server Device with the specified size.
AddDevice	Long Description	Successfully added device 'part9' (file '/testenv7/partition9.dat') with size 20 on site0 host 'rmawqiusitel.mo.sap.corp:14975'.
AddDevice	Task Start	Thu May 16 05:11:42 UTC 2019
AddDevice	Task End	Thu May 16 05:11:43 UTC 2019
AddDevice	Hostname	rmawqiusitel.mo.sap.corp
(9 rows affected)		

Example 2

If the command fails, find the error from the row where the `TYPE` value is `Failing Command Error` Message:

```
sap_add_device site0, part3, /testenv7/partition3.dat, 9
go
```

The returned result is:

TASKNAME	TYPE	VALUE
Add Replication Server Device	Start Time	Thu May 16 02:22:27 UTC 2019
Add Replication Server Device	Elapsed Time	00:00:00
AddDevice	Task Name	Add Replication Server Device
AddDevice	Task State	Error
AddDevice	Short Description	Add Replication Server Device with the specified size.
AddDevice	Long Description	Adding device 'part3' (file '/testenv7/partition3.dat') with size 9 on site0 host 'rmawqiusitel.mo.sap.corp:14975'.
AddDevice	Failing Command Error Message	Device logical name 'part3' already exists on site0 host 'rmawqiusitel.mo.sap.corp:14975'.
AddDevice	Task Start	Thu May 16 02:22:27 UTC 2019
AddDevice	Task End	Thu May 16 02:22:27 UTC 2019
AddDevice	Hostname	rmawqiusitel.mo.sap.corp

```
(10 rows affected)
```

Related Information

[sap_drop_device](#) [page 462]

[sap_list_device](#) [page 478]

12.1.2 sap_cancel

Use the `sap_cancel` command to cancel any asynchronous or background tasks.

Syntax

```
sap_cancel  
go
```

12.1.3 sap_collect_log

Use the `sap_collect_log` command to collect the ASE/RS/RMA logs of different hosts in a HADR system to one local directory.

Syntax

```
sap_collect_log [, <logical_host_name>] [, <number_of_days>]
```

Parameters

<logical_host_name>

Specifies the logical host name to collect the logs from.

<number_of_days>

Specifies the number of days for which the RMA maintains the logs. The valid value is between 1-20. The default value is 5.

Examples

Example 1

Collects logs from the site PR for the last 10 days:

```
Sap_collect_log PR 10
```

The returned result is:

TASKNAME	TYPE	VALUE
Collect Log	Start Time	Thu Feb 09 04:32:44 EST 2017
Collect Log	Elapsed Time	00:00:00
CollectLog	Task Name	Collect Log
CollectLog	Task State	Completed
CollectLog	Short Description	Collect ASE/RMA/RS logs at HADR hosts and put them to one directory.
CollectLog	Long Description	All the log files are collected at /opt/sap/DM/collected_logs.
CollectLog	Task Start	Thu Feb 09 04:32:44 EST 2017
CollectLog	Task End	Thu Feb 09 04:32:44 EST 2017
CollectLog	Hostname	site0

Example 2

Collects logs from the PR site for the last 5 (default) days.

```
Sap_collect_log PR
```

Example 3

Collects logs from all the logical hosts in HADR for the last 10 days.

```
Sap_collect_log 10
```

Example 4

Collects logs from all the logical hosts in HADR for the last 5 days.

```
Sap_collect_log
```

Usage

After you execute the `sap_collect_log` command, the result shows the directory path where the logs are saved. The name of the directory is same as the name of the logical host name.

The different types of log files are:

- ASE error log
- RMA log
- Replication Server error log
- Replication Server initialize log
- Replication Server standard output

- Replication Server standard error
- Replication Server RSSD output
- Replication Server RSSD ra output

12.1.4 sap_configure_rat

Allows you to configure the Replication Agent thread for SAP ASE (RepAgent for short).

Syntax

```

sap_configure_rat {all | <logical_host_name>}, {all | <database>} [, <property>,
<value>]
sap_configure_rat {redirect_to_er | redirect_to_ha}, {<database> | All}, <ER
admin user>, <ER admin password>

```

Parameters

redirect_to_er

Indicates primary RepAgent to redirect the connection to the external SAP Replication Server.

redirect_to_ha

Indicates primary RepAgent to redirect the connection to the standby SAP Replication Server.

<database> | All

Specifies the <database> parameter to execute the relevant operation to the specific database. Use All to execute the relevant operation to the whole HADR environment.

<ER admin user>

Specifies the admin user of the external SAP Replication Server to allow RMA to connect to it.

<ER admin password>

Specifies the admin user password of the external SAP Replication Server to allow RMA to connect to it.

Examples

Example 1

Configures the RAT properties:

```
sap_configure_rat site0, PI2,"max commit wait", 50
go
```

LOGICAL HOST DB NAME RESTART REQUIRED	PROPERTY NAME	CONFIG VALUE	RUN VALUE	DEFAULT VALUE
site0 10000000	PI2 NO	max commit wait	50	50

The returned result is:

LOGICAL HOST DB NAME RESTART REQUIRED	PROPERTY NAME	CONFIG VALUE	RUN VALUE	DEFAULT VALUE
site0 10000000	PI2 NO	max commit wait	50	50

Example 2

Displays the RAT properties:

i Note

In this example the value of the `scan timeout` parameter is 15 seconds. This is the default value and is recommended to maintain an optimum CPU usage. The value of `scan timeout` parameter must be 1 or more. 0 as the value of `scan timeout` parameter is not valid.

```
sap_configure_rat site0, PI2
go
```

The returned result is:

LOGICAL HOST DB NAME VALUE	HOST DB NAME DEFAULT VALUE	PROPERTY NAME RESTART REQUIRED	CONFIG VALUE	RUN
site0	PI2	rs servername	site1:4904	
site1:4904	n/a			
site0	PI2	rs username	PI2_RA_site1	
PI2_RA_site1	n/a			
site0	PI2	scan timeout	15	
15	15			
site0	PI2	retry timeout	60	
60	60			
site0	PI2	send warm standby xacts	true	
true	false			
site0	PI2	connect dataserver	PI2_site0_R2	
PI2_site0_R2	PI2			
site0	PI2	connect database	PI2	
PI2	PI2			
site0	PI2	send maint xacts to replicate	false	
false	false			

site0	PI2	trace log file	n/a	n/
a	n/a			
site0	PI2	security mechanism	n/a	n/
a	n/a			
site0	PI2	unified login	false	
false	false			
site0	PI2	msg confidentiality	false	
false	false			
site0	PI2	msg integrity	false	
false	false			
site0	PI2	msg replay detection	false	
false	false			
site0	PI2	msg origin check	false	
false	false			
site0	PI2	msg out-of-sequence check	false	
false	false			
site0	PI2	mutual authentication	false	
false	false			
site0	PI2	net password encryption	true	
true	true			
site0	PI2	priority	5	
5	5			
site0	PI2	auto start	true	
true	false			
site0	PI2	startup delay	0	
0	0			
site0	PI2	bind to engine	-1	
-1	-1			
site0	PI2	activate monitoring	false	
false	false			
site0	PI2	max schema cache per scanner	524288	
524288	524288			
site0	PI2	stream replication	true	
true	false			
site0	PI2	stream mode	async	
async	async			
site0	PI2	max stream retry	-1	
-1	-1			
site0	PI2	buffer pool size	80	
80	8			
site0	PI2	stream buffer size	1048576	
1048576	1048576			
site0	PI2	max commit wait	50	
50	10000000			
site0	PI2	max commands per package	20	
20	20			
site0	PI2	trace flags	n/a	n/
a	n/a			
(32 rows affected)				

Example 3

Redirect all databases in primary SAP ASE to connect to the external SAP Replication Server:

```
sap_configure_rat redirect_to_er, all, sa, Sybase123
```

The returned result is:

TASKNAME	TYPE	VALUE
Configure RAT Property	Start Time	Tue Oct 23 09:13:04 UTC 2018
Configure RAT Property	Elapsed Time	00:00:15
ConfigureRATProperty	Task Name	Configure RAT Property
ConfigureRATProperty	Task State	Completed
ConfigureRATProperty	Short Description	Configure an ASE Replication
Agent thread property.		
ConfigureRATProperty	Task Start	Tue Oct 23 09:13:04 UTC 2018
ConfigureRATProperty	Task End	Tue Oct 23 09:13:19 UTC 2018

```
ConfigureRATProperty      Hostname      rmawqiusite0.mo.sap.corp
```

Example 4

Redirect all databases in primary SAP ASE to connect to the standby SAP Replication Server:

```
sap_configure_rat redirect_to_ha, all, sa, Sybase123
```

The returned result is:

TASKNAME	TYPE	VALUE
Configure RAT Property	Start Time	Tue Oct 23 09:15:40 UTC 2018
Configure RAT Property	Elapsed Time	00:00:25
ConfigureRATProperty	Task Name	Configure RAT Property
ConfigureRATProperty	Task State	Completed
ConfigureRATProperty	Short Description	Configure an ASE Replication Agent thread property.
ConfigureRATProperty	Task Start	Tue Oct 23 09:15:40 UTC 2018
ConfigureRATProperty	Task End	Tue Oct 23 09:16:05 UTC 2018
ConfigureRATProperty	Hostname	rmawqiusite0.mo.sap.corp

12.1.5 sap_configure_rs

Use the `sap_configure_rs` command to list and configure the SAP Replication Server and its database connection properties. Use this command after the HADR environment is set up with the `sap_setup_replication` and `sap_materialize` commands

Syntax

```
sap_configure_rs { <logical_host_name> | 'all' } , { 'RS' | <database_name> | 'all' } [, <property_name>, <property_value> ]
```

Parameters

all

Specifies that all logical hosts must be configured.

<logical_host_name>

Specifies that only the specified logical host must be configured.

RS

Specifies that the command refers to the server-level properties.

all

Specifies that the command refers to the connection-level properties of all the database connections.

<database_name>

Specifies that the command refers to the connection-level properties of the specified database connection.

<property_name>

(optional) is the name of the property to be set.

<property_value>

(optional) is the new value for the property. If **<property_name>** and **<property_value>** are set, the specified property is set to the provided value. If **<property_name>** and **<property_value>** are not set, a list of all properties of the specified servers or connections is returned.

i Note

If a static property is set and a suspend or resume cycle has not been performed, the `Config Value` is different from the `Run Value` and the `Restart Required` value is set to `Yes`. To restart a connection, execute `sap_suspend_replication`, then `sap_resume_replication`.

Examples

Example 1

Lists all Replication Server properties on the logical host HA:

```

sap_configure_rs HA, RS
go
LOGICAL HOST RS NAME      DB NAME      PROPERTY NAME      CONFIG
VALUE              RUN VALUE      RESTART REQUIRED      DEFAULT
VALUE
-----
HA                PI2_REP_HA      stream_replication
false              false
false
HA                PI2_REP_HA      cap_prs_num_threads
2                  2
2
HA                PI2_REP_HA      cap_sqm_write_request_limit
8388608            8388608
8388608
HA                PI2_REP_HA      cap_sqm_write_msg_limit
5000                5000
5000
HA                PI2_REP_HA      spq_file_directory      The path of
the first partition      The path of the first partition      The path of
the first partition
HA                PI2_REP_HA      spq_min_size
2048                2048
1024

```

Example 2

Sets the Replication Server property `min_password_len` to 12 on all Replication Servers in the DR environment:

```

sap_configure_rs all, RS, min_password_len, 12

```

```

go
LOGICAL HOST RS NAME      DB NAME      PROPERTY NAME      CONFIG VALUE RUN VALUE
DEFAULT VALUE RESTART REQUIRED
-----
PR          PI2_REP_PR      min_password_len 12          12
6          NO
HA          PI2_REP_HA      min_password_len 12          12
6          NO
(2 rows affected)

```

Example 3

Sets the Replication Server Database Connection property `dsi_row_count_validation` to on for all database connections on all Replication Servers in the DR environment:

```

sap_configure_rs all, all, dsi_row_count_validation, on
go
LOGICAL HOST RS NAME      DB NAME      PROPERTY NAME      CONFIG VALUE
RUN VALUE  DEFAULT VALUE RESTART REQUIRED
-----
PR          PI2_REP_PR      PI2_PR.master  undefined        undefined
undefined  undefined      NO
PR          PI2_REP_PR      PI2_PR.PI2     undefined        undefined
undefined  undefined      NO
PR          PI2_REP_PR      PI2_PR.db1     undefined        undefined
undefined  undefined      NO
HA          PI2_REP_HA      PI2_HA.master  dsi_row_count_validation on
on         on         NO
HA          PI2_REP_HA      PI2_HA.PI2     dsi_row_count_validation on
on         on         NO
HA          PI2_REP_HA      PI2_HA.db1     dsi_row_count_validation on
on         on         NO
(6 rows affected)

```

Example 4

Sets the Replication Server Database Connection property `dsi_row_count_validation` to on for all PI2 database connections on all Replication Servers in the DR environment:

```

sap_configure_rs all, PI2, dsi_row_count_validation, on
go
LOGICAL HOST RS NAME      DB NAME      PROPERTY NAME      CONFIG VALUE RUN
VALUE  DEFAULT VALUE RESTART REQUIRED
-----
PR          PI2_REP_PR      PI2_PR.PI2  undefined        undefined
undefined  undefined      NO
HA          PI2_REP_HA      PI2_HA.PI2  dsi_row_count_validation on
on         on         NO

```

Usage

The command returns results in this format:

Logical Host, RS Name, DB Name, Property Name, Config Value, Run Value, Default Value, Restart Required

- Logical Host – logical host of the Replication Server being configured.

- `RS Name` – name of the Replication Server being configured.
- `DB Name` – name of the database connection being configured. This column is not populated when listing server properties.
- `Property Name` – the name of the property.
- `Config Value` – configured property value.
- `Run Value` – current property value used by the Replication Server.
- `Default Value` – default property value.
- `Restart Required` – Yes or No. If the value is Yes, then a restart is required for the property change to take effect. When listing properties, this column is blank or Yes.

12.1.6 sap_delay_replication

Delays replication to either a specified database or all the participating databases from the site. Delaying replication from the primary database provides time to recover from any undesirable event, such as when a table of records is dropped unexpectedly.

Although you can specify the default delay time with `sap_set delay_time_minutes`, issuing `sap_set` does not enable delayed replication. Instead, use `sap_delay_replication logical_host_name, on` after you configure the HADR environment with the `sap_setup_replication` and `sap_materialize` commands to enable delayed replication.

Syntax

```
sap_delay_replication <logical_hostname> [, <database_name>] [, on [, <delay_time_minutes>] |, off]
```

Parameters

<logical_host_name>

The logical host name of the current standby Replication Server.

<database_name>

The name of the database where replication is to be delayed.

<delay_time_minutes>

Delay time, in minutes. Valid values are 1 to 1439.

on

Enables delayed replication.

off

Disables delayed replication.

Examples

Example 1

Enables delayed replication on site HA for all the participating databases, using a default delay time value set by `sap_set delay_time_minutes`:

```
sap_delay_replication HA, on
go
TASKNAME          TYPE
VALUE

-----
-----
-----
Delay Replication Start Time      Mon Nov 23 05:03:00 EST
2015

Delay Replication Elapsed Time
00:00:05

DelayReplication Task Name        Delay
Replication

DelayReplication Task State
Completed

DelayReplication Short Description Toggle delayed replication to the ERP
database.

DelayReplication Long Description Turn on delayed replication for logical
host name 'HA', dataserver 'PI2_HA', and with delay minutes of '2'. Check
delayed database status by sap_delay_replication <logical_host_name>.
DelayReplication Task Start      Mon Nov 23 05:03:00 EST
2015

DelayReplication Task End        Mon Nov 23 05:03:05 EST
2015

DelayReplication Hostname        site0
(9 rows affected)
```

Example 2

Enables delayed replication and sets the delay time value to two minutes on all participating databases configured for site HA:

```
sap_delay_replication HA, on, 2
go
TASKNAME          TYPE
VALUE

-----
-----
-----
-----
```

```

Delay Replication Start Time      Mon Nov 23 05:03:00 EST
2015

Delay Replication Elapsed Time
00:00:05

DelayReplication  Task Name      Delay
Replication

DelayReplication  Task State
Completed

DelayReplication  Short Description Toggle delayed replication to the ERP
database.

DelayReplication  Long Description Turn on delayed replication for logical
host name 'HA', dataserver 'PI2_HA', and with delay minutes of '2'. Check
delayed database status by sap_delay_replication <logical_host_name>.
DelayReplication  Task Start      Mon Nov 23 05:03:00 EST
2015

DelayReplication  Task End        Mon Nov 23 05:03:05 EST
2015

DelayReplication  Hostname
site0

(9 rows affected)

```

Example 3

Enables delayed replication and sets the delay time value to two minutes on the PI2 database configured for site HA:

```

sap_delay_replication HA, PI2, on, 2
go
TASKNAME          TYPE
VALUE

-----
-----
-----
Delay Replication Start Time      Mon Nov 23 05:08:14 EST
2015

Delay Replication Elapsed Time
00:00:01

DelayReplication  Task Name      Delay
Replication

DelayReplication  Task State
Completed

DelayReplication  Short Description Toggle delayed replication to the ERP
database.

DelayReplication  Long Description Turn on delayed replication for logical
host name 'HA', dataserver 'PI2_HA', database name 'PI2', and with delay
minutes of '2'.

```



```

DelayReplication Task Start      Mon Nov 23 05:08:14 EST
2015

DelayReplication Task End        Mon Nov 23 05:08:15 EST
2015

DelayReplication Hostname
site0

(9 rows affected)

```

Example 4

Display the delay time (in minutes) on site HA:

```

sap_delay_replication HA
go
DATABASE NAME DELAY STATE RUNTIME VALUE DEFAULT VALUE
-----
PI2            ON           2           0
db1            ON           2           0
(2 rows affected)

```

Example 5

Disables the delayed replication on site HA for all participating databases:

```

sap_delay_replication HA, off
go
TASKNAME          TYPE
VALUE
-----
Delay Replication Start Time      Mon Nov 23 05:10:14 EST
2015

Delay Replication Elapsed Time
00:00:04

DelayReplication Task Name          Delay
Replication

DelayReplication Task State
Completed

DelayReplication Short Description Toggle delayed replication to the ERP
database.

DelayReplication Long Description Turn off delayed replication for logical
host name 'HA', dataserver 'PI2_HA', and with delay minutes of '0'. Check
delayed database status by sap_delay_replication <logical_host_name>.
DelayReplication Task Start      Mon Nov 23 05:10:14 EST
2015

DelayReplication Task End          Mon Nov 23 05:10:18 EST
2015

```

```
DelayReplication Hostname
site0
```

```
(9 rows affected)
```

Example 6

Disables the delayed replication for database PI2 on site HA:

```
sap_delay_replication HA, PI2, off
go
TASKNAME          TYPE
VALUE
-----
-----
Delay Replication Start Time      Mon Nov 23 05:09:17 EST
2015
Delay Replication Elapsed Time    00:00:02

DelayReplication Task Name        Delay
Replication

DelayReplication Task State
Completed

DelayReplication Short Description Toggle delayed replication to the ERP
database.
DelayReplication Long Description Turn off delayed replication for logical
host name 'HA', dataserver 'PI2_HA', database name 'PI2'.
DelayReplication Task Start       Mon Nov 23 05:09:17 EST
2015
DelayReplication Task End         Mon Nov 23 05:09:19 EST
2015
DelayReplication Hostname
site0

(9 rows affected)
```

12.1.7 sap_disable_external_replication

Disables replication to an external Replication Server for either a specific database or for all the databases.

Syntax

```
sap_disable_external_replication[, <database>]
```

Parameters

<database>

Specifies the name of the database for which external replication is to be disabled. If you do not specify this parameter, external replication is disabled for all the databases.

Examples

Example 1

Disables replication to the external Replication Server for database PI2:

```
sap_disable_external_replication PI2
```

The returned result is:

```
[Disable External Replication, Start Time, Tue Sep 13 06:05:25 EDT 2016]
[Disable External Replication, Elapsed Time, 00:00:08]
[DisableExternalReplication, Task Name, Disable External Replication]
[DisableExternalReplication, Task State, Completed]
[DisableExternalReplication, Short Description, Disable the flow of External
Replication]
[DisableExternalReplication, Long Description, Successfully disabled external
replication for database 'PI2'. Please execute
'sap_enable_external_replication PI2' to enable external replication for the
database.]
[DisableExternalReplication, Task Start, Tue Sep 13 06:05:25 EDT 2016]
[DisableExternalReplication, Task End, Tue Sep 13 06:05:33 EDT 2016]
[DisableExternalReplication, Hostname, site0]
```

12.1.8 sap_disable_replication

Stops replication for a specific database or all SAP databases.

Syntax

```
sap_disable_replication <primary_logical_host_name> [,
<standby_logical_host_name>] [, <database_name>]
```

Parameters

<primary_logical_host_name>

Specifies the name of the logical host that identifies the primary site.

<standby_logical_host_name>

Specifies the name of the logical host that identifies the standby site.

<database_name>

Specifies the database name.

Examples

Example 1

Disables the replication from primary host site0 to standby host site1 for database PI2:

```
sap_disable_replication site0,site1,PI2
go
```

The returned result is:

```
TASKNAME          TYPE
-----
Disable Replication Start Time
Tue Sep 15 02:42:08 EDT 2015
Disable Replication Elapsed Time
00:00:04
DisableReplication Task Name
Disable Replication
DisableReplication Task State
Completed
DisableReplication Short Description
Disable the flow of Replication
DisableReplication Long Description
Successfully disabled Replication for database 'PI2'.
DisableReplication Task Start
Tue Sep 15 02:42:08 EDT 2015
DisableReplication Task End
Tue Sep 15 02:42:12 EDT 2015
DisableReplication Hostname
site0
(9 rows affected)
```

Example 2

In this example, the command fails and it finds the error from the row where TYPE value is Failing
Command Error Message:

```
sap_disable_replication site0,site1,ERP
go
```

The returned result is:

```
TASKNAME          TYPE
-----
Disable Replication Start Time
Fri Nov 20 00:24:01 EST 2015
Disable Replication Elapsed Time
00:00:00
DisableReplication Task Name
Disable Replication
DisableReplication Task State
Error
DisableReplication Short Description
```

```

        Disable the flow of Replication
DisableReplication Long Description
        Validating user specified arguments.
DisableReplication Failing Command
        sap_disable_replication site0, site1, ERP,
DisableReplication Failing Command Error Message
        Database name 'ERP' is not one of the participating databases:
[master, PI2].
DisableReplication Corrective Action
        sap_disable_replication source_logical_hostname [,
target_logical_hostname] [, database]
DisableReplication Task Start
        Fri Nov 20 00:24:01 EST 2015
DisableReplication Task End
        Fri Nov 20 00:24:01 EST 2015
DisableReplication Hostname
        site0
(12 rows affected)

```

Example 3

Disables the replication from primary host site0 for all databases.

```

sap_disable_replication site0
go

```

The returned result is:

```

TASKNAME          TYPE
-----
VALUE
-----
Disable Replication Start Time
        Fri Nov 20 00:22:35 EST 2015
Disable Replication Elapsed Time
        00:00:13
DisableReplication Task Name
        Disable Replication
DisableReplication Task State
        Completed
DisableReplication Short Description
        Disable the flow of Replication
DisableReplication Long Description
        Successfully disabled Replication for participating databases
'[master, saptools, PI2]'.
DisableReplication Task Start
        Fri Nov 20 00:22:35 EST 2015
DisableReplication Task End
        Fri Nov 20 00:22:48 EST 2015
DisableReplication Hostname
        site0
(9 rows affected)

```

Usage

When the replication stops, you can only restart replication from the specified database or all databases by enabling and rematerializing the affected databases.

12.1.9 sap_drop_device

Use the `sap_drop_device` command to drop an SAP Replication Server device space.

Syntax

```
sap_drop_device <logical_host_name>, <logical_device_name>
```

Parameter

`<logical_host_name>`

Indicates the logical host name of the SAP Replication Server on which the device space is to be dropped.

`<logical_device_name>`

Indicates the logical name of the device to be dropped.

Example

Example 1

Drop device named `part3` from `site0` SAP Replication Server:

```
sap_drop_device site0, part3
go
```

The returned result is:

```
Drop Replication Server Device Start Time      Tue May 14 07:47:11 UTC 2019
Drop Replication Server Device Elapsed Time    00:00:00
DropDevice Task Name                          Drop Replication Server Device
DropDevice Task State                          Completed
DropDevice Short Description                   Drop the Replication Server
device with the specified name.
DropDevice Long Description                     Successfully dropped device
'part3' on host 'rmawqiusite1.mo.sap.corp:14975'.
DropDevice Task Start                          Tue May 14 07:47:11 UTC 2019
DropDevice Task End                            Tue May 14 07:47:11 UTC 2019
DropDevice Hostname                            rmawqiusite1.mo.sap.corp
```

Related Information

[sap_add_device \[page 444\]](#)

12.1.10 sap_drop_host

Before the setup of the HADR system, lets you drop a host that was previously defined by using the `sap_set_host` command.

i Note

You cannot drop the host with the `sap_drop_host` command after you set up the HADR system. Instead, use the `sap_teardown` command to tear down the HADR system, then run the `sap_set_host` command and re-create the logical host.

Syntax

```
sap_drop_host <logical_host_name>
```

Parameters

`<logical_host_name>`

Specifies the logical host name to be dropped.

Examples

Example 1

Drops the logical host site0:

```
sap_drop_host site0  
go
```

The returned result is:

```
TASKNAME      TYPE  
              VALUE  
-----  
-----  
Drop Host     Start Time  
              Mon Nov 16 00:39:33 EST 2015  
Drop Host     Elapsed Time  
              00:00:00  
DropHostApi   Task Name
```

```

Drop Host
DropHostApi Task State
Completed
DropHostApi Short Description
Drop the logical host from the environment.
DropHostApi Long Description
Submission of the design change for a model property was
successful.
DropHostApi Task Start
Mon Nov 16 00:39:33 EST 2015
DropHostApi Task End
Mon Nov 16 00:39:33 EST 2015
DropHostApi Hostname
site0
(9 rows affected)

```

If the command fails, finds the error from the row where TYPE value is Failing Command Error
 Message:

```

sap_drop_host site0
go

```

The returned result is:

```

TASKNAME      TYPE
VALUE
-----
-----
-----
-----
Drop Host      Start Time
                Mon Nov 16 04:38:05 EST
2015
Drop Host      Elapsed Time
00:00:00
DropHostApi Task Name
                Drop
Host
DropHostApi Task State
Error
DropHostApi Short Description
                Drop the logical host from the
environment.
DropHostApi Failing Command Error Message
                An internal Assertion Exception was caught during run task:
com.sybase.
                ua.plugins.dr.exceptions.DRAssertException: No host was
found for the l
                ogical name <site0>.
DropHostApi Failing Exception
                Message: 'No host was found for the logical name <site0>.'
Cause:
''
DropHostApi Task Start

```



```
2015          Mon Nov 16 04:38:05 EST
                DropHostApi Task End
                Mon Nov 16 04:38:05 EST
2015
                DropHostApi Hostname
site0
                (10 rows affected)
```

12.1.11 sap_enable_external_replication

Enables replication to an external Replication Server for either a specific database or for all the databases.

Syntax

```
sap_enable_external_replication[, <database>]
```

Parameters

<database>

Specifies the name of the database for which you are enabling external replication. If you do not specify this parameter, external replication is enabled for all the databases.

Examples

Example 1

Enables replication to the external Replication Server for database PI2:

```
sap_enable_external_replication PI2
```

The returned result is:

```
[Enable External Replication, Start Time, Tue Sep 13 06:07:15 EDT 2016]
[Enable External Replication, Elapsed Time, 00:00:02]
[EnableExternalReplication, Task Name, Enable External Replication]
[EnableExternalReplication, Task State, Completed]
[EnableExternalReplication, Short Description, Enable the flow of External
Replication]
[EnableExternalReplication, Long Description, Successfully enabled external
replication for database 'PI2'. The second truncation point of spq agent for
database 'PI2' has been reset.]
```

```
[EnableExternalReplication, Task Start, Tue Sep 13 06:07:15 EDT 2016]
[EnableExternalReplication, Task End, Tue Sep 13 06:07:17 EDT 2016]
[EnableExternalReplication, Hostname, site0]
```

12.1.12 sap_enable_replication

Enables replication for either a specific database or all SAP databases.

Syntax

```
sap_enable_replication <primary_logical_host_name> [,
<standby_logical_host_name>] [, <database_name>]
```

Parameters

<primary_logical_host_name>

Specifies the name of the logical host that identifies the primary site.

<standby_logical_host_name>

Specifies the name of the logical host that identifies the standby site.

<database_name>

Specifies the database name.

Examples

Example 1

Enables the replication from primary host site0 to standby host site1 for database PI2:

```
sap_enable_replication site0,site1,PI2
go
```

The returned result is:

```
TASKNAME TYPE
VALUE
-----
-----
Enable Replication Start Time
Fri Nov 20 00:41:19 EST 2015
Enable Replication Elapsed Time
00:01:36
EnableReplication Task Name
Enable Replication
```

```

EnableReplication Task State
Completed
EnableReplication Short Description
Enable the flow of Replication
EnableReplication Long Description
Successfully enabled Replication for database 'PI2'.
EnableReplication Task Start
Fri Nov 20 00:41:19 EST 2015
EnableReplication Task End
Fri Nov 20 00:42:55 EST 2015
EnableReplication Hostname
site0
(9 rows affected)

```

Example 2

In this example, the command fails and it finds the error from the row where TYPE value is Failing

Command Error Message:

```

sap_enable_replication site0,site1,ERP
go

```

The returned result is:

```

TASKNAME          TYPE
-----
VALUE
-----
Enable Replication Start Time
Fri Nov 20 00:27:00 EST 2015
Enable Replication Elapsed Time
00:00:00
EnableReplication Task Name
Enable Replication
EnableReplication Task State
Error
EnableReplication Short Description
Enable the flow of Replication
EnableReplication Long Description
Validating user specified arguments.
EnableReplication Failing Command
sap_enable_replication site0, site1, ERP,
EnableReplication Failing Command Error Message
Database name 'ERP' is not one of the participating databases:
[master, PI2].
EnableReplication Corrective Action
sap_enable_replication source_logical_hostname [,
target_logical_hostname] [, database]
EnableReplication Task Start
Fri Nov 20 00:27:00 EST 2015
EnableReplication Task End
Fri Nov 20 00:27:00 EST 2015
EnableReplication Hostname
site0
(12 rows affected)

```

Example 3

Enables the replication from primary host 'site0' for all databases.

```

sap_enable_replication site0
go

```

The returned result is:

```
TASKNAME          TYPE
      VALUE
-----
Enable Replication Start Time
      Fri Nov 20 00:46:50 EST 2015
Enable Replication Elapsed Time
      00:03:09
EnableReplication Task Name
      Enable Replication
EnableReplication Task State
      Completed
EnableReplication Short Description
      Enable the flow of Replication
EnableReplication Long Description
      Successfully enabled Replication for participating databases
'[master, db1, PI2]'.
EnableReplication Task Start
      Fri Nov 20 00:46:50 EST 2015
EnableReplication Task End
      Fri Nov 20 00:49:59 EST 2015
EnableReplication Hostname
      site0
(9 rows affected)
```

Usage

Materialize the databases after you execute the `sap_enable_replication` command.

12.1.13 sap_failover

Failover is switching activity to the standby site in the event of a failure on the primary site.

A planned failure occurs on a schedule. Typically as part of a test or other exercise, a planned failure allows for an orderly sequence of steps to be performed to move processing to the standby site.

An unplanned failure is unscheduled, occurring unintentionally and without warning. However, a similar sequence of events occur as in a planned failover.

Use the `sap_failover` command to perform planned and unplanned failovers. The `sap_failover` command:

- Monitors replication to verify all paths from the primary database to the standby are complete. No remaining in-flight data to be replicated exists for all SAP databases, master, and SAP_SID.
- Suspends the Replication Server at the standby site from applying any additional data from the primary.
- Configures and starts Replication Agent threads for each database in the standby server.
- Reconfigures the Replication Server to accept activity from the standby database.

i Note

You cannot perform two `sap_failover` commands in parallel. That is, the first `sap_failover` command must complete before you issue a second.

Syntax

```
sap_failover <primary_logical_host_name>, <standby_logical_host_name>,  
<deactivate_timeout> [, force | unplanned | 'null'[,<drain_timeout>]]
```

Parameters

<primary_logical_host_name>

The name of the logical host that identifies the primary site.

<standby_logical_host_name>

The name of the logical host that identifies the standby site.

<deactivate_timeout>

Specifies the number of seconds the process will wait while deactivating the primary data server. If the timeout reached, the failover process terminates.

force

(Optional) Causes the failover process to continue if the timeout value is reached. Applicable for `deactivate` step. However, the failover may not be successful for a number of reasons (for example, if there is a huge SPQ backlog)

<drain_timeout>

(Optional) Specifies the number of seconds the process waits while draining the transaction log from primary ASE to Replication Server. If the timeout is reached, the process terminates. If not set, the timeout equals to the value of `deactivate_timeout` by default.

unplanned

(Optional) Specifies an unplanned failover.

Examples

Example 1

Performs a planned failover, when all the servers are up in the HADR system:

```
sap_failover PR HA 60  
go
```

The returned result is:

```
TASKNAME          TYPE  
-----  
VALUE  
-----  
Failover Start Time  
Thu Nov 19 20:36:39 EST 2015  
Failover Elapsed Time
```

```

00:00:00
DRExecutorImpl Task Name
Failover
DRExecutorImpl Task State
Running
DRExecutorImpl Short Description
Failover moves primary responsibility from current logical source to logical
target.
DRExecutorImpl Long Description
Started task 'Failover' asynchronously.
DRExecutorImpl Additional Info
Please execute command 'sap_status task' to determine when task 'Failover'
is complete.
Failover Task Name
Failover
Failover Task State
Running
Failover Short Description
Failover moves primary responsibility from current logical source to logical
target.
Failover Long Description
Waiting 3 seconds: Waiting for the end of data marker for database 'master'
to be received.
Failover Current Task Number
6
Failover Total Number of Tasks
18
Failover Hostname
site0
(14 rows affected)

```

Checks the status:

```

sap_status
go

```

```

TASKNAME          TYPE
      VALUE
-----
Failover Start Time
hu Nov 19 20:36:37 EST 2015
Failover Elapsed Time
00:00:06
Failover Task Name
Failover
Failover Task State
Completed
Failover Short Description
Failover moves primary responsibility from current logical source to logical
target.
Failover Long Description
Failover from source 'PR' to target 'HA' is complete. The target may be
unquiesced.
Failover Additional Information
Please run command 'sap_host_available PR' to complete disabling replication
from the old source, now that the target 'HA' is the new primary.
Failover Current Task Number
14
Failover Total Number of Tasks
14
Failover Task Start
19 20:36:37 EST 2015
Failover Task End
19 20:36:43 EST 2015
Failover Hostname

```

```
site0
(12 rows affected)
```

Example 2

Performs an unplanned failover, when the primary ASE server is down in the HADR system:

```
sap_failover PR HA 60 unplanned
go
```

The returned result is:

```
TASKNAME          TYPE
-----
VALUE
-----
Failover Start Time
Thu Nov 19 20:57:20 EST 2015
Failover Elapsed Time
00:00:00
DRExecutorImpl Task Name
Failover
DRExecutorImpl Task State
Running
DRExecutorImpl Short Description
Failover moves primary responsibility from current logical source to logical
target.
DRExecutorImpl Long Description
Started task 'Failover' asynchronously.
DRExecutorImpl Additional Info
Please execute command 'sap_status task' to determine when task 'Failover'
is complete.
Failover Task Name
Failover
Failover Task State
Running
Failover Short Description
Failover moves primary responsibility from current logical source to logical
target.
Failover Long Description
Waiting 3 seconds: Waiting for the end of data marker for database 'master'
to be received.
Failover Additional Info 2
The primary Replication Server 'site1:5005' is configured for
synchronization mode and was found running in synchronization mode.
Failover Current Task Number
7
Failover Total Number of Tasks
8
Failover Hostname
site0
(15 rows affected)
```

Checks the status:

```
sap_status
go
```

```
TASKNAME          TYPE
-----
VALUE
-----
Failover Start Time
hu Nov 19 20:57:18 EST 2015
Failover Elapsed Time
```

```

00:00:06
Failover Task Name
Failover
Failover Task State
Completed
Failover Short Description
Failover moves primary responsibility from current logical source to logical
target.
Failover Long Description
Failover from source 'PR' to target 'HA' is complete. The target may be
unquiesced.
Failover Additional Info
When the source site for 'PR' is available, please run command
'sap_host_available PR' to disable replication from that source, now that the
target 'HĀ' is the new primary.
Failover Additional Info 2
The primary Replication Server 'site1:5005' is configured for
synchronization mode and was found running in synchronization mode.
Failover Current Task Number
12
Failover Total Number of Tasks
12
Failover Task Start
19 20:57:18 EST 2015
Failover Task End
19 20:36:24 EST 2015
Failover Hostname
site0
(13 rows affected)

```

12.1.14 sap_failover_drain_to_er

The `sap_failover_drain_to_er` command makes sure that the incremental backlogs from the HADR cluster are drained to the external replication system. Use this command while performing a failover within an HADR cluster with external replication.

Syntax

```
sap_failover_drain_to_er <time_out> | skip [<dbName>]
```

Parameters

<time_out>

Specifies the number of seconds the command waits for the remaining backlogs to be fully applied to the external system. If timeout is reached and the draining of backlogs to the external replication system is not finished and the `sap_failover_drain_to_er` command reports an error, retry this command with a higher `<time_out>` value.

skip

Forces the failover process to continue without applying the remaining backlogs to the external replication system. The `skip` option disables replication to the external replication system, causing the external replicate databases be out of sync with the HADR cluster.

<dbName>

Specifies the database on the external system to which replication is disabled from the HADR cluster. If you do not specify a database name following the `skip` parameter, you disable replication from the HADR cluster to all databases on the external system.

Examples

Example 1

Drains the transaction backlogs to the external replication system with a timeout of 120 seconds:

```
sap_failover_drain_to_er 120
go
```

Use `sap_status` to check the status.

```
sap_status
go

TASKNAME          TYPE
VALUE
-----
-----
Status            Start Time      Wed Sep 07 12:01:00 UTC
2016
Status            Elapsed Time    00:00:36

FailoverDrainToER Task Name      Failover drain to
ER.

FailoverDrainToER Task State
Completed

FailoverDrainToER Short Description  Failover drain to ER deactivate old
replication path and activate new replication path for external replication
system.
FailoverDrainToER Long Description   Failover drain to ER completed.
Update mark.
FailoverDrainToER Additional Info    Please run command
'sap_host_available' to complete the failover
operation.
FailoverDrainToER Current Task Number
5

FailoverDrainToER Total Number of Tasks
5

FailoverDrainToER Task Start        Wed Sep 07 12:01:00 UTC
2016
FailoverDrainToER Task End          Wed Sep 07 12:01:36 UTC
2016
```

```
FailoverDrainToER Hostname
site0
```

Example 2

Skips the transfer of the transaction backlog to the external replication system when performing a failover within an HADR cluster:

```
sap_failover_drain_to_er skip
go
```

Use `sap_status` to check the status:

```
sap_status
go

TASKNAME          TYPE
VALUE
-----
-----
Status            Start Time          Wed Sep 07 12:23:29 UTC
2016
Status            Elapsed Time
00:00:09

FailoverDrainToER Task Name          Failover drain to
ER.

FailoverDrainToER Task State
Completed

FailoverDrainToER Short Description  Failover drain to ER deactivate old
replication path and activate new replication path for external replication
system.
FailoverDrainToER Long Description  Failover drain to ER completed.
Update mark.
FailoverDrainToER Additional Info    Please run command
'sap_host_available' to complete the failover
operation.
FailoverDrainToER Current Task Number
2

FailoverDrainToER Total Number of Tasks
2

FailoverDrainToER Task Start          Wed Sep 07 12:23:29 UTC
2016
FailoverDrainToER Task End            Wed Sep 07 12:23:38 UTC
2016
FailoverDrainToER Hostname
site0
```

Example 3

Disables the replication to the database called `erp` on the external system when performing a failover within an HADR cluster:

```
sap_failover_drain_to_er skip erp
go
```

Use `sap_status` to check the status:

```
sap_status
go
TASKNAME                TYPE                VALUE
-----
-----
Status                  Start Time          Wed Sep 14 05:13:34 EDT 2016
Status                  Elapsed Time        00:00:11
FailoverDrainToER      Task Name           Failover drain to ER.
FailoverDrainToER      Task State          Completed
FailoverDrainToER      Short Description   Failover drain to ER deactivate old
replication path and activate new replication path for external replication
system.
FailoverDrainToER      Long Description   Issuing command to suspend SPQ Agents.
FailoverDrainToER      Additional Info    Please run command
'sap_failover_drain_to_er ' to continue the failover drain to external
replication operation.
FailoverDrainToER      Current Task Number 1
FailoverDrainToER      Total Number of Tasks 1
FailoverDrainToER      Task Start         Wed Sep 14 05:13:34 EDT 2016
FailoverDrainToER      Task End           Wed Sep 14 05:13:45 EDT 2016
FailoverDrainToER      Hostname           site0
```

Usage

After you use `sap_failover_drain_to_er skip <dbName>` to disable replication to a database on the external system, run `sap_failover_drain_to_er <timeout>` to make sure all backlogs on other databases are drained to the replicate databases.

12.1.15 sap_help

Displays a list of available commands, or detailed information for the specified command.

Syntax

```
sap_help [<command>]
go
```

12.1.16 sap_host_available

Use the `sap_host_available` command to reconfigure the primary database as the new backup for the activity occurring at the standby site.

The `sap_host_available` command:

- Disables the Replication Agents on the requested site for the master and SAP_SID databases in an SAP environment so that no data is replicated out from this site.
- Reconfigures the Replication Server to not accept activity from the requested site.
- Purges the Replication Server queues of any possible in-flight data.
- Resets the Replication Server at the current standby site to allow application of future activity, in the event a subsequent failover back to the primary site is needed.

Syntax

```
sap_host_available <primary_logical_host_name> [, suspend]
```

Parameters

`<primary_logical_host_name>`

Specifies the logical host name of the primary.

`suspend`

(Optional) If you do not want to start replication, suspend the replication path from the new primary server.

Examples

Example 1

Reconfigures the primary database after planned failover:

```
sap_host_available PR
go
```

The returned result is:

```
TASKNAME          TYPE
-----
HostAvailable Start Time
                  Thu Nov 19 20:47:34 EST 2015
HostAvailable Elapsed Time
```

```

00:01:31
HostAvailable Task Name
HostAvailable
HostAvailable Task State
Completed
HostAvailable Short Description
Resets the original source logical host when it is available after
failover.
HostAvailable Long Description
Completed the reset process of logical host 'PR' receiving
replication from logical host 'HA'.
HostAvailable Current Task Number
10
HostAvailable Total Number of Tasks
10
HostAvailable Task Start
Thu Nov 19 20:47:34 EST 2015
HostAvailable Task End
Thu Nov 19 20:49:05 EST 2015
HostAvailable Hostname
site0
(11 rows affected)

```

Example 2

Reconfigures the primary database after planned failover after unplanned failover and restarting the primary ASE server:

```

sap_host_available PR
go

```

The returned result is:

```

TASKNAME          TYPE
-----
HostAvailable Start Time
Thu Nov 19 21:07:57 EST 2015
HostAvailable Elapsed Time
00:01:07
HostAvailable Task Name
HostAvailable
HostAvailable Task State
Completed
HostAvailable Short Description
Resets the original source logical host when it is available after
failover.
HostAvailable Long Description
Completed the reset process of logical host 'PR' receiving
replication from logical host 'HA'.
HostAvailable Current Task Number
10
HostAvailable Total Number of Tasks
10
HostAvailable Task Start
Thu Nov 19 20:47:34 EST 2015
HostAvailable Task End
Thu Nov 19 21:09:04 EST 2015
HostAvailable Hostname
site0
(11 rows affected)

```

12.1.17 sap_list_device

Use the `sap_list_device` command to list information for devices added to SAP Replication Server.

Syntax

```
sap_list_device <logical_host_name> [,<logical_device_name>]
```

Parameter

<logical_host_name>

Indicates the logical host name of the SAP Replication Server for which the device information is to be listed.

<logical_device_name>

Indicates the logical device name for which the information is to be listed. If not indicated, RMA lists information for all devices on the specified SAP Replication Server.

Example

Example 1

List information for all devices added in `site0` SAP Replication Server:

```
sap_list_device site0
go
```

The returned result is:

```
LOGICAL HOST LOGICAL PARTITION NAME PHYSICAL PARTITION NAME TOTAL SIZE USED
SIZE STATE
-----
site0      part01      /testenv7/partition1.dat 1024
144      ON-LINE///
site0      part2      /testenv7/partition2.dat 16
0      ON-LINE///
```

Example 2

List information for the device named `part01` added in `site0` SAP Replication Server:

```
sap_list_device site0, part01
go
```

The returned result is:

```
LOGICAL HOST LOGICAL PARTITION NAME PHYSICAL PARTITION NAME TOTAL SIZE USED  
SIZE STATE  
-----  
site0      part01      /testenv7/partition1.dat 1024  
144      ON-LINE///
```

Related Information

[sap_add_device \[page 444\]](#)

[sap_drop_device \[page 462\]](#)

12.1.18 sap_materialize

Performs the initial copy of data from one site to the other.

Syntax

```
sap_materialize auto | retry [, start] | external | imprint | finish [, force],  
<source_logical_hostname>, <target_logical_hostname>, <database> [,  
<number_of_stripes>] [, <username>, <password>]
```

Parameters

auto

Performs automatic materialization. The `auto` option is the only option available for the master database materialization to manage consistently which tables are copied. For other databases, you can use `auto` or manual materialization method.

start

Configures the replication to anticipate the dump marker, generated by the `dump` command.

retry

Retries automatic materialization.

external

Materializes a standby database without using the `dump` and `load` solution. The `external` keyword skips the materialization process and sets up an active replication path between the primary and the standby databases.

imprint

Validates materialization before starting the external materialization process. If the external process requires the database to be offline, use `imprint` to add the verification row.

finish

Verifies that the verification row exists in the standby database.

force

Bypasses the verification test and finish the materialization process.

<source_logical_hostname>

Specifies the logical host name of the source.

<target_logical_hostname>

Specifies the logical host name of the target.

<database>

Specifies the name of the database.

<number_of_stripes>

Specifies the number of stripes.

<username>

Specifies the user name.

<password>

Specifies the password.

Examples

Example 1

Automatically performs the initial data copy from primary host site0 to standby host site1 for master database:

```
sap_materialize auto,site0,site1,master
go
```

The returned result is:

```
TASKNAME      TYPE
-----
Materialize
      Start Time
      Fri Nov 20 01:13:51 EST 2015
      Elapsed Time
      00:00:02
      Task Name
      Materialize
      Task State
      Running
      Short Description
```



```

Materialize database
DRExecutorImpl Long Description
Started task 'Materialize' asynchronously.
DRExecutorImpl Additional Info
Please execute command 'sap_status task' to determine
when task 'Materialize' is complete.
Materialize Task Name
Materialize Task State
Running
Materialize Short Description
Materialize database
Materialize Long Description
Starting materialization of the master database from
source 'site0' to
target 'site1'.
Materialize Task Start
Fri Nov 20 01:13:51 EST 2015
Materialize Hostname
site0
PerformMasterMaterialization Task Name
Materialize the Master database
PerformMasterMaterialization Task State
Running
PerformMasterMaterialization Short Description
Materializes the Master database by using a bulk copy
utility to copy necessary tables.
PerformMasterMaterialization Long Description
!
PerformMasterMaterialization.FIXUP_PREPARE_STAGE_PRIMARY!
PerformMasterMaterialization Current Task Number
1
PerformMasterMaterialization Total Number of Tasks
14
PerformMasterMaterialization Hostname
site0
(20 rows affected)

```

Checks the status by using the `sap_status task` command:

```

sap_status task
go
TASKNAME      TYPE
VALUE
-----
-----
Status        Start Time
              Fri Nov 20 01:13:51 EST 2015
Status        Elapsed Time
              00:00:28
Materialize   Task Name
              Materialize
Materialize   Task State
              Completed
Materialize   Short Description
              Materialize database
Materialize   Long Description
              Completed automatic materialization of database 'master'
from source 'site0' to target 'site1'.
Materialize   Task Start
              Fri Nov 20 01:13:51 EST 2015
Materialize   Task End

```

```

                Fri Nov 20 01:14:19 EST 2015
Materialize Hostname
                site0
(9 rows affected)

```

Example 2

If the command fails, finds the error from the row where TYPE value is Failing Command Error Message:

```

sap_status task
go

```

The returned result is:

```

TASKNAME                TYPE
VALUE
-----
-----
-----
-----
-----
--
                Status                Start Time
2015                Fri Nov 20 01:20:23 EST

                Status                Elapsed Time
00:00:13

                Materialize                Task Name
Materialize

                Materialize                Task State
Error

                Materialize                Short Description
                Materialize
database

                Materialize                Long Description
                Stop the Replication Agent for database 'master' on host
'site0:4901' a
                nd data server
'PI2_site0_R2'.

                Materialize                Task Start

```

```

2015                               Fri Nov 20 01:20:23 EST
Materialize                          Task End
                                   Fri Nov 20 01:20:36 EST
2015
Materialize                          Hostname
site0
DropSubscriptionWithForce Task Name
force                               Drop Subscription with
DropSubscriptionWithForce Task State
Error
DropSubscriptionWithForce Short Description
Server                               Drop subscription with force from the Replication
DropSubscriptionWithForce Long Description
database replicatio                 Dropping Subscription 'PI2_site0_R2_master_sub' for
                                   n definition
'PI2_site0_R2_master_repdef'.
DropSubscriptionWithForce Hostname
site0
ResumeDIST                          Task Name
Thread                               Resume DIST
ResumeDIST                          Task State
Error
ResumeDIST                          Short Description
database.                           Resume the Replication Server distribution thread for a
ResumeDIST                          Long Description
if the command e                     Waiting 10 seconds: Waiting 10 seconds before checking

```

```

successfully.
ResumeDIST          Failing Command
admin who,
no_trunc

ResumeDIST          Failing Command Return Code
0

ResumeDIST          Failing Command Error Message
The command to inspect the threads in the Replication
Server
failed.

ResumeDIST          Failing Exception
SQL ERROR -- Message: 'JZ006: Caught IOException:
com.sybase.jdbc4.jdbc
.SybConnectionDeadException: JZ0C0: Connection is
already closed.' SQLS
tate: 'JZ006' Remote Code: '0' Cause: ''
Message: 'JZ0C1:
An IOException occurred which closed the connection.'
SQLState: 'JZ0C1
' Remote Code: '0' Cause: ''.
ResumeDIST          Hostname
site0

(23 rows affected)

```

Example 3

Manually performs the initial data copy from primary host site0 to standby host site1 for PI2 database:

```

sap_materialize start site0, site1, PI2
go

```

The returned result is:

```

TASKNAME          TYPE          VALUE
-----
-----
Materialize       Start Time
                  Fri Nov 20 01:33:30 EST 2015
Materialize       Elapsed Time
                  00:00:01
DRExecutorImpl   Task Name
                  Materialize
DRExecutorImpl   Task State
                  Completed
DRExecutorImpl   Short Description
                  Materialize database
DRExecutorImpl   Long Description
                  Started task 'Materialize' asynchronously.

```

```

DRExecutorImpl Additional Info
                Please execute command 'sap_status task' to determine
when task 'Materialize' is complete.
Materialize Task Name
Materialize
Materialize Task State
Materialize Completed
Materialize Short Description
Materialize database
Materialize Long Description
                Adding the subscription required for materialization of
database 'PI2' to the Replication Server on host 'site1'.
Materialize Task Start
                Fri Nov 20 01:33:30 EST 2015
Materialize Task End
                Fri Nov 20 01:33:31 EST 2015
Materialize Hostname
                site0
(14 rows affected)

```

To bring the database online in the target SAP ASE system:

1. Use the SAP ASE `dump database` command to manually dump the PI2 database.
2. Copy the dump file to the target ASE system.
3. Manually load the dumped PI2 database by running the `load database ASE` command in the target ASE system.

```

sap_materialize finish site0, site1, PI2
go

```

The returned result is:

TASKNAME	TYPE	VALUE
Materialize	Start Time	Fri Nov 20 01:43:39 EST 2015
Materialize	Elapsed Time	00:00:02
DRExecutorImpl	Task Name	Materialize
DRExecutorImpl	Task State	Running
DRExecutorImpl	Short Description	Materialize database
DRExecutorImpl	Long Description	Started task 'Materialize' asynchronously.
DRExecutorImpl	Additional Info	Please execute command 'sap_status task' to determine when task 'Materialize' is complete.
Materialize	Task Name	Materialize
Materialize	Task State	Running
Materialize	Short Description	Materialize database
Materialize	Long Description	Validating user specified arguments.
Materialize	Task Start	Fri Nov 20 01:43:39 EST 2015

```

Materialize      Hostname
                site0
(13 rows affected)

```

Usage

- During materialization, `sap_materialize` drops the database users from the database. If it cannot drop the users after 20 attempts (waiting 10 seconds between each attempt), it forcibly removes them with a `kill with force` command.
- When you manually materialize the replicate database using the `sap_materialize start` command, RMA prompts you to dump the database on the primary node with a specified label. When you materialize the replicate database automatically using `sap_materialize auto`, RMA dumps the database with the specified label internally. This ensures that the replication restarts only after the labeled database dump is loaded.

For example:

```

sap_materialize start PR HA PI2
go

```

The returned result is:

```

TASKNAME          TYPE          VALUE
-----
-----
Materialize       Start Time    Wed Sep 28 21:56:47 EDT 2016
Materialize       Elapsed Time  00:00:02
DRExecutorImpl   Task Name     Materialize
DRExecutorImpl   Task State    Running
DRExecutorImpl   Short Description Materialize database
DRExecutorImpl   Long Description Started task 'Materialize'
asynchronously.
DRExecutorImpl   Additional Info Please execute command 'sap_status
task' to determine when task 'Materialize' is complete.
Materialize       Task Name     Materialize
Materialize       Task State    Running
Materialize       Short Description Materialize database
Materialize       Long Description Adding the subscription required for
materialization of database 'PI2' to the Replication Server on host 'site1'.
Materialize       Task Start    Wed Sep 28 21:56:47 EDT 2016
Materialize       Hostname      site0
CreateSubscription Task Name     Create Subscription
CreateSubscription Task State    Running
CreateSubscription Short Description Create or define a new subscription for
a database replication definition.
CreateSubscription Long Description Defining subscription
'PI2_PR_R2_PI2_sub' for database replication definition
'PI2_PR_R2_PI2_repdef'.
CreateSubscription Hostname      site0
(18 rows affected)

```

Check the status by using the `sap_status` command (see the bold text):

```

sap_status
go
TASKNAME      TYPE
VALUE

```

```

-----
-----
-----
-----
-----
-----
-----
-----
-----
-----
-----
-----
-----
-----
-----

Materialize Start Time              Wed Sep 28 21:56:47 EDT
2016

Materialize Elapsed Time
00:00:03

Materialize Task Name
Materialize

Materialize Task State
Completed

Materialize Short Description Materialize
database

Materialize Long Description The prerequisite work for manually dumping and
loading database PI2 is finished. You can use "dump database PI2 to ... with
compression = 1, label = 'RMA_DUMP_LABEL'" to dump the database. Be sure to
use the label option named 'RMA_DUMP_LABEL'.
Materialize Task Start              Wed Sep 28 21:56:47 EDT
2016

Materialize Task End                Wed Sep 28 21:56:50 EDT
2016

Materialize Hostname
site0

(9 rows affected)

```

The label name is defined in the `/DM/RMA-/instances/AgentContainer/config/bootstrap.prop` and `/DM/RMA-/config/bootstrap.prop` files. The label defined in the `/DM/RMA-/instances/AgentContainer/config/bootstrap.prop` file has higher priority. You may modify the default label in the files. An example of label definition in the files is:

```
rsge.bootstrap.rma_dump_label=RMA_DUMP_LABEL
```

12.1.19 sap_pre_setup_check

Use the `sap_pre_setup_check` command to test operating system permissions, database user roles and privileges, and host network port availability. Use the `sap_pre_setup_check` command before you use the `sap_setup_replication` or `sap_materialize` command.

Syntax

```
sap_pre_setup_check <env_type>, <primary_logical_host_name>,
<standby_logical_host_name>
```

Parameters

<env_type>

Specifies the type of environment the presetup check process validates. DR Agent supports the disaster recovery `dr` option, which configures the environment using the ASE HADR.

<primary_logical_host_name>

Specifies the name of the logical host that identifies the primary site.

<standby_logical_host_name>

Specifies the name of the logical host that identifies the standby site

Examples

Example 1

Performs a set-up check on logical host `site0` and `site1`:

```
sap_pre_setup_check dr, site0, site1
go
```

The returned result is:

```
NAME RESOURCE TYPE      DIAGNOSTICS
-----
site1  DR Agent Disaster recovery environment is ready for setup replication
site0  DR Agent Disaster recovery environment is ready for setup replication
(2 rows affected)
```

Usage

Before executing the replication setup command, correct any errors returned by `sap_pre_setup_check`. `sap_setup_replication` executes the same set of tests as part of the setup process.

12.1.20 sap_purge_trace

Purge `rs_ticket` information inserted into the `rs_ticket_history` table.

Syntax

```
sap_purge_trace <logical_host_name>, <days_to_keep> [, <database_name>]
```


Parameter

<logical_host_name>

Specifies the logical host name that identifies the primary, standby, or DR site.

<days_to_keep>

Specifies the number of days for which you want to retain the trace information that was inserted into `rs_ticket_history`. Trace information inserted before the specified number of days is purged from the table. The valid values are positive integers between 1 and 365.

<database_name>

Specifies the name of the database on the specified logical host. If not specified, all participating databases on the specified logical host are affected.

Example

Example 1

```
sap_purge_trace site0,7,master
go
```

The returned result is:

TASKNAME	TYPE	VALUE
Execute sap_purge_trace	Start Time	Fri Sep 20 05:04:44 UTC 2019
Execute sap_purge_trace	Elapsed Time	00:00:00
PurgeTrace	Task Name	Execute sap_purge_trace
PurgeTrace	Task State	Completed
PurgeTrace	Short Description	Execute sap_purge_trace to delete
rs_ticket in rs_ticket_history		
PurgeTrace	Long Description	Executing purge trace on database
master.		
PurgeTrace	Task Start	Fri Sep 20 05:04:44 UTC 2019
PurgeTrace	Task End	Fri Sep 20 05:04:44 UTC 2019
PurgeTrace	Hostname	mo-9c9b2e60d.mo.sap.corp
(9 rows affected)		

Example 2

```
sap_purge_trace sitel,7
go
```

The returned result is:

TASKNAME	TYPE	VALUE
Execute sap_purge_trace	Start Time	Fri Sep 20 05:05:12 UTC 2019
Execute sap_purge_trace	Elapsed Time	00:00:00
PurgeTrace	Task Name	Execute sap_purge_trace
PurgeTrace	Task State	Completed
PurgeTrace	Short Description	Execute sap_purge_trace to delete
rs_ticket in rs_ticket_history		

```
PurgeTrace      Long Description  Executing purge trace trace on all
participating databases.
PurgeTrace      Task Start        Fri Sep 20 05:05:12 UTC 2019
PurgeTrace      Task End          Fri Sep 20 05:05:12 UTC 2019
PurgeTrace      Hostname          mo-9c9b2e60d.mo.sap.corp
(9 rows affected)
```

12.1.21 sap_resume_component

Use the `sap_resume_component` command to resume components on a replication path.

Syntax

```
sap_resume_component <cpnt_name> <src_hostname> <tgt_hostname> <db_name>
```

Parameters

<cpnt_name>

Specifies the name of the component to be resumed.

<src_hostname>

Specifies the source host name.

<tgt_hostname>

Specifies the target host name.

<db_name>

Specifies the database name.

Examples

Example 1

Fails to resume the DSI component from PRHA on database master:

```
sap_resume_component DSI PR HA master
```

The returned result is: to

```
TASKNAME          TYPE          VALUE
-----
--
Resuming Rep Component. Start Time      Mon Mar 13 12:27:28 CST 2017
```

```

Resuming Rep Component. Elapsed Time      00:00:27
ResumeComponent         Task Name       Resuming Rep Component.
ResumeComponent         Task State      Error
ResumeComponent         Short Description Resume component on certain
replication path.
ResumeComponent         Long Description Failed to resume 'DSI' on
'PR.HA.master'.
ResumeComponent         Failing Command sap_resume_component DSI, PR, HA,
master,
ResumeComponent         Additional Info Thread of 'PIlocal_HA.master' typed
'DSI' on 'PVG50908038A' is in down status.
ResumeComponent         Additional Info 2 Process component 'DSI' from 'PR'
to 'HA' on database 'master'.
ResumeComponent         Task Start      Mon Mar 13 12:27:28 CST 2017
ResumeComponent         Task End        Mon Mar 13 12:27:55 CST 2017
ResumeComponent         Hostname        site1
(12 rows affected)
(As the thread of DSI is in the abnormal state of "Down", the command fails
to resume the DSI component.)

```

Example 2

Resume the RATCI component from site PR to HA on database master:

```

sap_resume_component RATCI PR HA master
go

```

The returned result is:

TASKNAME	TYPE	VALUE
Resuming Rep Component.	Start Time	Mon Mar 13 12:33:24 CST 2017
Resuming Rep Component.	Elapsed Time	00:00:01
ResumeComponent	Task Name	Resuming Rep Component.
ResumeComponent	Task State	Completed
ResumeComponent	Short Description	Resume component on certain replication path.
ResumeComponent	Long Description	Successfully resumed 'RATCI' on 'PR.HA.master'.
ResumeComponent	Additional Info	RATCI resumed, but it may still take upto 60 seconds to be active, please check it later.
ResumeComponent	Task Start	Mon Mar 13 12:33:24 CST 2017
ResumeComponent	Task End	Mon Mar 13 12:33:25 CST 2017
ResumeComponent	Hostname	site1

(10 rows affected)

Example 3

Resume the CAP component from site PR to HA on database master:

```

sap_resume_component CAP PR HA master

```

The returned result is:

TASKNAME	TYPE	VALUE
Resume Rep Component.	Start Time	Fri Mar 24 02:36:00 EDT 2017
Resume Rep Component.	Elapsed Time	00:00:02
ResumeComponent	Task Name	Resume Rep Component.
ResumeComponent	Task State	Completed

```

ResumeComponent      Short Description Resume component on certain
replication path.
ResumeComponent      Long Description  Successfully resumed 'CAP' on
'PR.HA.master'.
ResumeComponent      Task Start       Fri Mar 24 02:36:00 EDT 2017
ResumeComponent      Task End        Fri Mar 24 02:36:02 EDT 2017
ResumeComponent      Hostname        site1
(9 rows affected)

```

Usage

The valid values for the `cpnt` parameter are:

- RAT
- RATCI
- CAP
- DIST
- RSI
- DSI
- ALL

Set the `cpnt` parameter to `ALL` to resume all the supported components.

12.1.22 sap_resume_replication

Use the `sap_resume_replication` command to resume the replication to a specified database or all databases that are in the participating databases list (master and ERP).

Syntax

To resume replication to a specified database:

```
sap_resume_replication <standby_logical_host_name> [,<database_name>]
go
```

To resume replication to all SAP databases:

```
sap_resume_replication <standby_logical_host_name>, all
go
```

Parameters

`<standby_logical_host_name>`

Specifies the logical host name of the standby server.

<database_name>

Specifies the name of the database.

12.1.23 sap_send_trace

Use the `sap_send_trace` command to refresh the latency calculation time.

Latency calculations are based on the most recent trace flag sent through the system. Internally, this command inserts an `rs_ticket` into the source database or databases. Latency is calculated from the most recent entry in the target database's `rs_ticket_history` table by the `sap_status task` command. While executing the `sap_send_trace` command, specify a database name. If you do not specify a database name, a trace is sent to all participating databases for that host.

Syntax

```
sap_send_trace <primary_logical_host_name> [,<database_name>]
```

Parameters

<primary_logical_host_name>

Specifies the logical host name of the current primary Replication Server.

<database_name>

Specifies the name of the database where its latency is to be monitored.

Examples

Example 1

Sends trace on primary logical host `site0` so that the latency for all participating databases is calculated by using the `sap_status path` command:

```
sap_send_trace site0
go
```

The returned result is:

TASKNAME	TYPE	VALUE

		Execute sap_send_trace Start Time

```

                Fri Nov 13 04:37:28 EST 2015
Execute sap_send_trace Elapsed Time
00:00:01
DomainSendTrace      Task Name
Execute sap_send_trace
DomainSendTrace      Task State
Completed
DomainSendTrace      Short Description
Execute sap_send_trace to send a trace through the
Replication system u
                sing rs_ticket
DomainSendTrace      Task Start
                Fri Nov 13 04:37:28 EST 2015
DomainSendTrace      Task End
                Fri Nov 13 04:37:29 EST 2015
DomainSendTrace      Hostname
                site0
(8 rows affected)

```

If the command fails, finds the error from the row where TYPE value is Failing Command Error Message:

```

sap_send_trace site0
go

```

The returned result is:

```

TASKNAME          TYPE
VALUE
-----
-----
-----
Execute sap_send_trace Start Time
2015              Mon Nov 16 04:42:02 EST
Execute sap_send_trace Elapsed Time
00:00:00
DomainSendTrace   Task Name
Execute           sap_send_trace
DomainSendTrace   Task State
Error
DomainSendTrace   Short Description
Execute sap_send_trace to send a trace through the
Replication system u
                sing
rs_ticket
DomainSendTrace   Task Start
2015              Mon Nov 16 04:42:02 EST
DomainSendTrace   Task End
2015              Mon Nov 16 04:42:02 EST

```

```

DomainSendTrace      Hostname
site0

SendTrace            Task Name
Send
Trace

SendTrace            Task State
Error

SendTrace            Short Description
Send a trace through the Replication system using
rs_ticket

SendTrace            Long Description
Executing send trace on participating
databases.

SendTrace            Failing Command
SendTrace.execute

SendTrace            Failing Command Return Code
2812

SendTrace            Failing Command Error Message
rs_ticket.
Error encountered while executing send trace via

SendTrace            Failing Exception
SQL ERROR -- Message: 'Stored procedure 'rs_ticket' not
found. Specify
owner.objectname or use sp_help to check whether the object
exists (sp_
help may produce lots of output).' SQLState: 'ZZZZZ' Remote
Code: '281
2' Cause: ''.
SendTrace            Hostname
site0

(17 rows affected)

```

Example 2

Sends trace on primary logical host site0 for database PI2 so that the replication latency of PI2 is calculated by using the `sap_status path` command:

```
sap_send_trace site0, PI2
go
```

The returned result is:

```

TASKNAME              TYPE              VALUE
-----
Execute sap_send_trace Start Time
Fri Nov 13 04:36:07 EST 2015
Execute sap_send_trace Elapsed Time
00:00:00
DomainSendTrace      Task Name
Execute sap_send_trace
DomainSendTrace      Task State

```

```

Completed
DomainSendTrace      Short Description
Execute sap_send_trace to send a trace through the
Replication system u
sing rs_ticket
DomainSendTrace      Task Start
Fri Nov 13 04:36:07 EST 2015
DomainSendTrace      Task End
Fri Nov 13 04:36:07 EST 2015
DomainSendTrace      Hostname
site0
(8 rows affected)

```

12.1.24 sap_set

Use the `sap_set` command to read or set the initial configuration parameter values that are required for setting up the replication system for ASE HADR.

Syntax

```
sap_set [<global_level_property>, <value1>[, <value2>] | <logical_host_name>, <property_name>, <value>]
```

Parameters

<global_level_property>

Set the values for the properties that apply to the whole environment, the properties are as follows:

Global Level Property	Description
<code>sap_sid</code>	Denotes the Cluster ID. Consists of three characters.
<code>maintenance_user</code>	Specifies the login name and password of the maintenance user. Maintenance user is the login used by Replication Server to apply data changes to the standby SAP ASE databases.
<code>installation_mode</code>	Specifies the HADR system type. For Custom HADR, the mode is nonBS.
<code>participating_databases</code>	Specifies the list of user-participating databases that can be set up in the environment. In SAP Replication Server 16.0 SP03 PLO4 and later, HADR system allows you to set up the maximum of 20 databases in total.

Global Level Property	Description
<code>memory_size</code>	Specifies the memory limit for SAP Replication Server instances on all HADR nodes. The value of <code>memory_size</code> ranges from 1 GB to 2097151 GB. You can set this property at any time, either before or after the HADR system is setup. If unset, RMA tunes the memory limit for SAP Replication Server automatically.

<logical_host_name>, <property_name>

Set the values for the properties that apply to the specified logical host, the properties are as follows:

Logical Host Property	Description
<code>sybase_dir</code>	Specifies the \$SYBASE home directory for SAP Replication Server software installation. This is a read-only property.
<code>ase_port</code>	Specifies the SAP ASE server port number of the logical host.
<code>ase_user</code>	Specifies the login name used to connect to ASE server of this logical host. This is a read-only property.
<code>ase_backup_server_port</code>	Specifies the backup server port number of the logical host.
<code>db_dump_dir</code>	Specifies the filesystem location where the automatic dump file is to be saved.
<code>delay</code>	Specifies whether to delay a site. The values are: <ul style="list-style-type: none"> • 0 – does not delay the site • Any positive value – the number of minutes in which to delay the site
<code>ase_instance</code>	Specifies the names of the SAP ASE server. This is a read-only property.
<code>synchronization_mode</code>	Specifies the synchronization mode of the logical host. The value can be <code>sync</code> and <code>async</code> .
<code>distribution_mode</code>	Specifies whether the distribution target of the logical host is local or remote: <ul style="list-style-type: none"> • Local – the logical host distributes the data change to its SAP Replication Server • Remote – the logical host distributes data changes to the logical host of another SAP Replication Server
<code>replication_port</code>	Specifies the port number of the Replication Server.
<code>rssd_port</code>	Specifies the port number of the Replication Server ERSSD.

Logical Host Property	Description
<code>rssd_ra_port</code>	Specifies the port number for Replication Agent of ERSSD.
<code>replication_user</code>	Specifies the login name used to log in to the Replication Server of this logical host. This is a read-only property.
<code>device_buffer_dir</code>	Specifies the filesystem location where the Replication Server partition file is saved.
<code>device_buffer_size</code>	Specifies the partition device size, in MB.
<code>simple_persistent_queue_dir</code>	Specifies the filesystem location where the Replication Server simple persistent queue is saved.
<code>simple_persistent_queue_size</code>	Specifies the minimum size of the Replication Server simple persistent queue, in MB.
<code>simple_persistent_queue_max_size</code>	Specifies the maximum size of the Replication Server simple persistent queue, in MB.
<code>rs_instance</code>	Specifies the Replication Server instance name of the logical host. This is a read-only property.
<code>internal_name</code>	The identical to the logical host name. This is a read-only property.
<code>dr_plugin_port</code>	Specifies the RMA port number of the logical host. This is a read-only property.

Examples

Example 1

Configures PI2 as the ID:

```
sap_set sap_sid PI2
```

Example 2

Configures the username and password of the maintenance user:

```
sap_set maintenance_user DR_maint <password>
```

Example 3

Specifies the HADR system type. For Custom HADR, the mode is `nonBS`:

```
sap_set installation_mode nonBS
```

Example 5

Sets the SAP ASE server port number to 5000:

```
sap_set PR ase_port 5000
```

Example 7

Saves the automatic dump file to /dba work/dump:

```
sap_set PR db_dump_dir /dbawork/dump
```

Example 8

Delays the site by 60 minutes:

```
sap_set PR delay 60
```

Example 9

Sets the synchronization mode of the logical host to `sync`:

```
sap_set PR synchronization_mode sync
```

Example 10

Sets the distribution mode of the logical host to `remote`:

```
sap_set PR distribution_mode remote
```

Example 11

Sets the distribution target of the logical host to `HA`:

```
sap_set PR distribution_target HA
```

Example 12

Sets the port number of the Replication Server to 5005:

```
sap_set PR replication_port 5005
```

Example 13

Sets the port number of the Replication Server ERSSD to 5006:

```
sap_set PR rssid_port 5006
```

Example 14

Sets the port number for Replication Agent of ERSSD to 5007:

```
sap_set PR rssid_ra_port 5007
```

Example 15

Saves the Replication Server partition file to /dbawork/stablequeue:

```
sap_set PR device_buffer_dir /dbawork/stablequeue
```

Example 16

Sets the partition device size to 256 MB:

```
sap_set PR device_buffer_size 256
```

Example 17

Saves the Replication Server simple persistent queue to `/dbawork/spq`:

```
sap_set PR simple_persistent_queue_dir /dbawork/spq
```

Example 18

Sets the minimum size of the Replication Server simple persistent queue to 2000 MB:

```
sap_set PR simple_persistent_queue_size 2000
```

Example 19

Sets the maximum size of the Replication Server simple persistent queue to 8000 MB:

```
sap_set PR simple_persistent_queue_max_size 8000
```

Usage

- Execute `sap_set` without any parameter to query the configured value of the `memory_size` parameter.

12.1.24.1 sap_set device_buffer_size

Use the `sap_set device_buffer_size` command to adjust the amount of memory (partition disk space) that is allocated to the device buffer.

For optimum performance of the Replication Server, the valid range of the buffer size is 256 MB to 1 TB.

Syntax

```
sap_set <logical_hostname>, device_buffer_size, <size>
```

Parameters

logical_hostname

Specifies the name of logical Replication Server host to set the device buffer size

size

Specifies the size of the buffer that you want to set, in megabytes.

Examples

Example 1

Shows how to change partition disk space to 300 MB, where vegas is the logical host::

```
sap_set vegas, device_buffer_size, 300
```

12.1.24.2 sap_set simple_persistent_queue_size

Use the `sap_set simple_persistent_queue_size` command to specify the minimum disk space to allocate to the simple persistent queue (SPQ) in the SAP Replication Server.

i Note

SAP Replication Server creates one simple persistent queue of the size that you specify for each database that is to be replicated. The default minimum disk space that it allocates to each SPQ is 1000 MB and it is equal to the default maximum disk space size. Specify the minimum and maximum space of the SPQ based on the physical environment of the HADR system. Adjusting the persistent queue size is optional.

During initialization, the SPQ creates two data files with sizes equal to:

```
simple_persistent_queue_size/2
```

As the SPQ fills with data, it increases in size, and the system creates additional SPQ data files. The maximum size of these data files is also:

```
simple_persistent_queue_size/2
```

The SPQ is full when the total size of all SPQ data files reaches the size of the `simple_persistent_queue_max_size`, and any data that can be truncated is automatically truncated.

i Note

The minimum SPQ size for the master database is controlled by the `spq_min_size` entry in the `<${SYBASE}>/DM/RMA-16_0/instances/AgentContainer/config/RS_DB_master.properties`. Alter this value before configuring a new HADR system. The default value is 500 MB. For example (in **bold** below):

```
# This file defines RepServer database connection properties that are set
# when replication is set up.
#
db_packet_size=16384
dsi_cmd_batch_size=65536
spq_min_size=500M
. . .
dsi_serialization_method=wait_after_commit
```

Change the parameter on the Replication Server command line to alter the minimum SPQ size for an existing HADR system.

You can change the values for all databases in `<${SYBASE}>/DM/RMA-16_0/instances/AgentContainer/config/RS_DB.properties`. The default values for the master and CID databases is 500MB.

Syntax

```
sap_set <logical_host_name>, simple_persistent_queue_size, <spq_dir_size>
```

Parameters

<logical_host_name>

Specifies the name of logical replication server host to set the simple persistent queue max size.

<spq_dir_size>

Specifies the queue size that you want to set in megabytes.

Examples

Example 1

Changes the minimum SPQ disk space to 10,000 MB, where vegas is the logical host:

```
sap_set vegas, simple_persistent_queue_size, 10000
```

12.1.24.3 sap_set simple_persistent_queue_max_size

Use the `sap_set simple_persistent_queue_max_size` command to specify the maximum disk space to allocate to the simple persistent queue (SPQ) in the SAP Replication Server.

i Note

SAP Replication Server creates one simple persistent queue of the size that you specify for each database that is to be replicated. The default minimum disk space that it allocates to each SPQ is 1000 MB and it is equal to the default maximum disk space size. Specify the minimum and maximum space of the SPQ based on the physical environment of the HADR system. Adjusting the persistent queue size is optional.

During initialization, the SPQ creates two data files with sizes equal to:

```
simple_persistent_queue_size/2
```

As the SPQ fills with data, it increases in size, and the system creates additional SPQ data files. The maximum size of these data files is also:

```
simple_persistent_queue_size/2
```

The SPQ is full when the total size of all SPQ data files reaches the size of the `simple_persistent_queue_max_size`, and any data that can be truncated is automatically truncated.

i Note

The minimum SPQ size for the master database is controlled by the `spq_min_size` entry in the `<$$SYBASE>/DM/RMA-16_0/instances/AgentContainer/config/RS_DB_master.properties`. Alter this value before configuring a new HADR system. The default value is 500MB. For example (in **bold** below):

```
# This file defines RepServer database connection properties that are set
# when replication is set up.
#
db_packet_size=16384
dsi_cmd_batch_size=65536
spq_min_size=500M
. . .
dsi_serialization_method=wait_after_commit
```

Change the parameter on the Replication Server command line to alter the minimum SPQ size for an existing HADR system.

You can change the values for all databases in `<$$SYBASE>/DM/RMA-16_0/instances/AgentContainer/config/RS_DB.properties`. The default values for the master and CID databases is 500MB.

Syntax

```
sap_set <logical_host_name>, simple_persistent_queue_max_size, <spq_dir_size>
```

Parameters

`<logical_host_name>`

Specifies the name of logical replication server host to set the simple persistent queue max size

`<spq_dir_size>`

Specifies the queue size that you want to set in megabytes.

Examples

Example 1

Changes the maximum SPQ disk space to 10,000 MB, where vegas is the logical host:

```
sap_set vegas, simple_persistent_queue_max_size, 10000
```

12.1.24.4 sap_set memory_size

Use the `sap_set memory_size` command to set the memory limit for SAP Replication Server instances on all HADR nodes. The value of `memory_size` ranges from 1 GB to 2097151 GB. You can set this property either before or after the HADR system is set up. If not set, RMA tunes the memory limit for SAP Replication Server automatically.

Syntax

```
sap_set memory_size, <value>
```

Example

Set the memory limit for SAP Replication Server instances to 100 GB on all HADR nodes:

```
sap_set memory_size 100
```

The returned result is:

TASKNAME	TYPE	VALUE
Set Property	Start Time	Thu Mar 28 03:19:59 UTC 2019
Set Property	Elapsed Time	00:00:04
SetModelProperty	Task Name	Set Property
SetModelProperty	Task State	Completed
SetModelProperty	Short Description	Set a configuration property.
SetModelProperty	Long Description	Successfully set property memory_size to 100.
SetModelProperty	Task Start	Thu Mar 28 03:19:59 UTC 2019
SetModelProperty	Task End	Thu Mar 28 03:20:03 UTC 2019
SetModelProperty	Hostname	rmawqiusitel.mo.sap.corp
(9 rows affected)		

Check the settings for the `memory_size` parameter:

```
sap_set
```

The returned result is:

PROPERTY	VALUE
maintenance_user	ERP_maint
sap_sid	ERP
installation_mode	BS
participating_databases	[master,ERP]
connection_timeout	5
connection_alloc_once	true
memory_size	100

Usage

Setting memory limit for SAP Replication Server using the `memory_size` parameter does the following:

- If the `memory_size` parameter is not set before the initial setup, RMA automatically tunes the memory limit for SAP Replication Server when setting up the system or adding and removing databases and the DR host.
- If the `memory_size` parameter is set before initial setup, RMA uses the value as the memory limit for SAP Replication Server when setting up and does not tune the memory limit for SAP Replication Server when adding and removing databases and the DR host.
- If the `memory_size` parameter is not set after initial setup, RMA changes the memory limit from the previous set value to the auto-calculated value and tunes the memory limit when adding and removing databases and the DR host.
- If the `memory_size` parameter is set after the initial setup, RMA changes the memory limit for SAP Replication Server from the auto-calculated value to the specified value and does not tune the memory limit for SAP Replication Server when adding and removing databases and the DR host.
- After executing `sap_set memory_size` to change the `memory_size` parameter, RMA uses the new value as the memory limit for SAP Replication Server immediately.
- If the execution of `sap_set memory_size` fails to change the memory limit for some SAP Replication Server instances, the memory limit in SAP Replication Server instances becomes inconsistent. Resolve the errors that cause the failure manually and then execute the command `sap_set memory_size` again. Check if the new value is set successfully using the `sap_set` command.

12.1.25 sap_set_databases

Creates a list of database names for setting up a new replication environment.

Syntax

```
sap_set_databases <database_name> [, <additional_database_name> ...]
```

Parameters

<database_name>

Specifies the name of the database to replicate.

<additional_database_name>

Specifies a comma-separated list of additional databases to replicate.

Examples

Example 1

Sets `master`, `db1`, and `PI2` as the databases to participate in the new replication environment.

```
sap_set_databases master db1 PI2
go
TASKNAME          TYPE          VALUE
-----
Set databases     Start Time    Fri Oct 28 04:04:04 EDT 2016
Set databases     Elapsed Time  00:00:02
SetDatabases      Task Name     Set databases
SetDatabases      Task State    Completed
SetDatabases      Short Description Set the databases that will participate in a
new replication environment.
SetDatabases      Long Description Successfully verified and persisted
databases: [master, db1, PI2].
SetDatabases      Task Start    Fri Oct 28 04:04:04 EDT 2016
SetDatabases      Task End      Fri Oct 28 04:04:06 EDT 2016
SetDatabases      Hostname      site0

(9 rows affected)
```

Usage

The database list is verified, persisted, and used to validate any SAP commands specifying a database name. If a command's database name does not exist in the database list, the command is rejected.

Do not set the `saptools` database for replication; SAP ASE HADR does not support the replication for the `saptools` database.

12.1.26 sap_set_host

Use the `sap_set_host` command to register a new HADR logical host. The logical host consists of an ASE server, a Replication Server, and an RMA.

Syntax

To define the primary logical host:

```
sap_set_host <logical_host_name>, <sap_ase_host_name>, <sap_ase_port_num>
```

To define the SAP Replication Server on a separate computer:

```
sap_set_host <logical_host_name>, <sap_ase_host_name>, <sap_ase_port_num>,
<rs_host_name>, <rs_port_num>
```

```
go
```

To define the DR agent server port on a separate computer:

```
sap_set_host <logical_host_name>, <sap_ase_host_name>, <sap_ase_port_num>,  
<rs_host_name>, <rs_port_num>, <dr_agent_port_num>  
go
```

Parameters

<logical_host_name>

Specifies the logical host name to reference the site. The <logical_host_name> must have 10 characters or fewer and contain only digits or letters.

<ase_host_name>

Specifies the TCP/IP host name of the SAP ASE data server.

<ase_port_num>

Specifies the TCP/IP port number the SAP ASE data server is listening on.

<rs_host_name>

Specifies the TCP/IP host name of the SAP Replication Server.

<rs_port_num>

Specifies the TCP/IP port number the SAP Replication Server is listening on.

<dr_agent_port_num>

Specifies the TCP/IP port number the DR Agent is listening on (DR Agent is on the same host as the SAP Replication Server).

12.1.27 sap_set_password

Use the `sap_set_password` command to set or change the password for `DR_admin` on the primary SAP ASE, SAP Replication Server, and on the standby SAP ASE and SAP Replication Server.

The primary and standby SAP ASE and SAP Replication Server contain the system administrator `DR_admin` login. The Replication Management Agent authenticates its login by attempting to log in to one of these servers. This requires the `DR_admin` password to be the same on all servers. After the execution of the `sap_set_password` command, you must log out and then back in to the Replication Management Agent for the password change to take effect.

Syntax

```
sap_set_password <current_password>, <new_password>
```

Parameters

<current_password>

Specifies the current password.

<new_password>

Specifies the new password.

Examples

Example 1

Changes the password for DR_admin on both the primary and standby sites:

```
sap_set_password *****, *****
```

The returned result is:

```
SID: ERP
Participated DB: master, ERP_1,ERP_2,ERP
primary host: rmazwang2site0.mo.sap.corp
standby host: rmazwang2site1.mo.sap.corp
Executing 'set replication off' on the ASE at rmazwang2site1.mo.sap.corp:12510
Executing 'alter login 'DR_admin' with password '*****' modify password
'*****' on the ASE at rmazwang2site1.mo.sap.corp:12510.
Executing 'set replication on' on the ASE at rmazwang2site1.mo.sap.corp:12510.
Executing 'set replication off' on the ASE at rmazwang2site0.mo.sap.corp:
12895.
Executing 'alter login 'DR_admin' with password '*****' modify password
'*****' on the ASE at rmazwang2site0.mo.sap.corp:12895.
Executing 'set replication on' on the ASE at rmazwang2site0.mo.sap.corp:12895.
Executing 'alter user DR_admin set password '*****' verify password
'*****' on the Replication Server at rmazwang2site1.mo.sap.corp:12505.
Executing 'alter user DR_admin set password '*****' verify password
'*****' on the Replication Server at rmazwang2site0.mo.sap.corp:12890.
Executing 'alter user ERP_RA_site1 set password '*****' verify password
'*****' on the Replication Server at rmazwang2site1.mo.sap.corp:12505.
Executing 'alter user ERP_RA_site0 set password '*****' verify password
'*****' on the Replication Server at rmazwang2site0.mo.sap.corp:12890.
Executing 'use master' on the ASE at rmazwang2site1.mo.sap.corp:12510.
Executing 'set replication off' on the ASE at rmazwang2site1.mo.sap.corp:
12510.
Executing 'sp_config_rep_agent master,'rs password', '*****' on the ASE
at rmazwang2site1.mo.sap.corp:12510.
Executing 'set replication on' on the ASE at rmazwang2site1.mo.sap.corp:12510.
Executing 'use ERP_2' on the ASE at rmazwang2site1.mo.sap.corp:12510.
Executing 'set replication off' on the ASE at rmazwang2site1.mo.sap.corp:
12510.
Executing 'sp_config_rep_agent ERP_2,'rs password', '*****' on the ASE at
rmazwang2site1.mo.sap.corp:12510.
Executing 'set replication on' on the ASE at rmazwang2site1.mo.sap.corp:12510.
Executing 'use ERP' on the ASE at rmazwang2site1.mo.sap.corp:12510.
Executing 'set replication off' on the ASE at rmazwang2site1.mo.sap.corp:
12510.
Executing 'sp_config_rep_agent ERP,'rs password', '*****' on the ASE at
rmazwang2site1.mo.sap.corp:12510.
Executing 'set replication on' on the ASE at rmazwang2site1.mo.sap.corp:12510.
Executing 'use ERP_1' on the ASE at rmazwang2site1.mo.sap.corp:12510.
Executing 'set replication off' on the ASE at rmazwang2site1.mo.sap.corp:
12510.
```

```

Executing 'sp_config_rep_agent ERP_1,'rs password', '*****' on the ASE at
rmazwang2site1.mo.sap.corp:12510.
Executing 'set replication on' on the ASE at rmazwang2site1.mo.sap.corp:12510.
Executing 'use master' on the ASE at rmazwang2site0.mo.sap.corp:12895.
Executing 'set replication off' on the ASE at rmazwang2site0.mo.sap.corp:
12895.
Executing 'sp_config_rep_agent master,'rs password', '*****' on the ASE
at rmazwang2site0.mo.sap.corp:12895.
Executing 'set replication on' on the ASE at rmazwang2site0.mo.sap.corp:12895.
restart RAT
Executing 'select rep_agent_admin('is running', (select db_id('master')))' on
the ASE at rmazwang2site0.mo.sap.corp:12895.
Executing 'use ERP_2' on the ASE at rmazwang2site0.mo.sap.corp:12895.
Executing 'set replication off' on the ASE at rmazwang2site0.mo.sap.corp:
12895.
Executing 'sp_config_rep_agent ERP_2,'rs password', '*****' on the ASE at
rmazwang2site0.mo.sap.corp:12895.
Executing 'set replication on' on the ASE at rmazwang2site0.mo.sap.corp:12895.
restart RAT
Executing 'select rep_agent_admin('is running', (select db_id('ERP_2')))' on
the ASE at rmazwang2site0.mo.sap.corp:12895.
Executing 'use ERP' on the ASE at rmazwang2site0.mo.sap.corp:12895.
Executing 'set replication off' on the ASE at rmazwang2site0.mo.sap.corp:
12895.
Executing 'sp_config_rep_agent ERP,'rs password', '*****' on the ASE at
rmazwang2site0.mo.sap.corp:12895.
Executing 'set replication on' on the ASE at rmazwang2site0.mo.sap.corp:12895.
restart RAT
Executing 'select rep_agent_admin('is running', (select db_id('ERP')))' on
the ASE at rmazwang2site0.mo.sap.corp:12895.
Executing 'use ERP_1' on the ASE at rmazwang2site0.mo.sap.corp:12895.
Executing 'set replication off' on the ASE at rmazwang2site0.mo.sap.corp:
12895.
Executing 'sp_config_rep_agent ERP_1,'rs password', '*****' on the ASE at
rmazwang2site0.mo.sap.corp:12895.
Executing 'set replication on' on the ASE at rmazwang2site0.mo.sap.corp:12895.
restart RAT
Executing 'select rep_agent_admin('is running', (select db_id('ERP_1')))' on
the ASE at rmazwang2site0.mo.sap.corp:12895.
Executing 'alter user ERP_REP_site1_RSSD_prim set password '*****' verify
password '*****' on the Replication Server at rmazwang2site1.mo.sap.corp:
12505.
Executing 'alter user ERP_REP_site0_RSSD_prim set password '*****' verify
password '*****' on the Replication Server at rmazwang2site0.mo.sap.corp:
12890.
Executing 'alter user ERP_REP_site1_RSSD_maint identified by '*****' on
the Replication Server at rmazwang2site1.mo.sap.corp:12505.
Executing 'alter user ERP_REP_site0_RSSD_maint identified by '*****' on
the Replication Server at rmazwang2site0.mo.sap.corp:12890.
Executing 'alter connection to ERP_REP_site1_RSSD.ERP_REP_site1_RSSD set
password to '*****' on the Replication Server at
rmazwang2site1.mo.sap.corp:12505.
Executing 'alter connection to ERP_REP_site0_RSSD.ERP_REP_site0_RSSD set
password to '*****' on the Replication Server at
rmazwang2site0.mo.sap.corp:12890.
Executing 'sp_addexternlogin 'ERP_site0DR', 'DR_admin', 'DR_admin',
'*****' on the ASE at rmazwang2site1.mo.sap.corp:12510.
Executing 'sp_addexternlogin 'ERP_site1DR', 'DR_admin', 'DR_admin',
'*****' on the ASE at rmazwang2site0.mo.sap.corp:12895.
Executing 'set replication off' on the ASE at rmazwang2site1.mo.sap.corp:
12510.
Executing 'sp_addexternlogin ERP_site1_DRA, null, DR_admin, '*****',
replication_role' on the ASE at rmazwang2site1.mo.sap.corp:12510.
Executing 'set replication on' on the ASE at rmazwang2site1.mo.sap.corp:12510.
Executing 'set replication off' on the ASE at rmazwang2site0.mo.sap.corp:
12895.
Executing 'sp_addexternlogin ERP_site0_DRA, null, DR_admin, '*****',
replication_role' on the ASE at rmazwang2site0.mo.sap.corp:12895.

```

```
Executing 'set replication on' on the ASE at rmazwang2site0.mo.sap.corp:12895.
```

Usage

- The following login accounts use the same DR_admin password, so the execution of `sap_set_password` also resets the passwords for these login accounts:
 - Rep Agent user in SAP Replication Server and Replication Agent thread for SAP ASE configuration
 - RSSD primary user
 - RSSD maint user
 - HADR external logins
- If the password change fails, do not try to change the password again, or a password inconsistency issue might occur. Reset the passwords manually using the following procedures:

1. Check if you can log in to the primary SAP ASE server as DR_admin using the new password. If the login fails, it means that the password was not changed successfully. Use the following command to change the password:

```
set replication off
go
alter login 'DR_admin'
  with password <old password>
  modify password <new password>
go
set replication on
go
```

2. Log in to the standby SAP ASE server using the same process as for the primary server.
3. Using the new external login password, check if you can log in to the remote SAP ASE server by executing `<remote_ASE>...sp_help`. If the login fails, it means that the password is not changed successfully, use the following command to change the password:

```
set replication off
go
sp_addexternlogin <remote ASE server name>, DR_admin, DR_admin,
  <new password>
go

set replication on
go
```

4. Execute the same operation on the standby SAP ASE server as you performed on the primary server.
5. Using the new external login password, check if you can log in to the local RMA by executing `<local ASE Server name>_DRA...hadrstatuspath`. If the login fails, it means that the password was not changed successfully. Use the following command to change the password:

```
set replication off
go
sp_addexternlogin <local ASE Server name>_DRA, null, DR_admin,
  <new password>, replication_role
go
set replication on
go
```

6. Execute the same operation on the standby SAP ASE server as you performed on the primary server.

7. For all HADR participating databases on the primary SAP ASE, execute the following command:

```
set replication off
go
sp_config_rep_agent <database>, rs_password, <new password>
go
set replication on
go
```

8. Execute the same operation on the standby SAP ASE server as you performed on the primary server.
9. Restart RepAgent for the primary SAP ASE server. This step is not required for the standby SAP ASE server.
10. Using the new password, check if you can log in to the SAP Replication Server on the primary SAP ASE server as DR_admin. If the login fails, it means that the password was not changed successfully. Use the following command to change the password:

```
alter user DR_admin set password <new password> verify password <old password>
```

11. Execute the same operation on the SAP Replication Server on the standby SAP ASE server as you performed on the primary server.
12. Using the new password, check if you can log in to the SAP Replication Server on the primary SAP ASE server as the RepAgent user <SID>_RA_<logical host name>. If the login fails, it means that the password was not changed successfully. Use the following command to change the password:

```
alter user <SID>_RA_<logical host name> set password <new password> verify password <old password>
```

13. Execute the same operation on the SAP Replication Server on the standby SAP ASE server as you performed on the primary server.
14. Using the new password, check if you can log in to the SAP Replication Server on the primary SAP ASE server as the RSSD primary user <RS server name>_RSSD_prim. If the login fails, it means that the password was not changed successfully. Use the following command to change the password:

```
alter user <RS server name>_RSSD_prim set password <new password> verify password <old password>
```

15. Execute the same operation on the SAP Replication Server on the standby SAP ASE server as you performed on the primary server.
16. Using the new password, check if you can log in to RSSD on the primary SAP ASE server as the RSSD maintenance user <SID>_REP_<primary logical host name>_RSSD_maint. If the login fails, it means that the password was not changed successfully. Use the following command to change the password:

```
alter user <SID>_REP_<primary logical host name>_RSSD_maint identified by <new password>
```

17. Execute the same operation on the standby SAP ASE server as you performed on the primary server.
18. Check if the connection to <RS server name>_RSSD.<RS server name>_RSSD is available by logging in to the SAP Replication Server on the primary SAP ASE server. If the connection is unavailable, it means that the password was not changed successfully. Log in to RSSD as the RSSD primary user <RS server name>_RSSD_prim and execute the following command to change the password:

```
alter connection to <RS server name>_RSSD. <RS server name>_RSSD
set password to <new password>
```

19. Execute the same operation on the standby SAP ASE server as you performed on the primary server.

12.1.28 sap_set_replication_service

Use the `sap_set_replication_service` command to update the Replication Server Windows Service credentials.

Using the `sap_set_replication_service` command you can also restart the Replication Server using the Windows Service. The Windows Service does not exist until the `rs_init` utility is executed to create the Replication Server.

Syntax

```
sap_set_replication_service <logical_host_name>, [create | restart |, <domain_name>\<user_name>, <password>]
```

Parameters

<logical_host_name>

Specifies the name of the logical host.

create

Defines the Windows service on the logical host.

restart

Restarts the Replication Server on the logical host using the Windows Service for the Replication Server belonging to that logical host.

<domain_name>

Specifies the domain name.

<user_name>

Specifies the username.

<password>

Specifies the password.

Examples

Example 1

Defines a Windows Service on the local machine with the credentials provided:

```
sap_set_replication_service local, sapuser, sappassword
```


Example 2

Defines a Windows Service on myhost logical host for the Replication Server belonging to that logical host:

```
sap_set_replication_service myhost, create
```

By default, credentials default to the local system account login. You can change the login credentials using subsequent call to change the credentials.

Example 3

Restarts Replication Server on myhost logical host using the Windows Service for the Replication Server belonging to that logical host:

```
sap_set_replication_service myhost, restart
```

Example 4

Sets or changes the login credentials for the Windows Service on myhost logical host to use SAP username and SAP password:

```
sap_set_replication_service myhost, sapuser, sappassword
```

Usage

- If the user_name or password includes non alphanumeric characters, such as "@", enclose them in double quotation marks. For example:

```
sap_set_replication_service myhost, "sapuser", "sappassword@123"
```

- If you are in the local host domain, use `.\<user name>` instead of `localhost\<user name>`. For example, if `sapuser` is in the local host domain, the command is:

```
sap_set_replication_service myhost, ".\sapuser", "sappassword@123"
```

12.1.29 sap_setup_replication

Use the `sap_setup_replication` command to create a disaster recovery replication environment.

Syntax

```
sap_setup_replication, <env_type>, <primary_logical_host_name>,  
<standby_logical_host_name>
```

Parameters

<env_type>

Specifies the type of environment the presetup check process validates. DR Agent supports the disaster recovery "dr" option, which configures the environment using the SAP ASE HADR.

<primary_logical_host_name>

Specifies the name of the logical host that identifies the primary site.

<standby_logical_host_name>

Specifies the name of the logical host that identifies the standby site.

Usage

This command is executed asynchronously (in the background) and may take 30 or more minutes to complete. The order of the logical host names dictates the direction of replication (from primary to standby). The setup command returns immediately, indicating the setup task has been successfully started and is running asynchronously.

12.1.30 sap_sql_replication

Use the `sap_sql_replication` command to enable and disable SQL statement replication, configure threshold, and display SQL statement settings.

Syntax

```
sap_sql_replication {<database> | All}, {on | off} [,<option>[<option>][...]  
[,<table>[,<table>][,...]]]  
  
<option> ::= { U | D | I | S }  
sap_sql_replication {<database> | All}, display  
sap_sql_replication {<database> | All}, threshold, <value>
```

Parameter

<database>

Specifies to enable and disable SQL statement replication, or to configure and display SQL statement settings for a specific database.

All

Indicates the system to enable and disable SQL statement replication, or to configure and display SQL statement settings for the whole HADR environment.

on | off

Enables or disables SQL statement replication.

<option>[<option>][...]

Specifies the DML operations you want to enable or disable in SQL statement replication. The options are:

- U – update
- D – delete
- I – insert select
- S – select into

<table>[,<table>][,...]

Specifies to enable or disable SQL statement replication for specific tables by name.

- You can specify up to 100 tables in a single `sap_sql_replication` command. Execute `sap_sql_replication` multiple times if you have more than 100 tables.
- Use an asterisk (*) as a wildcard to enable or disable SQL statement replication for all tables.
- The table name settings are as follows:
 - Must start with an alphabetic character, an underscore or an asterisk.
 - Less than 255 bytes in length.

threshold, <value>

Specifies threshold to define when to trigger SQL statement replication. The threshold value is the minimum number of rows that affected by a SQL statement when SQL statement replication is triggered.

By default, SQL statement replication is triggered when the SQL statement affects more than 50 rows. You can adjust the value of threshold according to your needs. You can only set different threshold values at the database level.

To set the threshold for a specific database, specify the `<database>` parameter. Use `<ALL>` to set the threshold for the whole HADR environment.

The `<value>` parameter defines the minimum number of rows affected by a SQL statement when the SQL statement replication is triggered.

display

Displays SQL statement settings, such as the value of threshold and the tables that are enabled or disabled with SQL statement replication.

To display settings for a specific database, specify the `<database>` parameter. Use `<All>` to display settings for the whole HADR environment.

Example

Example 1

This example replicates update and delete statements as SQL statements for the ERP database:

```
sap_sql_replication ERP, on, UD
```

The command returns:

```
TASKNAME          TYPE
VALUE
-----
SQL Replication Start Time      Thu Sep 13 02:24:50 UTC
2018
SQL Replication Elapsed Time    00:00:01
SQLReplication Task Name        SQL
Replication Task State         Completed
SQLReplication Short Description Toggle SQL Replication in the
system
SQLReplication Long Description Enable the SQL Replication on Replication
Server.
SQLReplication Current Task Number
2
SQLReplication Total Number of Tasks
2
SQLReplication Task Start       Thu Sep 13 02:24:50 UTC
2018
SQLReplication Task End         Thu Sep 13 02:24:51 UTC
2018
SQLReplication Hostname         rmazwang2site0.mo.sap.corp
(11 rows affected)
```

Example 2

This example disables the replication of delete and insert select statements as SQL statements for all databases:

```
sap_sql_replication All, off, DI
```

The command returns:

```
TASKNAME          TYPE          VALUE
-----
SQL Replication Start Time      Thu Sep 13 02:24:19 UTC 2018
SQL Replication Elapsed Time    00:00:00
SQLReplication Task Name        SQL Replication
SQLReplication Task State         Completed
SQLReplication Short Description Toggle SQL Replication in the system
SQLReplication Long Description Disable the SQL Replication on ASE.
SQLReplication Current Task Number
2
SQLReplication Total Number of Tasks
2
SQLReplication Task Start       Thu Sep 13 02:24:19 UTC 2018
SQLReplication Task End         Thu Sep 13 02:24:19 UTC 2018
SQLReplication Hostname         rmazwang2site0.mo.sap.corp
(11 rows affected)
```

Example 3

This example replicates update, delete, and insert select statements as SQL statements for specific tables:

```
sap_sql_replication ERP_1, on, UDI, owner1.*, *x*y.a*b, *.sqldmltest_tb1,
owner3.sqldmltest_tb2, sqldmltest_tb3
```

The command returns:

```
TASKNAME          TYPE
VALUE
-----
SQL Replication Start Time      Thu Sep 13 02:28:33 UTC
2018
SQL Replication Elapsed Time    00:00:00
SQLReplication Task Name        SQL
Replication
SQLReplication Task State       Completed
SQLReplication Short Description Toggle SQL Replication in the
system
SQLReplication Long Description Enable the SQL Replication on Replication
Server.
SQLReplication Current Task Number
2
SQLReplication Total Number of Tasks
2
SQLReplication Task Start       Thu Sep 13 02:28:33 UTC
2018
SQLReplication Task End         Thu Sep 13 02:28:33 UTC
2018
SQLReplication Hostname         rmazwang2site0.mo.sap.corp
(11 rows affected)
```

Example 4

This example triggers SQL statement replication when the SQL statement affects more than 99 rows for the ERP database:

```
sap_sql_replication ERP, threshold, 99
```

The command returns:

```
TASKNAME          TYPE          VALUE
-----
SQL Replication Start Time      Thu Sep 13 02:29:36 UTC 2018
SQL Replication Elapsed Time    00:00:00
SQLReplication Task Name        SQL Replication
SQLReplication Task State       Completed
SQLReplication Short Description Toggle SQL Replication in the system
SQLReplication Long Description Set the SQL Replication threshold on ASE.
SQLReplication Current Task Number
1
SQLReplication Total Number of Tasks
1
SQLReplication Task Start       Thu Sep 13 02:29:36 UTC 2018
SQLReplication Task End         Thu Sep 13 02:29:36 UTC 2018
SQLReplication Hostname         rmazwang2site0.mo.sap.corp
(11 rows affected)
```

Example 5

This example triggers SQL statement replication when the SQL statement affects more than 50 rows for all databases:

```
sap_sql_replication All, threshold, 50
```

The command returns:

```
TASKNAME          TYPE                VALUE
-----
SQL Replication  Start Time          Thu Sep 13 02:24:19 UTC 2018
SQL Replication  Elapsed Time        00:00:00
SQLReplication   Task Name           SQL Replication
SQLReplication   Task State          Completed
SQLReplication   Short Description    Toggle SQL Replication in the system
SQLReplication   Long Description     Set the SQL Replication threshold on ASE.
SQLReplication   Current Task Number 1
SQLReplication   Total Number of Tasks 1
SQLReplication   Task Start          Thu Sep 13 02:24:19 UTC 2018
SQLReplication   Task End            Thu Sep 13 02:24:19 UTC 2018
SQLReplication   Hostname            rmazwang2site0.mo.sap.corp
(11 rows affected)
```

Example 6

This example displays the SQL statement settings for the database ERP:

```
sap_sql_replication ERP, display
```

The command returns:

```
DB_NAME    THRESHOLD  SQL_TYPE  LIST_TYPE  TABLE_LIST
-----
ERP        50         D         All        []
ERP        50         U         All        []
ERP        50         S         None       []
ERP        50         I         None       []
(4 rows affected)
```

Example 7

This example displays the SQL statement settings for the database ERP_1:

```
sap_sql_replication ERP_1, display
```

The command returns:

```
DB_NAME    THRESHOLD  SQL_TYPE  LIST_TYPE  TABLE_LIST
-----
ERP_1      50         D         In-List    [owner1.*, *x*y.a*b,
*.sqldmltest_tb1, owner3.sqldmltest_tb2, *.sqldmltest_tb3]
ERP_1      50         U         In-List    [owner1.*, *x*y.a*b,
*.sqldmltest_tb1, owner3.sqldmltest_tb2, *.sqldmltest_tb3]
ERP_1      50         S         None       []
ERP_1      50         I         In-List    [owner1.*, *x*y.a*b,
*.sqldmltest_tb1, owner3.sqldmltest_tb2, *.sqldmltest_tb3]
(4 rows affected)
```

Example 8

This example displays the SQL statement settings for the database ERP:

```
sap_sql_replication ERP, display
```

The command returns:

```
DB_NAME      THRESHOLD  SQL_TYPE  LIST_TYPE
TABLE_LIST
-----
ERP          50         D         None
[]
ERP          50         U         Out-List  [*.sqldmltest_tb1, *.sqldmltest_tb2,
owner3.sqldmltest_tb3, dbo.sqldmltest_tb4]
ERP          50         S         None
[]
ERP          50         I         None
[]
(4 rows affected)
```

Usage

The following table describes the output columns when specifying the `<display>` parameter:

Column	Description
DB_NAME	The name of the database to be displayed with the corresponding SQL statement settings.
THRESHOLD	The value of threshold that is being set for the indicated database. The threshold value is the minimum number of rows that affected by a SQL statement when SQL statement replication is triggered.
SQL_TYPE	The type of SQL operations, including: <ul style="list-style-type: none"> • U – update • D – delete • I – insert select • S – select into
LIST_TYPE	The type of list shown in the TABLE_LIST column: <ul style="list-style-type: none"> • All – SQL statement replication for all tables are enabled for the corresponding DML operation. The TABLE_LIST column does not show table details. • None – SQL statement replication for all tables are disabled for the corresponding DML operation. The TABLE_LIST column does not show table details. • In-List – the TABLE_LIST column lists tables that are enabled with SQL statement replication for corresponding DML operation. • Out-List – the TABLE_LIST column lists tables that are disabled with SQL statement replication for corresponding DML operation.
TABLE_LIST	The list of tables that are enabled or disabled with SQL statement replication for corresponding DML operation.

Related Information

[SQL Statement Replication \[page 349\]](#)

12.1.31 sap_status

Use the `sap_status` command to monitor the detailed status of a replication path, such as source and target ASE connection, source and target Replication Server connection, source ASE Replication agent, source Replication Server route thread (RSI) to target Replication Server, and target replication DSI thread to target ASE.

Syntax

```
sap_status [task | path | resource | clear | active_path | route |
synchronization | details, <source_logical_hostname>, <target_logical_hostname>,
<database>]
```

Parameters

task

Displays the asynchronous task status.

path

Checks the health and latency status of the replication paths.

resource

Monitors the Replication Server device usage and queue backlog and SAP ASE transaction log size and backlog.

clear

Clears the asynchronous task status output.

active_path

Displays the status of the active path from the primary site to the standby site.

route

Monitors the status of the servers, threads and queues of the replication path.

synchronization

Displays the synchronization mode and the synchronization state of the primary site.

details

Displays the detailed status of a replication path.

<source_logical_hostname>

Specifies the logical host name of the source.

<target_logical_hostname>

Specifies the logical host name of the target.

<database>

Specifies the name of the database.

Examples

Example 1

Output from checking the status of `sap_failover`.

```
Status      Start Time      Fri Aug 28 06:02:24 EDT
2015

Status      Elapsed Time
00:00:05

Failover    Task Name
Failover

Failover    Task State
Completed

Failover    Short Description  Failover moves primary responsibility from
current logical source to logical
target.
Failover    Long Description  Failover from source 'PR' to target 'HA' is
complete. The target may be
unquiesced.
Failover    Additional Info   Please run command 'sap_host_available PR'
to complete disabling replication from the old source, now that the target
'HA' is the new primary.
Failover    Current Task Number
14

Failover    Total Number of Tasks
14

Failover    Task Start        Fri Aug 28 06:02:24 EDT
2015

Failover    Task End          Fri Aug 28 06:02:29 EDT
2015

Failover    Hostname          site0
```

12.1.31.1 sap_status active_path

The `sap_status active_path` displays the replication paths that are active from the primary site to the standby site.

Syntax

```
sap_status active_path
```

Examples

Example 1

Displays the replication paths that are active from the primary site to the standby site:

```
sap_status active_path
go
```

The returned result is:

PATH	NAME	VALUE	
PR	Start Time	2015-11-19 21:19:12.764	Time command started executing.
PR	Elapsed Time	00:00:01	Command execution time.
PR	Hostname	site0	Logical host name.
PR	HADR Status	Primary : Active	Identify the primary and standby sites.
PR	Synchronization Mode	Synchronous	The configured Synchronization Mode value.
PR	Synchronization State	Synchronous	Synchronization Mode in which replication is currently operating.
PR	Distribution Mode	Remote	Configured value for the distribution_mode replication model property.
PR	Replication Server Status	Active	The status of Replication Server.
HA	Hostname	site1	Logical host name.
HA	HADR Status	Standby : Inactive	Identify the primary and standby sites.
HA	Synchronization Mode	Synchronous	The configured Synchronization Mode value.
HA	Synchronization State	Inactive	Synchronization Mode in which replication is currently operating.
HA	Distribution Mode	Remote	Configured value for the distribution_mode replication model property.
HA	Replication Server Status	Active	The status of Replication Server.
PR.HA.PI2	State	Active	Path is active and replication can occur.

```

PR.HA.PI2      Latency Time           Unknown           No latency
information for database 'PI2'.
PR.HA.PI2      Latency                 Unknown           No latency
information for database 'PI2'.
PR.HA.PI2      Commit Time           2015-09-23 04:38:45.852 Time last
commit replicated
PR.HA.PI2      Distribution Path      HA               The path of
Replication Server through which transactions travel.
PR.HA.PI2      Drain Status           Not Drained      The drain
status of the transaction logs of the primary database server.
PR.HA.db1      State                 Active           Path is
active and replication can occur.
PR.HA.db1      Latency Time           Unknown           No latency
information for database 'db1'.
PR.HA.db1      Latency                 Unknown           No latency
information for database 'db1'.
PR.HA.db1      Commit Time           2015-09-23 04:42:41.520 Time last
commit replicated
PR.HA.db1      Distribution Path      HA               The path of
Replication Server through which transactions travel.
PR.HA.db1      Drain Status           Not Drained      The drain
status of the transaction logs of the primary database server.
PR.HA.master  State                 Active           Path is
active and replication can occur.
PR.HA.master  Latency Time           Unknown           No latency
information for database 'master'.
PR.HA.master  Latency                 Unknown           No latency
information for database 'master'.
PR.HA.master  Commit Time           2015-09-23 04:24:52.112 Time last
commit replicated
PR.HA.master  Distribution Path      HA               The path of
Replication Server through which transactions travel.
PR.HA.master  Drain Status           Not Drained      The drain
status of the transaction logs of the primary database server.
(32 rows affected)

```

Usage

The returned information of this command is listed in the following table.

Table 9: Return Information for sap_status active_path

Information	Description
Start Time	The time at which the command starts to run.
Elapsed Time	The running time of the command.
Hostname	Logical host name.
HADR Status	HADR mode: HADR state or Unknown.
Synchronization Mode	The replication synchronization mode you have configured between a database and the SAP Replication Server, which can be one of: <ul style="list-style-type: none"> • Synchronous • Asynchronous

Information	Description
Synchronization State	<p>The current replication synchronization mode between a database and the SAP Replication Server, which can differ from the mode you have configured.</p> <div style="border: 1px solid #ccc; background-color: #f9f9f9; padding: 10px; margin-top: 10px;"> <p>i Note</p> <p>The synchronization state returned by the <code>sap_status active_path</code> command represents the state of all databases that are replicated by the primary site. If the synchronization state of the different databases is not the same (for example, if one database is in the 'synchronous' state and another is in the 'asynchronous' state), the result displayed by the <code>sap_status active_path</code> command for the site is 'Inconsistent', indicating the databases do not all have the same synchronization state at this time.</p> </div>
Distribution Mode	<p>The replication distribution mode you have configured between a database and the Replication Server, which can be one of:</p> <ul style="list-style-type: none"> • Local • Remote
Replication Server Status	<p>The status of the Replication Server, which can be one of:</p> <ul style="list-style-type: none"> • Active • Down • Unknown
State	<p>The status of the replication path, which can be one of:</p> <ul style="list-style-type: none"> • Defined: The state expected after setup and before materialization. • Suspended: The state expected during materialization, when data flow is suspended while waiting for load activities to complete. • Active: The replication path is supporting the replication. • Unknown: The server situation when encountering problems or errors.
Latency Time	<p>The <code>timestamp</code> of the most recent <code>trace</code> command that was applied to the target database and used for the latency calculation.</p>
Latency	<p>The approximate length of time it takes for an update on the source system to reach the target, based on the last <code>trace</code> command sent.</p>
Commit Time	<p>The local <code>timestamp</code> of a command applied to the target database.</p>
Distribution Path	<p>The logical host name of the distribution target server.</p>
Drain Status	<p>The status of draining the primary database server's transaction logs. Values are:</p> <ul style="list-style-type: none"> • Drained: The primary database server's transaction logs are completely transferred to Replication Server. • Not Drained: The primary database server's transaction logs are only partially transferred to Replication Server. • Unknown: The status cannot be queried.

i Note

To get the `<Latency Time>`, `<Latency>` and `<Commit Time>` parameter values, first execute the `sap_send_trace <primary logical host name>` command, then execute the `sap_status active_path` command.

12.1.31.2 sap_status path

The `sap_status path` command monitors information on the replication modes you have configured, the current replication states in the HADR with DR node environment, distribution mode and path, Replication Server status, and latency.

Syntax

```
sap_status path
```

Examples

Example 1

Monitors and returns the information on the replication modes you have configured, the current replication states in the HADR with DR node environment, distribution mode and path, Replication Server status, and latency:

```
sap_status path  
go
```

The returned result is:

PATH INFO	NAME	VALUE	
----- -----			
	Start Time	2016-04-27 22:33:04.026	Time command
started executing.	Elapsed Time	00:00:03	Command
execution time.	DR Hostname	site2	Logical host
name.	DR HADR Status	DR Standby : Inactive	Identify the
primary and standby sites.	DR Synchronization Mode	Asynchronous	The configured
Synchronization Mode value.	DR Synchronization State	Inactive	
DR Synchronization Mode in which replication is currently			
operating.	DR Distribution Mode	Local	Configured
value for the distribution_mode replication model property.	DR Replication Server Status	Active	The status of
DR Replication Server.	HA Hostname	site1	Logical host
HA name.	HA HADR Status	Standby : Inactive	Identify the
primary and standby sites.	HA Synchronization Mode	Synchronous	The configured
HA Synchronization Mode value.	HA Synchronization State	Inactive	
HA Synchronization Mode in which replication is currently			
operating.			

HA	Distribution Mode	Remote	Configured
value for the distribution_mode replication model property.			
HA	Replication Server Status	Active	The status of
Replication Server.			
PR	Hostname	site0	Logical host
name.			
PR	HADR Status	Primary : Active	Identify the
primary and standby sites.			
PR	Synchronization Mode	Synchronous	The configured
Synchronization Mode value.			
PR	Synchronization State	Synchronous	
Synchronization Mode in which replication is currently			
operating.			
PR	Distribution Mode	Remote	Configured
value for the distribution_mode replication model property.			
PR	Replication Server Status	Active	The status of
Replication Server.			
HA.DR.PI2	State	Suspended	Path is
suspended (Replication Agent Thread). Transactions are not being replicated.			
HA.DR.PI2	Latency Time	Unknown	No latency
information for database 'PI2'.			
HA.DR.PI2	Latency	Unknown	No latency
information for database 'PI2'.			
HA.DR.PI2	Commit Time	Unknown	No last commit
time for the database 'PI2'.			
HA.DR.PI2	Distribution Path	PR	The path of
Replication Server through which transactions travel.			
HA.DR.PI2	Drain Status	Unknown	The drain
status of the transaction logs of the primary database server.			
HA.DR.db1	State	Suspended	Path is
suspended (Replication Agent Thread). Transactions are not being replicated.			
HA.DR.db1	Latency Time	Unknown	No latency
information for database 'db1'.			
HA.DR.db1	Latency	Unknown	No latency
information for database 'db1'.			
HA.DR.db1	Commit Time	Unknown	No last commit
time for the database 'db1'.			
HA.DR.db1	Distribution Path	PR	The path of
Replication Server through which transactions travel.			
HA.DR.db1	Drain Status	Unknown	The drain
status of the transaction logs of the primary database server.			
HA.DR.master	State	Suspended	Path is
suspended (Replication Agent Thread). Transactions are not being replicated.			
HA.DR.master	Latency Time	Unknown	No latency
information for database 'master'.			
HA.DR.master	Latency	Unknown	No latency
information for database 'master'.			
HA.DR.master	Commit Time	Unknown	No last commit
time for the database 'master'.			
HA.DR.master	Distribution Path	PR	The path of
Replication Server through which transactions travel.			
HA.DR.master	Drain Status	Unknown	The drain
status of the transaction logs of the primary database server.			
HA.PR.PI2	State	Suspended	Path is
suspended (Replication Agent Thread). Transactions are not being replicated.			
HA.PR.PI2	Latency Time	Unknown	No latency
information for database 'PI2'.			
HA.PR.PI2	Latency	Unknown	No latency
information for database 'PI2'.			
HA.PR.PI2	Commit Time	Unknown	No last commit
time for the database 'PI2'.			
HA.PR.PI2	Distribution Path	PR	The path of
Replication Server through which transactions travel.			
HA.PR.PI2	Drain Status	Unknown	The drain
status of the transaction logs of the primary database server.			
HA.PR.db1	State	Suspended	Path is
suspended (Replication Agent Thread). Transactions are not being replicated.			

HA.PR.db1	Latency Time information for database 'db1'.	Unknown	No latency
HA.PR.db1	Latency information for database 'db1'.	Unknown	No latency
HA.PR.db1	Commit Time time for the database 'db1'.	Unknown	No last commit
HA.PR.db1	Distribution Path Replication Server through which transactions travel.	PR	The path of
HA.PR.db1	Drain Status status of the transaction logs of the primary database server.	Unknown	The drain
HA.PR.master	State suspended (Replication Agent Thread). Transactions are not being replicated.	Suspended	Path is
HA.PR.master	Latency Time information for database 'master'.	Unknown	No latency
HA.PR.master	Latency information for database 'master'.	Unknown	No latency
HA.PR.master	Commit Time time for the database 'master'.	Unknown	No last commit
HA.PR.master	Distribution Path Replication Server through which transactions travel.	PR	The path of
HA.PR.master	Drain Status status of the transaction logs of the primary database server.	Unknown	The drain
PR.DR.PI2	State and replication can occur.	Active	Path is active
PR.DR.PI2	Latency Time last calculated	2016-04-27 22:33:00.840	Time latency
PR.DR.PI2	Latency (ms)	707	Latency
PR.DR.PI2	Commit Time commit replicated	2016-04-27 22:33:00.840	Time last
PR.DR.PI2	Distribution Path Replication Server through which transactions travel.	HA	The path of
PR.DR.PI2	Drain Status status of the transaction logs of the primary database server.	Not Drained	The drain
PR.DR.db1	State and replication can occur.	Active	Path is active
PR.DR.db1	Latency Time last calculated	2016-04-27 22:33:00.832	Time latency
PR.DR.db1	Latency (ms)	703	Latency
PR.DR.db1	Commit Time commit replicated	2016-04-27 22:33:00.832	Time last
PR.DR.db1	Distribution Path Replication Server through which transactions travel.	HA	The path of
PR.DR.db1	Drain Status status of the transaction logs of the primary database server.	Not Drained	The drain
PR.DR.master	State and replication can occur.	Active	Path is active
PR.DR.master	Latency Time last calculated	2016-04-27 22:33:00.832	Time latency
PR.DR.master	Latency (ms)	703	Latency
PR.DR.master	Commit Time commit replicated	2016-04-27 22:33:00.832	Time last
PR.DR.master	Distribution Path Replication Server through which transactions travel.	HA	The path of
PR.DR.master	Drain Status status of the transaction logs of the primary database server.	Not Drained	The drain
PR.HA.PI2	State and replication can occur.	Active	Path is active
PR.HA.PI2	Latency Time last calculated	2016-04-27 22:33:00.640	Time latency
PR.HA.PI2	Latency (ms)	607	Latency
PR.HA.PI2	Commit Time commit replicated	2016-04-27 22:33:00.646	Time last
PR.HA.PI2	Distribution Path Replication Server through which transactions travel.	HA	The path of

```

PR.HA.PI2      Drain Status      Not Drained      The drain
status of the transaction logs of the primary database server.
PR.HA.db1     State      Active      Path is active
and replication can occur.
PR.HA.db1     Latency Time      2016-04-27 22:33:00.646 Time latency
last calculated
PR.HA.db1     Latency      610      Latency
(ms)
PR.HA.db1     Commit Time      2016-04-27 22:33:00.646 Time last
commit replicated
PR.HA.db1     Distribution Path      HA      The path of
Replication Server through which transactions travel.
PR.HA.db1     Drain Status      Not Drained      The drain
status of the transaction logs of the primary database server.
PR.HA.master State      Active      Path is active
and replication can occur.
PR.HA.master Latency Time      2016-04-27 22:33:00.632 Time latency
last calculated
PR.HA.master Latency      603      Latency
(ms)
PR.HA.master Commit Time      2016-04-27 22:33:00.632 Time last
commit replicated
PR.HA.master Distribution Path      HA      The path of
Replication Server through which transactions travel.
PR.HA.master Drain Status      Not Drained      The drain
status of the transaction logs of the primary database server.
(92 rows affected)

```

Usage

The returned information of this command is listed in the following table.

Table 10: Result Information for `sap_status path`

Information	Description
Start Time	The time point that the command starts to run.
Elapsed Time	The running time of the command.
Hostname	Logical host name.
HADR Status	HADR mode: HADR state or Unknown.
Synchronization Mode	One of two replication synchronization modes you have configured between a database and the SAP Replication Server: <ul style="list-style-type: none"> • Synchronous • Asynchronous

Information	Description
Synchronization State	<p>The current replication synchronization mode between a database and the SAP Replication Server, which can be different from the mode you have configured.</p> <div style="border: 1px solid #ccc; background-color: #f9f9f9; padding: 10px; margin-top: 10px;"> <p>i Note</p> <p>The synchronization state returned by the <code>sap_status path</code> command represents the state of all databases that are replicated by the primary site. If the synchronization state of the different databases is not the same (for example, if one database is in the synchronous state and another is in the asynchronous state), the result displayed by the <code>sap_status path</code> command for the site is <code>Inconsistent</code> - indicating the databases do not all have the same synchronization state at this time.</p> </div>
Distribution Mode	<p>One of two replication distribution modes you have configured between a database and the Replication Server:</p> <ul style="list-style-type: none"> • Local • Remote
Replication Server Status	<p>The status of the Replication Server, which can be one of:</p> <ul style="list-style-type: none"> • Active • Down • Unknown
State	<p>The status of the replication path, which can be one of:</p> <ul style="list-style-type: none"> • Defined: The state expected after setup and before materialization. • Suspended: The state expected during materialization, when data flow is suspended while waiting for load activities to complete. • Active: The replication path is supporting the replication. • Unknown: The server situation when encountering problems or errors.
Latency Time	<p>The timestamp of the most recent <code>trace</code> command that was applied to the target database and used to calculate latency.</p>
Latency	<p>The approximate length of time it takes for an update on the source system to reach the target, based on the last <code>trace</code> command sent.</p> <div style="border: 1px solid #ccc; background-color: #f9f9f9; padding: 10px; margin-top: 10px;"> <p>i Note</p> <p>During bulk materialization, the Replication Server holds the transactions in the outbound queue (OBQ) until the subscription marker is processed. The <code>sap_status path</code> command may report some latency in replication during this time. It can be ignored as it is just a difference between the previous <code>rs_ticket</code> and the current time.</p> <p>The <code>rs_ticket</code> stored procedure works with replicate database stored procedure <code>rs_ticket_report</code> to measure the amount of time it takes for a command to move from the primary database to the replicate database.</p> </div>
Commit Time	<p>The local timestamp of a command applied to the target database.</p>
Distribution Path	<p>The logical host name of the distribution target server.</p>

Information	Description
Drain Status	<p>The status of draining the primary database server's transaction logs. Values are:</p> <ul style="list-style-type: none"> • Drained: The primary database server's transaction logs are completely transferred to Replication Server. • Not Drained: The primary database server's transaction logs are only partially transferred to Replication Server. • Unknown: The status cannot be queried.

i Note

To get the `<Latency Time>`, `<Latency>` and `<Commit Time>` parameter values, first execute the `sap_send_trace <primary logical host name>` command, then execute the `sap_status active_path` command.

12.1.31.3 sap_status resource

Monitors the estimated failover time, Replication Server device size, simple persistent queue (SPQ) size, usage, backlog, replication truncation backlog (inbound queue and outbound queue), replication route queue truncation backlog, SAP ASE transaction log size and backlog, as well as stable queue backlogs.

Syntax

```
sap_status resource
```

Examples

Example 1

Returns resource information:

```
sap_status resource
go
```

The returned result is:

```
sap_status resource
go
NAME          TYPE
VALUE
-----
                Start Time                2016-06-02
22:38:20.891  Elapsed Time
00:00:02
```

	Estimated Failover Time
-1	
PR	Replication device size (MB)
256	
PR	Replication device usage (MB)
208	
DR	Replication device size (MB)
512	
DR	Replication device usage (MB)
64	
HA	Replication device size (MB)
256	
HA	Replication device usage (MB)
208	
HA.master	Replication simple persistent queue size (MB)
500	
HA.PI2	Replication simple persistent queue size (MB)
2000	
HA.db1	Replication simple persistent queue size (MB)
2000	
PR.master	ASE transaction log size (MB)
100	
PR.master	ASE transaction log backlog (MB)
0	
PR.db1	ASE transaction log size (MB)
20	
PR.db1	ASE transaction log backlog (MB)
0	
PR.PI2	ASE transaction log size (MB)
10	
PR.PI2	ASE transaction log backlog (MB)
0	
HA.master.SPQ	Replication simple persistent queue backlog (MB)
0	
HA.master.IBQ	Replication inbound queue backlog (MB)
0	
HA.master.IBQ	Replication inbound queue truncation backlog (MB)
0	
HA.DR.RQ	Replication route queue backlog (MB)
0	
HA.DR.RQ	Replication route queue truncation backlog (MB)
0	
DR.master.OBQ	Replication outbound queue backlog (MB)
0	
DR.master.OBQ	Replication outbound queue truncation backlog (MB)
0	
HA.db1.SPQ	Replication simple persistent queue backlog (MB)
0	
HA.db1.IBQ	Replication inbound queue backlog (MB)
0	
HA.db1.IBQ	Replication inbound queue truncation backlog (MB)
0	
DR.db1.OBQ	Replication outbound queue backlog (MB)
0	
DR.db1.OBQ	Replication outbound queue truncation backlog (MB)
0	
HA.PI2.SPQ	Replication simple persistent queue backlog (MB)
0	
HA.PI2.IBQ	Replication inbound queue backlog (MB)
0	
HA.PI2.IBQ	Replication inbound queue truncation backlog (MB)
0	
DR.PI2.OBQ	Replication outbound queue backlog (MB)
0	
DR.PI2.OBQ	Replication outbound queue truncation backlog (MB)
0	
HA.master.OBQ	Replication outbound queue backlog (MB)
0	

```

HA.master.OBQ Replication outbound queue truncation backlog (MB)
0
HA.db1.OBQ    Replication outbound queue backlog (MB)
0
HA.db1.OBQ    Replication outbound queue truncation backlog (MB)
0
HA.PI2.OBQ    Replication outbound queue backlog (MB)
0
HA.PI2.OBQ    Replication outbound queue truncation backlog (MB)
0
DR.PI2        Replication queue backlog (MB)
0
HA.PI2        Replication queue backlog (MB)
0
HA.master     Replication queue backlog (MB)
0
DR.db1        Replication queue backlog (MB)
0
DR.master     Replication queue backlog (MB)
0
HA.db1        Replication queue backlog (MB)
0
(47 rows affected)

```

Usage

This table lists the information the command returns:

Table 11: Return Information for `sap_status resource`

Information	Description	Value
Start Time	The time at which the command starts to run.	
Elapsed Time	The running time of the command.	
Estimated Failover Time	The failover time estimated by the system.	In the following conditions, the value is -1: <ul style="list-style-type: none"> • Replication Server has recently started and initialization is still underway. • The data server interface (DSI) thread in the Replication Server is inactive. • DR Agent has communication errors with Replication Server.
Replication device size (MB)	The disk space allocated for the Replication Server.	Displays "Unable to monitor the replication devices" if the Replication Server cannot be reached.

Information	Description	Value
Replication device usage (MB)	The disk space used by the Replication Server.	Displays "Unable to monitor the replication devices" if the Replication Server cannot be reached.
<div style="background-color: #e0e0e0; padding: 10px; border: 1px solid #ccc;"> <p>i Note</p> <p>If the device usage percentages returned from the command are high, consider adding device space to the replication paths to reduce the risk that the primary ASE transaction log will run out of space.</p> </div>		
Replication simple persistent queue size (MB)	The disk space allocated for the simple persistent queue.	Displays "Unable to monitor the replication devices" if the Replication Server cannot be reached.
ASE transaction log size (MB)	The disk space allocated for saving the transaction logs in the primary SAP ASE.	Displays "Unable to monitor the ASE transaction log" if the primary SAP ASE cannot be reached.
ASE transaction log backlog (MB)	The accumulated logs to be processed in the primary SAP ASE.	Displays "Unable to monitor the ASE transaction log" if the primary SAP ASE cannot be reached.
Replication simple persistent queue backlog (MB)	The accumulated logs to be processed in the simple persistent queue.	Displays "Unable to monitor the replication queues" if the Replication Server cannot be reached.
Replication inbound queue backlog (MB)	The accumulated logs to be processed in the inbound queue.	Displays "Unable to monitor the replication queues" if the Replication Server cannot be reached.
Replication route queue backlog (MB)	The accumulated logs to be processed in the route queue.	Displays "Unable to monitor the replication queues" if the Replication Server cannot be reached.
Replication outbound queue backlog (MB)	The accumulated logs to be processed in the outbound queue.	Displays "Unable to monitor the replication queues" if the Replication Server cannot be reached.
Replication queue backlog (MB)	The sum of the simple persistent queue backlog, inbound queue backlog, and outbound queue backlog.	Displays "Unable to monitor the replication queues" if the Replication Server cannot be reached.
Replication truncation backlog (MB)	The data in the Replication Server queues inbound queue (IBQ), outbound queue -(OBQ), and route queue (RQ) that cannot be truncated.	Displays "Unable to monitor the replication queues" if the Replication Server cannot be reached.

12.1.31.4 sap_status route

The `sap_status route` command monitors the sequence of queues, threads, and servers that the data is transacting in the replication path.

Syntax

```
sap_status route
```

Examples

Example 1

Returns information about the queues, threads, and servers:

```
sap_status route
go
```

The returned result is:

```
sap_status route
go
PATH          SEQUENCE  NAME      TYPE      QID      SPID
SITE          STATE     BACKLOG
-----
PR.DR.master 1      ASE       S         NULL     58312
site0         Active
PR.DR.master 2      RAT       T         NULL     63
site0         Active
PR.DR.master 3      RATCI     T         NULL     NULL
site1         Active (Active)
PR.DR.master 4      SPQ       Q         106     NULL
site1         NULL
PR.DR.master 5      CAP       T         NULL     53
site1         Active (Awaiting Command)
PR.DR.master 6      SQM       T         NULL     22
site1         Active (Awaiting Message)
PR.DR.master 7      IBQ       Q         106     NULL
site1         NULL
PR.DR.master 8      SQT       T         NULL     73
site1         Active (Awaiting Wakeup)
PR.DR.master 9      DIST     T         NULL     41
site1         Active (Awaiting Wakeup)
PR.DR.master 10     SQM       T         NULL     95
site1         Active (Awaiting Message)
PR.DR.master 11     RouteQ    Q         16777319 NULL
site1         NULL
PR.DR.master 12     RSI       T         NULL     96
site1         Active (Awaiting Wakeup)
PR.DR.master 13     SQM       T         NULL     96
site2         Active (Awaiting Message)
PR.DR.master 14     OBQ       Q         116     NULL
site2         NULL
```

PR.DR.master	15	DSI	T	NULL	171
site2	Active	(Awaiting Message)	NULL		
PR.DR.master	16	ASE	S	NULL	31638
site2	Active		NULL		
PR.DR.PI2	1	ASE	S	NULL	58312
site0	Active		0		
PR.DR.PI2	2	RAT	T	NULL	64
site0	Active		NULL		
PR.DR.PI2	3	RATCI	T	NULL	NULL
site1	Active	(Active)	NULL		
PR.DR.PI2	4	SPQ	Q	110	NULL
site1	NULL		0		
PR.DR.PI2	5	CAP	T	NULL	59
site1	Active	(Awaiting Command)	NULL		
PR.DR.PI2	6	SQM	T	NULL	26
site1	Active	(Awaiting Message)	NULL		
PR.DR.PI2	7	IBQ	Q	110	NULL
site1	NULL		0		
PR.DR.PI2	8	SQT	T	NULL	71
site1	Active	(Awaiting Wakeup)	NULL		
PR.DR.PI2	9	DIST	T	NULL	43
site1	Active	(Awaiting Wakeup)	NULL		
PR.DR.PI2	10	SQM	T	NULL	95
site1	Active	(Awaiting Message)	NULL		
PR.DR.PI2	11	RouteQ	Q	16777319	NULL
site1	NULL		0		
PR.DR.PI2	12	RSI	T	NULL	96
site1	Active	(Awaiting Wakeup)	NULL		
PR.DR.PI2	13	SQM	T	NULL	115
site2	Active	(Awaiting Message)	NULL		
PR.DR.PI2	14	OBQ	Q	117	NULL
site2	NULL		0		
PR.DR.PI2	15	DSI	T	NULL	227
site2	Active	(Awaiting Message)	NULL		
PR.DR.PI2	16	ASE	S	NULL	31638
site2	Active		NULL		
PR.DR.db1	1	ASE	S	NULL	58312
site0	Active		0		
PR.DR.db1	2	RAT	T	NULL	65
site0	Active		NULL		
PR.DR.db1	3	RATCI	T	NULL	NULL
site1	Active	(Active)	NULL		
PR.DR.db1	4	SPQ	Q	114	NULL
site1	NULL		0		
PR.DR.db1	5	CAP	T	NULL	65
site1	Active	(Awaiting Command)	NULL		
PR.DR.db1	6	SQM	T	NULL	30
site1	Active	(Awaiting Message)	NULL		
PR.DR.db1	7	IBQ	Q	114	NULL
site1	NULL		0		
PR.DR.db1	8	SQT	T	NULL	69
site1	Active	(Awaiting Wakeup)	NULL		
PR.DR.db1	9	DIST	T	NULL	45
site1	Active	(Awaiting Wakeup)	NULL		
PR.DR.db1	10	SQM	T	NULL	95
site1	Active	(Awaiting Message)	NULL		
PR.DR.db1	11	RouteQ	Q	16777319	NULL
site1	NULL		0		
PR.DR.db1	12	RSI	T	NULL	96
site1	Active	(Awaiting Wakeup)	NULL		
PR.DR.db1	13	SQM	T	NULL	134
site2	Active	(Awaiting Message)	NULL		
PR.DR.db1	14	OBQ	Q	118	NULL
site2	NULL		0		
PR.DR.db1	15	DSI	T	NULL	198
site2	Active	(Awaiting Message)	NULL		
PR.DR.db1	16	ASE	S	NULL	31638
site2	Active		NULL		

PR.HA.master 1	ASE	S	NULL	58312
site0 Active		0		
PR.HA.master 2	RAT	T	NULL	63
site0 Active		NULL		
PR.HA.master 3	RATCI	T	NULL	NULL
site1 Active (Active)		NULL		
PR.HA.master 4	SPQ	Q	106	NULL
site1 NULL		0		
PR.HA.master 5	CAP	T	NULL	53
site1 Active (Awaiting Command)		NULL		
PR.HA.master 6	SQM	T	NULL	22
site1 Active (Awaiting Message)		NULL		
PR.HA.master 7	IBQ	Q	106	NULL
site1 NULL		0		
PR.HA.master 8	SQT	T	NULL	73
site1 Active (Awaiting Wakeup)		NULL		
PR.HA.master 9	DIST	T	NULL	41
site1 Active (Awaiting Wakeup)		NULL		
PR.HA.master 10	SQM	T	NULL	19
site1 Active (Awaiting Message)		NULL		
PR.HA.master 11	OBQ	Q	105	NULL
site1 NULL		0		
PR.HA.master 12	DSI	T	NULL	33
site1 Active (Awaiting Message)		NULL		
PR.HA.master 13	ASE	S	NULL	55597
site1 Active		NULL		
PR.HA.PI2 1	ASE	S	NULL	58312
site0 Active		0		
PR.HA.PI2 2	RAT	T	NULL	64
site0 Active		NULL		
PR.HA.PI2 3	RATCI	T	NULL	NULL
site1 Active (Active)		NULL		
PR.HA.PI2 4	SPQ	Q	110	NULL
site1 NULL		0		
PR.HA.PI2 5	CAP	T	NULL	59
site1 Active (Awaiting Command)		NULL		
PR.HA.PI2 6	SQM	T	NULL	26
site1 Active (Awaiting Message)		NULL		
PR.HA.PI2 7	IBQ	Q	110	NULL
site1 NULL		0		
PR.HA.PI2 8	SQT	T	NULL	71
site1 Active (Awaiting Wakeup)		NULL		
PR.HA.PI2 9	DIST	T	NULL	43
site1 Active (Awaiting Wakeup)		NULL		
PR.HA.PI2 10	SQM	T	NULL	23
site1 Active (Awaiting Message)		NULL		
PR.HA.PI2 11	OBQ	Q	109	NULL
site1 NULL		0		
PR.HA.PI2 12	DSI	T	NULL	34
site1 Active (Awaiting Message)		NULL		
PR.HA.PI2 13	ASE	S	NULL	55597
site1 Active		NULL		
PR.HA.db1 1	ASE	S	NULL	58312
site0 Active		0		
PR.HA.db1 2	RAT	T	NULL	65
site0 Active		NULL		
PR.HA.db1 3	RATCI	T	NULL	NULL
site1 Active (Active)		NULL		
PR.HA.db1 4	SPQ	Q	114	NULL
site1 NULL		0		
PR.HA.db1 5	CAP	T	NULL	65
site1 Active (Awaiting Command)		NULL		
PR.HA.db1 6	SQM	T	NULL	30
site1 Active (Awaiting Message)		NULL		
PR.HA.db1 7	IBQ	Q	114	NULL
site1 NULL		0		
PR.HA.db1 8	SQT	T	NULL	69
site1 Active (Awaiting Wakeup)		NULL		

PR.HA.db1	9	DIST	T	NULL	45
site1	Active	(Awaiting Wakeup)	NULL		
PR.HA.db1	10	SQM	T	NULL	27
site1	Active	(Awaiting Message)	NULL		
PR.HA.db1	11	OBQ	Q	113	NULL
site1	NULL		0		
PR.HA.db1	12	DSI	T	NULL	36
site1	Active	(Awaiting Message)	NULL		
PR.HA.db1	13	ASE	S	NULL	55597
site1	Active		NULL		
(87 rows affected)					

Usage

The returned information of this command is listed in the following table.

Table 12: Result Set Column Description

Information	Description
Path	The replication path.
Sequence	The order number of the current queue, thread or server in the sequence. See the <i>Result Set Row Description</i> table, below, for detailed information.
Name	The name of the queue, thread, or server.
Type	The type can be one of: <ul style="list-style-type: none"> • T - Thread • Q - Queue • S - Server
QID	The ID number of the current queue.
SPID	The ID number of the current thread or the process ID of the server.
Site	The host name of the server in which the thread or queue is located.
State	The status of the thread and server. <ul style="list-style-type: none"> • Active • Down • NULL - represents SQL<NULL>, which means the information cannot be queried.
<div style="background-color: #f0f0f0; padding: 5px; border: 1px solid #ccc;"> <p>i Note Threads also have some other specific states.</p> </div>	
Backlog	The accumulated logs to be processed. Displays 0 when there are no logs to be processed. Displays NULL when the information cannot be queried.
<div style="background-color: #f0f0f0; padding: 5px; border: 1px solid #ccc;"> <p>i Note Backlogs are only available for queues and the primary ASE, so NULL is displayed for threads and the standby ASE.</p> </div>	

Table 13: Result Set Row Description

Sequence	Name	Description
1	ASE	Primary SAP ASE.
2	RAT	Replication Agent thread - read and analyze the transaction logs of the primary SAP ASE.
3	RATCI	Replication Agent CI thread.
4	SPQ	Simple persistent queue.
5	CAP	Capture - receive information from the RAT.
6	SQM	Stable queue management - manage inbound queue.
7	IBQ	Inbound queue.
8	SQT	Stable queue transaction: Sort logs according to the commit time.
9	DIST	Distributor: Distribute logs to different route queues.
10	SQM	(Only for local distribution mode) Stable queue management - manage route queue.
11	RouteQ	(Only for local distribution mode) Route queue
12	RSI	(Only for local distribution mode) Replication Server interface - the interface between Replication Servers.
13	SQM	Stable queue management: Manage outbound queue.
14	OBQ	Outbound queue.
15	DSI	Data server interface: The interface that connects to the standby database.
16	ASE	The standby SAP ASE.

12.1.31.5 sap_status spq_agent

Displays information about the SPQ Agent.

Syntax

```
sap_status spq_agent
```

Examples

Example 1

Displays information about the SPQ Agent. In this example the participating databases are; master, PI2, and db1. Database db1 is configured for SPQ Agent with `ACTIVE` status from remote connection

PI2_RP_R2 and with INACTIVE status from remote connection PI2_HA_R1. The local connections (PI2_PR and PI2_HA) of db1 are not configured for SPQ Agent.

```
sap_status spq_agent
go
```

The returned result is:

HOST	INFO	STATE	BACKLOG
PR	108 PI2_HA_R1.db1	INACTIVE	0
PR	107 PI2_PR.db1	INACTIVE	NA
HA	109 PI2_HA.db1	INACTIVE	NA
HA	110 PI2_PR_R2.db1	ACTIVE	0

Usage

This table lists the information the command returns.

Table 14: Return Information for sap_status spq_agent

Information	Description
Host	Logical host name.
Info	Information about the SPQ.
State	The state of the SPQ, which can be either active or inactive: <ul style="list-style-type: none"> Active: Indicates that the SPQ Agent, is configured and the external replication is functional on this path. Inactive: Indicates that the external replication is not functional on this path.
Backlog	The size of the SPQ agent backlog. If the backlog size is not available, it indicates that the SPQ Agent (external replication) is not configured on this path.

12.1.31.6 sap_status synchronization

Use the `sap_status synchronization` command to monitor information on the replication modes you have configured, and the current replication states in the HADR environment.

Syntax

```
sap_status synchronization
```

Examples

Example 1

Monitors and returns the information on the replication modes you have configured, and the current replication states in the HADR environment:

```
sap_status synchronization
go
```

The returned result is:

```
PATH          NAME          VALUE
INFO
-----
PR           Hostname          site0          Logical host
name.
PR           HADR Status      Primary : Active Identify the primary and
standby sites.
PR           Synchronization Mode Synchronous    The configured
Synchronization Mode value.
PR           Synchronization State Synchronous    Synchronization Mode in
which replication is currently operating.
PR.master   Synchronization State Synchronous    Synchronization Mode in
which replication is currently operating.
PR.db1      Synchronization State Synchronous    Synchronization Mode in
which replication is currently operating.
PR.PI2      Synchronization State Synchronous    Synchronization Mode in
which replication is currently operating.
(7 rows affected)
```

Usage

The result information of this command is listed in the following table.

Table 15: Result Information for `sap_status synchronization`

Information	Description
Hostname	Primary logical host name.
HADR Status	HADR mode: HADR state or Unknown.
Synchronization Mode	The replication synchronization mode you have configured between a database and the SAP Replication Server, which can be one of: <ul style="list-style-type: none">• Synchronous• Asynchronous

Information	Description
Synchronization State	The current replication synchronization mode between a database and the SAP Replication Server, which can be different from the mode you have configured. Also provides the synchronization state for each database.

i Note

The synchronization state returned by the `sap_status synchronization` command represents the state of all databases that are replicated by the primary site. If the synchronization state of the different databases is not the same (for example, if one database is in the 'synchronous' state and another is in the 'asynchronous' state), the result displayed by the `sap_status synchronization` command for the site is 'Inconsistent' — indicating the databases do not all have the same synchronization state at this time.

12.1.32 sap_suspend_component

Use the `sap_suspend_component` command to suspend components on a replication path.

Syntax

```
sap_suspend_component <cpnt_name> <src_hostname> <tgt_hostname> <db_name>
```

Parameters

<cpnt_name>

Specifies the name of the component to be suspended.

<src_hostname>

Specifies the source host name.

<tgt_hostname>

Specifies the target host name.

<db_name>

Specifies the database name.

Examples

Example 1

Suspends the active RAT component from PR to HA on database master:

```
sap_suspend_component RAT PR HA master
```

The returned result is:

TASKNAME	TYPE	VALUE
Suspending Rep Component.	Start Time	Fri Feb 10 14:02:23 CST 2017
Suspending Rep Component.	Elapsed Time	00:00:06
SuspendComponent	Task Name	Suspending Rep Component.
SuspendComponent	Task State	Completed
SuspendComponent	Short Description	Suspending component on replication path.
SuspendComponent	Long Description	Suspending component 'RAT' on replication path from 'PR.HA'.
SuspendComponent	Task Start	Fri Feb 10 14:02:23 CST 2017
SuspendComponent	Task End	Fri Feb 10 14:02:29 CST 2017
SuspendComponent	Hostname	site1

(9 rows affected)

Execute the sap_status route command:

```
sap_status route
```

The returned result is:

"PR.HA.db1	2	RAT	T	NULL	NULL
site0	Down		NULL"		
"PR.HA.db1	3	RATCI	T	NULL	NULL
site1	Inactive (Down)		NULL"		

(This command directly takes RAT down and indirectly makes RATCI inactive)

Example 2

Suspends the inactive RAT component from site "PR" to "HA" on database "master":

```
sap_suspend_component RAT PR HA master
```

The returned result is:

TASKNAME	TYPE	VALUE
Suspending Rep Component.	Start Time	Tue Feb 14 18:53:17 CST 2017
Suspending Rep Component.	Elapsed Time	00:00:00
SuspendComponent	Task Name	Suspending Rep Component.
SuspendComponent	Task State	Completed
SuspendComponent	Short Description	Suspending component on replication path.
SuspendComponent	Long Description	Suspending component 'RAT' on replication path from 'PR.HA'.
SuspendComponent	Task Start	Tue Feb 14 18:53:17 CST 2017
SuspendComponent	Task End	Tue Feb 14 18:53:17 CST 2017
SuspendComponent	Hostname	site1

(As RAT is down, this command returns immediately)

Example 3

Suspends all components from site "PR" to "HA" on database "master".

```
sap_suspend_component ALL PR HA master
```

The returned result is:

TASKNAME	TYPE	VALUE
Suspending Rep Component.	Start Time	Tue Feb 14 18:53:17 CST 2017
Suspending Rep Component.	Elapsed Time	00:00:00
SuspendComponent	Task Name	Suspending Rep Component.
SuspendComponent	Task State	Completed
SuspendComponent	Short Description	Suspending component on replication path.
SuspendComponent	Long Description	Suspending component 'ALL' on replication path from 'PR,HA'.
SuspendComponent	Task Start	Tue Feb 14 18:53:17 CST 2017
SuspendComponent	Task End	Tue Feb 14 18:53:17 CST 2017
SuspendComponent	Hostname	site1

(9 rows affected)

Example 4

Suspends the "CAP" component from site "HA" to "PR" on database "master".

```
sap_suspend_component CAP HA PR master
```

The returned result is:

TASKNAME	TYPE	VALUE
Suspending Rep Component.	Start Time	Mon Mar 13 13:48:19 CST 2017
Suspending Rep Component.	Elapsed Time	00:00:00
SuspendComponent	Task Name	Suspending Rep Component.
SuspendComponent	Task State	Error
SuspendComponent	Short Description	Suspend component on certain replication path.
SuspendComponent	Long Description	Validating arguments.
SuspendComponent	Failing Command	sap_suspend_component CAP, HA, PR, master,
SuspendComponent	Additional Info	Source host 'HA' is not a primary site, please check your input.
SuspendComponent	Corrective Action	sap_suspend_component
SuspendComponent	component_name src_host_logical_name	tgt_host_logical_name database_name
SuspendComponent	Task Start	Mon Mar 13 13:48:19 CST 2017
SuspendComponent	Task End	Mon Mar 13 13:48:19 CST 2017
SuspendComponent	Hostname	site1

(12 rows affected)

Usage

The valid values for the cpnt parameter are:

- RAT
- RATCI
- CAP
- DIST
- RSI
- DSI
- ALL

Set the `cpnt` parameter to `ALL` to suspend all the supported components.

i Note

When you set the `<cpnt>` parameter to `ALL`, RMA initiates a batch job to suspend all the supported components. However, the `RSI` component is not suspended as suspending the `RSI` component on one path may affect some other too.

12.1.33 sap_suspend_replication

Use the `sap_suspend_replication` command to suspend replication to a specified database or all databases that are in the participating databases list (master and ERP).

Syntax

To suspend replication to a specified database:

```
sap_suspend_replication <standby_logical_host_name> {all, [, <database_name>]}
go
```

Parameters

<standby_logical_host_name>

Specifies the standby logical host name.

all

Suspends replication to all SAP databases.

<database_name>

Suspends replication to the specified database.

12.1.34 sap_teardown

Tearing down a replication environment includes disabling replication in the SAP ASE servers, stopping the SAP Replication Servers, and deleting all directories and files created during setup, including the SAP Replication Server instances.

Use the `sap_teardown` command to tear down the replication environment. The command does not modify any data that has been replicated to the standby databases. Additionally, the databases on both the primary and standby hosts are marked for replication. The command does not remove any software, but it does remove SAP Replication Servers and the configurations that support replication.

Syntax

```
sap_teardown
```

Examples

Example 1

Tears down the replication environment:

```
sap_teardown
go
```

The returned result is:

```
TASKNAME          TYPE
VALUE
-----
-----
Tear Down Replication Start Time      Mon Nov 23 05:24:25 EST
2015
Tear Down Replication Elapsed Time
00:00:22
TearDownRS          Task Name      Tear Down
Replication
TearDownRS          Task State
Completed
TearDownRS          Short Description Tear down the Replication
Environment
TearDownRS          Long Description Tear Down of the Replication
environment is complete.
TearDownRS          Task Start      Mon Nov 23 05:24:25 EST
2015
TearDownRS          Task End        Mon Nov 23 05:24:47 EST
2015
TearDownRS          Hostname
site1
Tear Down Replication Start Time      Mon Nov 23 05:23:41 EST
2015
Tear Down Replication Elapsed Time
00:00:44
TearDownRS          Task Name      Tear Down
Replication
```

```

TearDownRS          Task State
Completed
TearDownRS          Short Description Tear down the Replication
Environment
TearDownRS          Long Description  Tear Down of the Replication
environment is complete.
TearDownRS          Task Start      Mon Nov 23 05:23:41 EST
2015
TearDownRS          Task End      Mon Nov 23 05:24:25 EST
2015
TearDownRS          Hostname
site0
(18 rows affected)

```

Usage

The `sap_tearardown` command performs the following tasks:

- Stops the Replication Server and deletes its instance directory, partition files, simple persistent queue directories, and kills all Replication Server related processes.
- Demotes the source SAP ASE, if the source host (the machine on which SAP ASE runs) is available.
- Drops all servers from the HADR server list on both SAP ASE servers.
- Drops the HADR group from both servers.
- Disables HADR on both servers.
- Disables CIS RPC Handling.

12.1.35 sap_tune_rat

Tunes and configures the Replication Agent thread for SAP ASE (RepAgent for short) in an HADR environment.

Syntax

```
sap_tune_rat {<database> | all}, <memory_limit>
```

Parameters

<database> | all

Specifies the name of the database to be tuned. Specifies `all` for RMA to tune all the participating databases with the same memory input.

<memory_limit>

Specifies the memory limit in SAP Replication Server that the RepAgent connects to. The valid range is 4 GB to 256 GB.

Examples

Example 1

Execute the command to tune the RepAgent that connects to an SAP Replication Server, with the memory limit set to 8:

```
sap_tune_rat PI2, 8
go
```

The returned result is:

```
TASKNAME                                TYPE                                VALUE
-----                                -
Tune Replication Agent Thread Start Time    Wed Sep 21 01:07:09 EDT 2016
Tune Replication Agent Thread Elapsed Time    00:00:22
TuneRAT                                  Task Name      Tune Replication Agent Thread
TuneRAT                                  Task State     Completed
TuneRAT                                  Short Description Tune Tune Replication Agent
Thread related configurations.
TuneRAT                                  Task Start     Wed Sep 21 01:07:09 EDT 2016
TuneRAT                                  Task End       Wed Sep 21 01:07:31 EDT 2016
TuneRAT                                  Hostname       site0
(8 rows affected)
```

Usage

Although RMA automatically runs `sap_tune_rat` when you add a database in the HADR system, you can manually configure the parameters if the automatic execution of the command fails.

To view the configuration parameters, use the `admin config Replication Server` command. For more details, see the *Replication Server Reference Manual*.

Table 16: ASE Configuration Parameters for Both the Primary and Standby ASE that Configured RAT, Within an HA Cluster (Primary and Standby)

Parameter Name	Value (User databases) (ASE command <code>sp_config_rep_agent <user db>...</code>)	Value (master databases) (ASE command <code>sp_config_rep_agent master...</code>)
<code>max commands per package</code>	100	100
<code>stream buffer size</code>	1572864	1572864
<code>buffer pool size</code>	100* <memory> / 4 For Windows: 50* <memory> / 4	8

Parameter Name	Value (User databases) (ASE command <code>sp_config_rep_agent</code> <user db>...)	Value (master databases) (ASE command <code>sp_config_rep_agent</code> master...)
<code>replication agent memory</code>	$51200 + (<USER_DB_1><stream\ buffer\ size> * <buffer\ pool\ size>) + <USER_DB_2><stream\ buffer\ size> * <buffer\ pool\ size> + \dots + <USER_DB_N><stream\ buffer\ size> * <buffer\ pool\ size>) / 2048$ <USER_DB_N> refers to all the participating databases in the HADR system, excluding the master database.	
	<div style="border-left: 2px solid #0070C0; padding-left: 10px;"> <p>i Note</p> <p>RepAgent memory is tuned only when the current memory is less than the calculated value. If the current memory is larger than the calculated value, RepAgent does not change its memory.</p> </div>	
<code>max memory</code>	The <code>max memory</code> value varies according to the <code>replication agent memory</code> value. If the <code>replication agent memory</code> value decreases, the value of <code>max memory</code> remains the same. If the <code>replication agent memory</code> increases, additional memory is added to <code>max memory</code> .	

12.1.36 sap_tune_rs

Specifies the number of the CPUs and the maximum size of the memory to tune the Replication Server instance after the HADR replication system is set up and Replication Server instances have been created on all the hosts.

Syntax

```
sap_tune_rs <logical_host_name>, <memory_limit>, <cpu_number> {,display | apply_only | apply_restart}
```

Parameters

<logical_host_name>

Specifies the logical host name of the primary Replication Server to be configured.

<memory_limit>

Specifies the available memory limit.

<cpu_number>

Specifies the available number of CPUs.

display

Displays the configurations to be applied, and the status of the configurations.

apply_only

Applies the configuration without restarting the Replication Server.

apply_restart

Applies the configuratio after restarting the Replication Server.

Examples

Example 1

Tunes SAP Replication Server on logical host site1 with 4 GB of memory and 2 CPUs:

```
sap_tune_rs site1,4,2
go
TASKNAME TYPE VALUE
-----
Tune Replication Server Start Time Wed Apr 01 05:06:05 EDT 2015
Tune Replication Server Elapsed Time 00:00:33
TuneRS Task Name Tune Replication Server
TuneRS Task State Completed
TuneRS Short Description Tune Replication Server configurations.
TuneRS Task Start Wed Apr 01 05:06:05 EDT 2015
TuneRS Task End Wed Apr 01 05:06:38 EDT 2015
TuneRS Hostname site0
```

Example 2

Displays the configurations applied on the SAP Replication Server on logical host site1 with 4 GB of memory and 2 CPUs:

```
sap_tune_rs site1,4,2,display
go
LOGICAL_HOST_NAME RSNAME CONFIG_LEVEL PROPNAME CONFIG_VALUE RUN_VALUE
DEFAULT_VALUE RESTART_REQ
-----
site1 PI2_REP_site1 site1 memory_limit 4096 4096 2047 Restart not required
site1 PI2_REP_site1 site1 dsi_cdb_max_size 341 341 1024 Connection/route
restart required
site1 PI2_REP_site1 site1 sqm_cmd_cache_size 402653184 402653184 20971520
Restart required
site1 PI2_REP_site1 site1 sqm_max_cmd_in_block 1024 1024 320 Restart required
site1 PI2_REP_site1 site1 sqt_max_cache_size 400000000 400000000 20971520
Restart not required
site1 PI2_REP_site1 105:0 PI2 sqm_cache_size 128 128 16 Restart required
site1 PI2_REP_site1 105:1 PI2 sqm_cache_size 128 128 16 Restart required
site1 PI2_REP_site1 106:0 PI2 sqm_cache_size 128 128 16 Restart required
site1 PI2_REP_site1 106:1 PI2 sqm_cache_size 128 128 16 Restart required
```

Usage

- If the `sap_tune_rs` command fails, you can manually configure the parameters. To view the configuration values, use the `admin config Replication Server` command. For more details, refer to the *Replication Server Reference Manual*.

To manually configure the parameter values, refer to the formulas mentioned in the below table. If your calculated value exceeds the upper limit of the parameter, the command applies the upper limit value to the configuration.

Table 17: Server Level Configuration Values

Parameter Name	Server Level Value (SRS command <code>configure replication server...</code>)
<code>memory_limit</code>	<code><memory> * 1024</code>

- (Only for Linux with `numactl` command) The `<cpu_number>` parameter affects the Replication Server instance run file by binding the Replication Server processes running on specific CPUs with the `numactl` system command. The Replication Server run file with the appending `numactl` command is saved in the `$SYBASE/<CID_database_name>_REP_<logical_host_name>/RUN_<CID_database_name>_REP_<logical_host_name>.sh` directory. SAP Replication Server run file without the appending `numactl` command is saved as a backup in the same directory with the name `RUN_<CID_database_name>_REP_<logical_host_name>.sh.prev`. The RMA only backs up the run scripts in which there are no lines appended with CPU number `numactl`. For example, the `RUN_PI2_REP_site1.sh` run script is updated if the RS instance directory is `/sybase/DM/PI2_REP_site1`.

```
Old run script content:
# Environment variables
# .....
#
# Runserver file for Replication Server 'PI2_REP_site0'. Created by rs_init.
#
/sybase/DM/REP-16_0/bin/repserver -SPI2_REP_site0 \
-C/sybase/DM/PI2_REP_site0/PI2_REP_site0.cfg \
-E/sybase/DM/PI2_REP_site0/PI2_REP_site0.log -I/sybase/DM/interfaces
Updated run script content:
# Environment variables
# .....
#
# Runserver file for Replication Server 'PI2_REP_site0'. Created by rs_init.
#
echo 'Binding repserver process on CPU 0-1'
<USER_BIN_DIRECTORY>/numactl --physcpubind=0-1 /sybase/DM/REP-16_0/bin/
repserver -SPI2_REP_site0 \
-C/sybase/DM/PI2_REP_site0/PI2_REP_site0.cfg \
-E/sybase/DM/PI2_REP_site0/PI2_REP_site0.log -I/sybase/DM/interfaces
// If the cpu number is x, numactl --physcpubind=0-<x-1> is appended. For
example for cpu number = 16, numactl --physcpubind=0-15 is appended
```

By default, the Replication Server is bound with CPUs starting with `CPU0`. To bind the SAP Replication Server to other CPUs, modify the `<USER_BIN_DIRECTORY>/numactl --physcpubind=0-1` line of the run script.

The RMA does not modify the Replication Server run file or create a backup of the `.prev` file if the `numactl` command is not found under the following `<USER_BIN_DIRECTORY>` directories:

- /usr/bin
- /usr/local/bin
- /usr/sbin
- /bin
- /sbin

12.1.37 sap_update_replication

Use the `sap_update_replication` command to adjust the replication system at runtime.

Syntax

```
sap_update_replication add_db <database_name> | remove_db <database_name> |
add_additional_device, <logical_host_name>, <device_directory>, <device_size> |
synchronization_mode <logical_host_name>, {async | sync | ltl} |
distribution_mode <logicalHost>, {local | remote} [, <remoteTargetHost> ]
```

Parameters

add_db

Adds the specified database to the HADR replication system.

<database_name>

Specifies the name of the database.

remove_db

Removes the specified database from the HADR replication system.

add_additional_device

Indicates that you are adding a device to the HADR system.

<logical_host_name>

Name of the host on which you are adding the device.

<device_directory>

Path to the directory in which you are creating the device.

<device_size>

Size of the device.

12.1.37.1 sap_update_replication add_db | remove_db

Use the `sap_update_replication add_db` command to add or remove a database from the HADR replication system.

Syntax

```
sap_update_replication add_db <database_name>| remove_db <database_name>
```

Parameters

<database_name>

The name of the database to be added or removed.

add_db

Adds a database to the HADR replication system.

remove_db

Removes a database from the HADR replication system.

Examples

Example 1

Removes database PI2 from the HADR system:

```
sap_update_replication remove_db,PI2
go
```

The returned result is:

```
TASKNAME          TYPE
                VALUE
-----
-----
Update Replication Start Time
                   Fri Nov 20 02:18:17 EST 2015
Update Replication Elapsed Time
                   00:00:02
DRExecutorImpl    Task Name
                   Update Replication
DRExecutorImpl    Task State
                   Running
DRExecutorImpl    Short Description
                   Update configuration for a currently replicating site.
DRExecutorImpl    Long Description
```



```

        Started task 'Update Replication' asynchronously.
DRExecutorImpl      Additional Info
        Please execute command 'sap_status task' to determine when
task 'Update
        Replication' is complete.
UpdateReplication  Task Name
        Update Replication
UpdateReplication  Task State
        Running
UpdateReplication  Short Description
        Update configuration for a currently replicating site.
UpdateReplication  Long Description
        Disabling the incoming replication data for database 'PI2'.
UpdateReplication  Current Task Number
        2
UpdateReplication  Total Number of Tasks
        10
UpdateReplication  Task Start
        Fri Nov 20 02:18:17 EST 2015
UpdateReplication  Hostname
        site0
(15 rows affected)

```

Checks the status of the task:

```

sap_status task
go

```

The returned result is:

```

TASKNAME          TYPE
                VALUE
-----
-----
--
        Status          Start Time
        Fri Nov 20 02:18:17 EST 2015
        Status          Elapsed Time
        00:04:58
UpdateReplication Task Name
        Update Replication
UpdateReplication Task State
        Completed
UpdateReplication Short Description
        Update configuration for a currently replicating site.
UpdateReplication Long Description
        Update replication request to remove database 'PI2'
completed succ
        essfully.
UpdateReplication Current Task Number
        10
UpdateReplication Total Number of Tasks
        10
UpdateReplication Task Start
        Fri Nov 20 02:18:17 EST 2015
UpdateReplication Task End
        Fri Nov 20 02:23:15 EST 2015
UpdateReplication Hostname
        site0
(11 rows affected)

```

Finds the error message:

```

sap_update_replication remove_db,PI2

```

```
go
```

The returned result is:

```

TASKNAME          TYPE
          VALUE
-----
-----
Update Replication Start Time
          Fri Nov 20 02:24:52 EST 2015
Update Replication Elapsed Time
          00:00:00
UpdateReplication Task Name
          Update Replication
UpdateReplication Task State
          Error
UpdateReplication Short Description
          Update configuration for a currently replicating site.
UpdateReplication Failing Command Error Message
          The database 'PI2' does not exist in the HADR system, so
no change
          was performed.
UpdateReplication Task Start
          Fri Nov 20 02:24:52 EST 2015
UpdateReplication Task End
          Fri Nov 20 02:24:52 EST 2015
UpdateReplication Hostname
          site0
(9 rows affected)

```

Example 2

Adds database PI2 into the HADR system:

```
sap_update_replication add_db,PI2
go
```

The returned result is:

```

TASKNAME          TYPE
          VALUE
-----
-----
Update Replication Start Time
          Fri Nov 20 02:26:11 EST 2015
Update Replication Elapsed Time
          00:00:02
DRExecutorImpl    Task Name
          Update Replication
DRExecutorImpl    Task State
          Running
DRExecutorImpl    Short Description
          Update configuration for a currently replicating site.
DRExecutorImpl    Long Description
          Started task 'Update Replication' asynchronously.
DRExecutorImpl    Additional Info
          Please execute command 'sap_status task' to determine
when task 'Update
          Replication' is complete.
UpdateReplication Task Name
          Update Replication
UpdateReplication Task State

```

```

Running
UpdateReplication Short Description
Update configuration for a currently replicating site.
UpdateReplication Long Description
Add database 'PI2' into Replication Servers.
UpdateReplication Current Task Number
1
UpdateReplication Total Number of Tasks
3
UpdateReplication Task Start
Fri Nov 20 02:26:11 EST 2015
UpdateReplication Hostname
site0
AddASEDatabaseForDisasterRecovery Task Name
Add Database to Replication
AddASEDatabaseForDisasterRecovery Task State
Running
AddASEDatabaseForDisasterRecovery Short Description
Add an ASE database to the Replication System for
Disaster Recovery sup
port.
AddASEDatabaseForDisasterRecovery Long Description
Adding database 'PI2' to the replication server on host
'site0'.
AddASEDatabaseForDisasterRecovery Current Task Number
0
AddASEDatabaseForDisasterRecovery Total Number of Tasks
20
AddASEDatabaseForDisasterRecovery Hostname
site0
AddLocalASEDatabaseForDisasterRecovery Task Name
Add Local Database to Replication
AddLocalASEDatabaseForDisasterRecovery Task State
Running
AddLocalASEDatabaseForDisasterRecovery Short Description
Add an ASE database to the local Replication System for
Disaster Recove
ry support.
AddLocalASEDatabaseForDisasterRecovery Long Description
Adding database 'PI2' to the replication server on host
'site0'.
AddLocalASEDatabaseForDisasterRecovery Hostname
site0
ResumeDSI Task Name
Resume DSI Connection
ResumeDSI Task State
Running
ResumeDSI Short Description
Resume the Replication Server connection to a database.
ResumeDSI Long Description
Waiting 1 seconds: Waiting 10 seconds before checking if
the command ex
ecuted successfully.
ResumeDSI Hostname
site0
(32 rows affected)

```

Usage

- After a new database is added to the HADR replication system, use the `sap_materialize` command to materialize the database with the current primary database.
- After a database is removed from the HADR replication system, the replication path between them is uninstalled and no synchronization happens.

- This command cannot add or remove SAP ASE system databases, including the master database, which is needed for database logins and user information synchronization.
- This command cannot add or remove SAP CID databases.

12.1.37.2 sap_update_replication distribution_mode

Use the `sap_update_replication distribution_mode` command to change the distribution mode for the logical host.

Syntax

```
sap_update_replication distribution_mode <source_logical_host_name> {local |
remote, <target_logical_host_name>}
```

Parameters

<source_logical_host_name>

Specifies the logical host name of the source.

local

Changes the distribution mode of the site to local.

remote

Changes the distribution mode of the site to remote.

<target_logical_host_name>

Specifies the logical host name of the target.

Examples

Example 1

Changes the distribution mode of site0 to remote:

```
sap_update_replication distribution_mode,site0,local
go
```

The returned result is:

TASKNAME	TYPE	VALUE
-----	-----	-----
-----	-----	-----

```

Update Replication Start Time
      Fri Nov 20 02:56:53 EST 2015
Update Replication Elapsed Time
      00:00:33
UpdateReplication Task Name
      Update Replication
UpdateReplication Task State
      Completed
UpdateReplication Short Description
      Update configuration for a currently replicating site.
UpdateReplication Long Description
      Completed changing host 'site0' to distribution mode of
'local'.
UpdateReplication Current Task Number
      16
UpdateReplication Total Number of Tasks
      17
UpdateReplication Task Start
      Fri Nov 20 02:56:53 EST 2015
UpdateReplication Task End
      Fri Nov 20 02:57:26 EST 2015
UpdateReplication Hostname
      site0

```

Example 2

Finds the error from row where TYPE value is Failing Command Error Message:

```

sap_update_replication distribution_mode, site2, local
go

```

The returned result is:

```

TASKNAME          TYPE
VALUE
-----
-----
-----
-----
-----
-----
-----
Update Replication Start Time
      Fri Nov 20 03:09:36 EST
2015

Update Replication Elapsed Time
00:00:00

UpdateReplication Task Name
      Update
Replication

UpdateReplication Task State
Error

```

```

UpdateReplication Short Description
Update configuration for a currently replicating
site.

UpdateReplication Failing Command
sap_update_replication distribution_mode, site2,
local,

UpdateReplication Failing Command Error Message
The logical name 'site2' is not
valid.

UpdateReplication Corrective Action
sap_update_replication add <logicalHost> | remove
<logicalHost> | add_d
b <DBName>, [force] | remove_db <DBName>,
[force] | add_additional_devi
ce <LogicalHost>, <additional device dir>,
<additional device size> | sy
nchronization_mode <logicalHost>, [async | sync |
ltl ] | di
istribution_mode <logicalHost> [ local | remote,
<remoteTargetHost> ]
UpdateReplication Task Start
Fri Nov 20 03:09:36 EST
2015

UpdateReplication Task End
Fri Nov 20 03:09:36 EST
2015

UpdateReplication Hostname
site0

(11 rows affected)

```

12.1.37.3 sap_update_replication synchronization_mode

Use the `sap_update_replication synchronization_mode` command to change the HADR system synchronization mode.

Syntax

Variable declaration:

```
synchronization_mode <logical_host_name>, {async | sync | ltl}
```

Parameters

`<logical_host_name>`

Specifies the name of the host on which you are changing the synchronization mode.

Examples

Example 1

Changes the synchronization mode of `site0` to `sync`:

```
sap_update_replication synchronization_mode, site0, sync
```

The returned result is:

```
TASKNAME          TYPE
                VALUE
-----
-----
Update Replication Start Time
                Fri Nov 20 03:24:52 EST 2015
Update Replication Elapsed Time
                00:00:22
UpdateReplication Task Name
                Update Replication
UpdateReplication Task State
                Completed
UpdateReplication Short Description
                Update configuration for a currently replicating site.
UpdateReplication Long Description
                Successfully submitted the design changes for local host
'site0' for se
                rver on host 'null'.
UpdateReplication Task Start
                Fri Nov 20 03:24:52 EST 2015
UpdateReplication Task End
```

```
Fri Nov 20 03:25:14 EST 2015
UpdateReplication Hostname
                  site0
(9 rows affected)
```

12.1.37.4 sap_update_replication add_additional_device

Use the `sap_update_replication add_additional_device` command to add additional devices to the HADR system. You can also specify the logical device name, the device file name, and the size using the `sap_update_replication add_additional_device` command.

Syntax

```
sap_update_replication add_additional_device <logical_host_name>,
<logical_device_name>, <device_file>, <device_size>
```

Parameters

`add_additional_device`

Indicates that you are adding a device to the HADR system.

`<logical_host_name>`

Specifies the logical host name of the SAP Replication Server on which the device space is to be added.

`<logical_device_name>`

Specifies the logical name of the device to be added.

`<device_file>`

Specifies the path to the directory in which you are creating the device and the name of the device file.

`<device_size>`

Specifies the size of the device.

Examples

Example 1

Add 20 MB device space named `part9` on `site0` SAP Replication Server under `/testenv7/partition9.dat`:

```
sap_update_replication add_additional_device site0, part9, /testenv7/
partition9.dat, 20
go
```

The returned result is:

TASKNAME	TYPE	VALUE
Add Replication Server Device	Start Time	Thu May 16 05:11:42 UTC 2019
Add Replication Server Device	Elapsed Time	00:00:01
AddDevice	Task Name	Add Replication Server Device
AddDevice	Task State	Completed
AddDevice	Short Description	Add Replication Server Device with the specified size.
AddDevice	Long Description	Successfully added device 'part9' (file '/testenv7/partition9.dat') with size 20 on site0 host 'rmawqiusitel.mo.sap.corp:14975'.
AddDevice	Task Start	Thu May 16 05:11:42 UTC 2019
AddDevice	Task End	Thu May 16 05:11:43 UTC 2019
AddDevice	Hostname	rmawqiusitel.mo.sap.corp
(9 rows affected)		

Example 2

If the command fails, find the error from the row where the `TYPE` value is `Failing Command Error` Message:

```
sap_update_replication add_additional_device site0, part3, /testenv7/
partition3.dat, 9
go
```

The returned result is:

TASKNAME	TYPE	VALUE
Add Replication Server Device	Start Time	Thu May 16 02:22:27 UTC 2019
Add Replication Server Device	Elapsed Time	00:00:00
AddDevice	Task Name	Add Replication Server Device
AddDevice	Task State	Error
AddDevice	Short Description	Add Replication Server Device with the specified size.
AddDevice	Long Description	Adding device 'part3' (file '/testenv7/partition3.dat') with size 9 on site0 host 'rmawqiusitel.mo.sap.corp:14975'.
AddDevice	Failing Command Error Message	Device logical name 'part3' already exists on site0 host 'rmawqiusitel.mo.sap.corp:14975'.
AddDevice	Task Start	Thu May 16 02:22:27 UTC 2019
AddDevice	Task End	Thu May 16 02:22:27 UTC 2019

```
AddDevice      Hostname
rmawqiusitel.mo.sap.corp
(10 rows affected)
```

Related Information

[sap_add_device \[page 444\]](#)

12.1.38 sap_upgrade_server

The `sap_upgrade_server RS` and `sap_upgrade_server ASE` commands are used while you upgrade the Replication Server and SAP ASE respectively.

Syntax

```
sap_upgrade_server {SRS|ASE}, {start|finish}, <upgrade_logical_hostname>
[,suspend]
```

i Note

See the *Configuration Guide* for information about upgrading Replication Server.

Parameters

SRS

Upgrades the Replication Server.

i Note

The Replication Server runserver file, `RUN_rs instance name.sh`, is regenerated during an upgrade and any user changes to this file are lost. If your site requires these changes, edit the runserver file after the upgrade is complete and restart Replication Server to make the environment settings take effect.

ASE

Upgrades the SAP ASE server.

start

Starts the upgrade. Execute this command before you install the new server release.

finish

Finishes the upgrade. Execute this command after you install the new server release.

<upgrade_logical_hostname>

Specifies the logical host name where the server to be upgraded is present.

suspend

Indicates the replication path has been suspended. No ticket is sent to verify the replication status during the upgrade process.

Examples

Example 1

Upgrades the standby SAP Replication Server:

1. At the RMA located at the same site as the standby Replication Server, execute:

```
sap_upgrade_server SRS, start, standby_logical_host_name  
go
```

2. Shut down the RMA for the site:

```
Shutdown  
go
```

3. Install the new Replication Server release.
4. Start the new RMA and execute:

```
sap_upgrade_server SRS, finish, standby_logical_host_name  
go
```

Example 2

Upgrades the standby SAP ASE:

Upgrades the standby SAP Replication Server:

1. At the RMA located at the same site as the standby SAP ASE server, execute:

```
sap_upgrade_server ASE, start, standby_logical_host_name  
go
```

2. Shut down the SAP ASE server and Backup Server for the site:

```
Shutdown  
go
```

3. Install the new SAP ASE release and upgrade the site1 SAP ASE including starting the upgraded data server and Backup Server, and running the post-installation tasks.
4. Start the new RMA and execute:

```
sap_upgrade_server ASE, finish, standby_logical_host_name  
go
```

12.1.39 sap_verify_replication

Use the `sap_verify_replication` command to verify if you can change the synchronization mode or the distribution mode of a running HADR system. After the check is passed, you can execute the `sap_update_replication` command to change the synchronization or the distribution mode of the HADR system.

Syntax

```
sap_verify_replication synchronization_mode <logical_host>, [async | sync | ltl] | distribution_mode <logical_host> [ local | remote <remote_logical_host> ]
```

Parameters

<logical_host>

Specifies the logical host name to be changed.

sync

Specifies that the transaction commit is blocked until the log records of the transaction are received and stored persistently in the standby memory.

async

Specifies that the transaction commit is blocked until the log records of the transaction are received in the standby memory.

ltl

Specifies that the transaction is transferred by the Log Transfer Language.

<remote_logical_host>

Specifies the remote logical host name when changing the `distribution_mode` to remote.

Examples

Example 1

Changes the site0 synchronization mode to `sync`:

```
sap_verify_replication synchronization_mode site0,sync  
go
```

The returned result is:

TASKNAME	TYPE VALUE
----------	---------------

```

-----
Verify Replication Start Time
Sun Nov 15 22:28:13 EST 2015
Verify Replication Elapsed Time
00:00:00
VerifyReplication Task Name
Verify Replication
VerifyReplication Task State
Completed
VerifyReplication Short Description
Verify configuration for a currently replicating site.
VerifyReplication Task Start
Sun Nov 15 22:28:13 EST 2015
VerifyReplication Task End
Sun Nov 15 22:28:13 EST 2015
VerifyReplication Hostname
site0
(8 rows affected)

```

If the command fails, finds the error from the row where TYPE value is Failing Command Error Message:

```

sap_verify_replication synchronization_mode site0, sync
go

```

The returned result is:

```

TASKNAME          TYPE
VALUE
-----
-----
-----
Verify Replication Start Time
2015              Mon Nov 16 04:34:49 EST
Verify Replication Elapsed Time
00:00:00
VerifyReplication Task Name
Replication      Verify
VerifyReplication Task State
Error
VerifyReplication Short Description
site.            Verify configuration for a currently replicating
VerifyReplication Failing Command
sync,            sap_verify_replication synchronization_mode, site0,
VerifyReplication Failing Command Error Message
valid.          The logical name 'site0' is not
VerifyReplication Corrective Action

```

```

sap_verify_replication add <logicalHost> | remove
<logicalHost> | synch
                    ronization_mode <logicalHost>, [async | sync | ltl ] | distr
                    istribution_mode <logicalHost> [ local | remote,
<remoteTargetHost> ]
VerifyReplication  Task Start
                    Mon Nov 16 04:34:49 EST
2015
VerifyReplication  Task End
                    Mon Nov 16 04:34:49 EST
2015
VerifyReplication  Hostname
site0
(11 rows affected)

```

Example 2

Changes the site0 distribution mode to `remote` and specifies the remote target as `site1`:

```

sap_verify_replication distribution_mode,site0,remote,site1
go

```

The returned result is:

```

TASKNAME          TYPE
                VALUE
-----
Verify Replication Start Time
                    Sun Nov 15 22:27:15 EST 2015
Verify Replication Elapsed Time
                    00:00:02
VerifyReplication Task Name
                    Verify Replication
VerifyReplication Task State
                    Completed
VerifyReplication Short Description
                    Verify configuration for a currently replicating site.
VerifyReplication Task Start
                    Sun Nov 15 22:27:15 EST 2015
VerifyReplication Task End
                    Sun Nov 15 22:27:17 EST 2015
VerifyReplication Hostname
                    site0
(8 rows affected)

```

12.1.40 sap_version

Use the `sap_version` command to display the version of the DR agent plug-in. The `all` keyword displays a list of all the servers known and usable by the DR Agent in this replication environment and their versions.

Syntax

```
sap_version [all]
```

12.2 SAP ASE Commands

Although you should use the SAP ASE Cockpit and the RMA, you may need to occasionally use `sp_hadr_admin` to manage and monitor the HADR system from the command line.

12.2.1 Use Cases for SAP ASE Commands

Use SAP ASE `sp_hadr_admin` command instead of RMA commands in the following situations.

Primary Server Starting as Standby Server

If the standby server goes down, and you subsequently restart the primary server, the HADR system starts the old primary server as a standby server to avoid a split-brain scenario.

You cannot use a RMA command to rectify the situation. Instead, log in to the current primary server as the system administrator and perform these steps:

1. (If the server retains the standby role, otherwise skip this step) Promote the server to the primary role (`sp_hadr_admin` must be run from the master database):

```
sp_hadr_admin primary
```

If the execution fails because the standby node is down, verify that neither node running as the primary and issue:

```
sp_hadr_admin primary, 'force'
```

2. Recover the databases for which recovery failed:

```
dbcc dbrecover('<database_name>')
```

3. Start all Replication Agents by issuing:

```
sp_start_rep_agent '<database_name>
```

4. Resume user application transaction activity

```
sp_hadr_admin activate
```

Performing Activity on Primary Without Planned Failover

Use the `sp_hadr_admin deactivate` parameter to perform maintenance or other activity on the primary server without performing a planned failover (from which you want to restrict transactional activities).

`sp_hadr_admin deactivate` can deactivate the server and demote it to standby, if necessary. The syntax is:

```
sp_hadr_admin deactivate, '<timeout>', '<label>'
sp_hadr_admin standby
```

Cancel an Ongoing Deactivation

Use this syntax to cancel an ongoing deactivation and restore the server to the primary active state:

```
sp_hadr_admin cancel
```

Drain the Transaction Logs

Use this syntax if you used the `nodrain` parameter during the deactivation but must subsequently drain the transaction logs to Replication Server:

```
sp_hadr_admin drain, '<label>'
```

Estimate Time Required for Rolling Back Active Transactions

Use this syntax to estimate the amount of time SAP ASE requires to roll back active transactions before you can issue `sp_hadr_admin deactivate` or `sp_hadr_admin drain` to drain the transaction log:

```
sp_hadr_admin failover_timeestimate
```


Viewing the Log Drain Activity

Use this syntax to view the status of the log drain activity on the primary server undergoing deactivation or log drain:

```
sp_hadr_admin status
```

Viewing the Transaction Replication Status

Use this syntax to view the transaction replication status on the standby server:

```
sp_hadr_admin status, '<label>'
```

Adding, Listing, or Dropping the Application Interface

Use this syntax to add an application interface:

```
sp_hadr_admin add_application_interface, '<HADR_server_name>' [, '<HADR_app_interface_netname>']
```

Use this syntax to list an application interface:

```
sp_hadr_admin 'list_application_interface'
```

Use this syntax to drop an application interface:

```
sp_hadr_admin 'drop_application_interface', '<HADR_server_name>'
```

12.2.2 sp_hadr_admin Syntax

Configures and maintains the HADR system.

Syntax

- Adds a server to the HADR member list:

```
sp_hadr_admin addserver, '<HADR_server_name>' [, [<pname>] [, 'noproagate' ]]
```

- Removes the server name from the member list:

```
sp_hadr_admin dropserver, '<HADR_server_name>' [, 'noproagate']
```

- Promotes the standby server to a primary server:

```
sp_hadr_admin primary, ['force']
```

You can run `sp_hadr_admin primary` only on the standby server.

- Resumes user application transaction activity on the new primary member after failover completes:

```
sp_hadr_admin activate
```

You can run `sp_hadr_admin activate` only on the inactive primary server.

- Moves the primary server to an inactive state:

```
sp_hadr_admin 'deactivate', '<timeout_period>' [, {'<label>' | NULL} ,
{'force' | NULL}, {'nodrain' | 'drain' | NULL} [, '<drain_timeout>']
```

If you do not provide `<label>`, it is auto-generated using this format:

```
HADR_LABEL_YYYYMMDD_HH:MM:SS:MS
```

The `deactivate` parameter triggers a transition from the active to the deactivating state, and then to the inactive state.

- Stops the deactivation process and moves the server back to an active state:

```
sp_hadr_admin cancel
```

You can execute `sp_hadr_admin cancel` only on primary servers in a deactivating state (that is, while `sp_hadr_admin deactivate` is executing).

- Changes the mode of an inactive HADR member from primary to standby mode. If the state of the primary server is not already inactive, use the `deactivate` parameter to change the state to inactive before issuing `standby`. Use the `force` parameter to force the mode change if the server is deactivated with `no_drain`:

```
sp_hadr_admin standby [, 'force']
```

- Initiates a log drain to Replication Server:

```
sp_hadr_admin drain [, {'<label>' | NULL} , '<drain_timeout>']
```

If you do not provide `<label>`, it is auto-generated using this format:

```
HADR_LABEL_YYYYMMDD_HH:MM:SS:MS
```

- Checks the progress of the deactivation process from primary and standby servers:

```
sp_hadr_admin status [, '<label>']
```

- Removes the specified HADR group and stops the local node from participating in an HADR system:

```
sp_hadr_admin dropgroup, '<group_name>'
```

- Estimates of the amount of time it takes to roll back transactions:

```
sp_hadr_admin failover_timeestimate [, <standby_server_name>]
```

- Returns the mode or state in which the HADR system is running:

```
sp_hadr_admin [mode | state]
```

- Provides the application network interface to the HADR cluster. Applications use this interface to connect to SAP ASE. The interface is added with the `addserver` parameter if you do not specify it with the `add_application_interface` parameter:

```
sp_hadr_admin [add_application_interface, '<HADR_server_name>',
              '<HADR_app_interface_netname>']
```

- Drops the application network interface from the HADR cluster:

```
sp_hadr_admin 'drop_application_interface', '<HADR_server_name>'
```

- Lists the application network interfaces on the HADR cluster:

```
sp_hadr_admin 'list_application_interface'
```

Parameters

<group_name>

specifies the HADR group you are adding or dropping.

<local_HADR_server_name>

specifies the HADR server name. The default is `<@@servername>`.

addserver

adds a server to the HADR system and member list.

<HADR_server_name>

is the server you are adding to – or dropping from – the HADR system or member list.

<pname>

is the name specified in the interfaces file for the server named `<HADR_server_name>`.

nopropagate

disables automatic propagation of the server name to the member list for standby servers and to user connections enabled with `hadr_list_cap` (user connections with HADR capabilities).

dropserver

drops a server from the member list.

primary

promotes the standby server to primary mode.

force

forces either of the following:

- A standby server to the role of a primary server if some members of the HADR system are unreachable
- A transition to an inactive state if there are ongoing (unprivileged or otherwise) transactions when the `<timeout>` expires

activate

activates the inactive primary server, and allows it to resume transaction activity by user applications.

deactivate

moves the primary server to an inactive state. All transactional activity is gracefully terminated on HADR and non-HADR databases by all connections (including privileged connections), and the transaction logs for all HADR databases are drained.

<timeout_period>

specifies the duration, in seconds, during which the primary server is allowed to remain in the deactivating state. If **<timeout_period>** expires before the server moves to the inactive state, it is moved back to active mode. You must use quotes for the numerical value.

<label>

specifies either of the following:

- The string used by the log drain mechanism to mark the transaction logs of those HADR databases that have been successfully drained by Replication Agent during the deactivation or drain process
- The string used to retrieve the status of transaction replication on the standby server. **<label>** is ignored if you issue `sp_hadr_admin status [, <label>]` on the primary server

nodrain

allows deactivation to occur without initiating a log drain. By default, if you do not include the `nodrain` parameter, the server initiates a log drain.

<drain_timeout>

is the length of time, in seconds, allocated for the log to drain. The drain is terminated after the timeout period elapses. If you include **<drain_timeout>** period with the `deactivate` parameter, the server is set to active mode after the timeout period ends.

cancel

aborts the deactivation process.

standby

moves the primary server to standby mode.

drain

drains the transaction log.

status

checks the progress of the deactivation process. The information `status` reports depends on whether the server is in primary or standby mode.

dropgroup

drops the HADR group. Do not execute `dropgroup` until you drop all the HADR members.

failover_timeestimate

estimates the amount of time it takes to roll back open transactions. Run `sp_hadr_admin failover_timeestimate` from the primary server.

<standby_server_name>

is unused in this version.

<Backup_Server_host>

is the hostname or IP address of the machine running Backup Server. If you specify NULL, `sp_hadr_admin` uses the host name of the listener that received the connection.

<host_to_authorize>

is the name of the SAP ASE host you are authorizing.

mode | state

returns the mode and state of the HADR system.

mode return values are:

- `NoHADR` – HADR system is disabled.
- `Primary` – HADR system is enabled. This is a primary server.
- `Standby` – HADR system is enabled. This is a standby server.
- `Unreachable` – HADR system is enabled, but the server is unreachable.
- `Starting` – HADR system is enabled, and the server is ready for initialization.

state return values are:

- `Unknown` – HADR system is in an unknown state.
- `Active` – Primary server allows transaction processing from user applications.
- `Inactive` – Server is inactive, and does not allow transaction processing from user applications.
- `Deactivating` – The server is changing from an active to an inactive state, and the log is being drained.

nowait_hadr

indicates that you are gracefully shutting down a primary server that is in an active state or has undrained transaction logs.

<HADR_app_interface_netname>

the `srvnetname` pointing to the entry in the `interfaces` file or an entry in the format `<host:port>`. The `srvnetname` is listed in `syssservers`.

Examples

Example 1

Adds a server named PARIS to the member list:

```
sp_hadr_admin addserver, PARIS
(return status = 0)
Adding server 'PARISDR', physical name 'PARIS'
Server added.
Command 'addserver' successful.
(1 row affected)
```

Example 2

Adds the PARISDR server to the member list using the <pname> format <hostname:port>:

```
sp_hadr_admin addserver, PARIS, 'daisy:4901'  
(return status = 0)  
Adding server 'PARISDR', physical name 'daisy:4901'  
Server added.  
Command 'addserver' successful.  
(1 row affected)
```

Example 3

Drops the PARIS server from the HADR group without propagating the change to the other members in the group:

```
sp_hadr_admin dropserver, PARIS, 'nopropagate'
```

Example 4

Attempts to promote the LONDON server to primary mode, but the HADR system cannot connect to the PARISDR server:

```
sp_hadr_admin primary  
Msg 11206, Level 16, State 1:  
Server 'JAIDB', Line 1:  
Unable to connect to server 'PARISDR'.  
Cannot promote the server to PRIMARY mode due to split brain check error.  
Use the 'force' option to force the promotion.  
Msg 19842, Level 16, State 1:  
Server 'MYSERVER', Procedure 'sp_hadr_admin', Line 531:  
'primary' encountered an error and could not succeed.
```

Example 5

Forcefully promotes the LONDON server to primary:

```
sp_hadr_admin primary, 'force'  
Command 'primary' successful.
```

Example 6

Activates the LONDON primary server:

```
sp_hadr_admin activate  
(return status = 0)  
Command 'event' successful.  
(1 row affected)  
(0 rows affected)  
(return status = 0)  
Command 'activate' successful.  
(1 row affected)
```

Example 7

Deactivates the current server LONDON with a timeout period of 30 seconds:

```
sp_hadr_admin 'deactivate','30','scheduled_offline'  
User connections statistics:: 0 in xact, 0 in chained mode, 9 in unchained  
mode, 0 holding server side cursors.  
Server reached INACTIVE state.  
Initiating log drain mechanism.  
Command 'deactivate' successful.  
(return status = 0)
```

Example 8

Cancels an ongoing deactivation process:

```
sp_hadr_admin cancel
Command 'cancel' successful.
(return status = 0)
```

Example 9

Changes the mode of the inactive primary server LONDON to standby:

```
sp_hadr_admin standby
Command 'standby' successful.
(return status = 0)
```

Example 10

Initiates the log drain using the string `scheduled_offline_07092013` as the label:

```
sp_hadr_admin drain, 'scheduled_offline_07092013'
Initiating log drain mechanism.
Command 'drain' successful.
(return status = 0)
```

Example 11

Drops the `IT_GROUP` group:

```
sp_hadr_admin dropgroup, 'IT_GROUP'
Command 'dropgroup' successful.
(return status = 0)
```

Example 12

Estimates the amount of time required to roll back transactions:

```
sp_hadr_admin fail over_timeestimate
total potential rollback time (mins)
-----
                                0
(1 row affected)
dbid  rep_drain_time
-----
(0 rows affected)
total potential rep drain time
-----
NULL
(1 row affected)
Command 'fail over_timeestimate' successful.
(return status = 0)
```

Example 13

Returns the current mode and state of the HADR system:

```
sp_hadr_admin mode
HADR Mode is :
-----
Starting
```

Example 14

Checks the state of the system:

```
sp_hadr_admin state
HADR State is :
-----
Inactive
```

Example 15

Adds the application interfaces on port 30015 to the LONDON and PARIS cluster hosts:

```
sp_hadr_admin 'add_application_interface','LONDON',lily:30015'
sp_hadr_admin 'add_application_interface','PARIS',daisy:30015'
```

Lists the application interfaces:

```
sp_hadr_admin 'list_application_interface'
name          network_name
-----
LONDON        lily:30015
PARIS         daisy:30015
```

Drops the application interfaces from the LONDON server:

```
sp_hadr_admin 'drop_application_interface', LONDON
```

Usage

- If you specify `<pname>`, `sp_hadr_admin` uses `<HADR_server_name>`. Use this format to specify the host name or IP address and port for the `<HADR_server_name>` server:
 - `<hostname:port>`
 - `<ipaddress:port>`
- `force` does not promote the standby server to a primary server if the HADR system detects an existing primary server. The administrator first demotes the existing primary server before reissuing the `force` parameter.
- The `deactivate` parameter triggers a transition from the active to the deactivating state, and then to the inactive state.
- `sp_hadr_admin deactivate` ignores the `<label>` parameter when you include the `nodrain` parameter.
- If the deactivation cannot complete in the period of time indicated by `<timeout_period>`, the server returns to active mode. Monitor the progress of replication by searching for the label `'scheduled offline'` in the Replication Server log.
- To initiate a log drain, the server must be in an inactive state. Execute `sp_hadr_admin drain` after executing `sp_hadr_admin deactivate nodrain`.
- To check the label on the standby server and monitor the replication status, use the `sp_hadr_admin status` parameter. To avoid confusion, use a different label for each deactivation process.
- Running `sp_hadr_admin status [, <label>']` on the primary server displays the database name and the log drain status for each HADR database. The `label` parameter is ignored on the primary server. The log drain status is one of:

- `completed` – Replication Agent has processed the deactivation checkpoint marker.
- `in progress` – a drain operation has started, and Replication Agent has not yet processed the checkpoint marker.
- `inactive` – there is no drain operation in progress.

For example:

```
sp_hadr_admin status
Database Name      Log Drain Status      Log Pages Left
-----
db1                completed              0
db2                completed              0
master             completed              0
```

Running `sp_hadr_admin status [, '<label>']` on the standby server shows:

- The status of the replication process corresponding to the specified label, allowing a privileged user to determine if replication is complete.
- The database name and status of replication for each HADR database. Status listed as `completed` or `pending`. `completed` means the replicated data is applied to the standby database up to the deactivation checkpoint marker and its associated label. `pending` means this process is in progress.

This example displays the status of a standby server using the label `'deactivate 2012/10/10'`:

```
sp_hadr_admin status, 'deactivate 2012/10/10'
Database Name      Replication Status
-----
db1                pending
db2                pending
master             completed
```

12.3 RMA Proxy Tables

When the HADR system has been set up, a number of proxy tables are created in the master database of the primary server.

The proxy tables include:

- `hadrGetLog`
- `hadrGetTicketHistory`
- `hadrStatusActivePath`
- `hadrStatusPath`
- `hadrStatusResource`
- `hadrStatusRoute`
- `hadrStatusSynchronization`

12.3.1 hadrGetLog

Provides information for the RMA and Replication Server logs.

Syntax

```
select * from master..hadrGetLog where _logicalHost='<host_name>' and
_serverType='[RS | RMA]' and _startDate='<YYYY-MM-DD HH:MM:SS>'
```

Parameters

_logicalHost

The logical hostname.

_serverType

Type of server. Either **RS** (Replication Server) or **RMA** (Replication Management Agent).

_startDate

The start date of the log. The format is 'YYYY-MM-DD HH:MM:SS'

Table 18: Column Information for hadrGetLog

Information	Description
Severity	Severity level of the log.
Date	Date of the log.
Message	Content of the log.
_logicalHost	Logical host name.
_serverType	Type of server. Either RS (Replication Server) or RMA (Replication Management Agent).
_startDate	The start date of the log. The format is 'YYYY-MM-DD HH:MM:SS'.

Examples

Provides information about the local host **PR** for the specified date.

```
select * from master..hadrGetLog where _logicalHost='PR' and _serverType='RS'
and _startDate='2015-11-20 11:00:00'
go
```

```
I 2015-11-20 02:09:40.000 The sqm_seg_prealloc option is
depreciated. SFHADR1 RS 2015-02-02 11:00:00
I 2015-11-20 02:09:40.000 Open partition '/opt/sap/partition1.dat' with
sqm_write_flush 'on'. SFHADR1 RS 2015-02-02 11:00:00
W 2015-11-20 02:09:44.000 WARNING #5059 GLOBAL RS(GLOBAL RS) - neric/dsi/
dsiqmext.c(432) SFHADR1 RS 2015-02-02 11:00:00
W 2015-11-20 02:09:44.000 WARNING #5059 GLOBAL RS(GLOBAL RS) - neric/dsi/
dsiqmext.c(432) SFHADR1 RS 2015-02-02 11:00:01
```

```

T 2015-11-20 02:09:44.000 (5): Create task manager with 2
tasks. SFHADR1 RS 2015-02-02 11:00:01
I 2015-11-20 02:09:44.000 Replication Server 'PI2_REP_PR' is
started. SFHADR1 RS 2015-02-02 11:00:02
E 2015-11-20 02:09:44.000 ERROR #21 USER( ) - /generic/exec/
exec.c(525) SFHADR1 RS 2015-02-02 11:00:02

```

12.3.2 hadrGetTicketHistory

The `hadrGetTicketHistory` table is a proxy table that is created after the HADR system is set up. After you log into the primary ASE with user `DR_admin`, use the `hadrGetTicketHistory` command to retrieve multiple ticket rows (one ticket per row) where `pdb_t` is bigger than a user specified start date.

`pdb_t` indicates the time stamp when the ticket was injected in corresponding primary ASE database.

Syntax

```

select * from master..hadrGetTicketHistory where _logicalHost='<hostname>' and
_DBName='<database>' and _startDate='<YYYY-MM-DD HH:MM:SS>'

```

Parameters

`_logicalHost`

The name of logical ASE host, to query ticket information.

`_DBName`

The name of the database.

`_startDate`

The start date since when ticket information must be returned, the format is 'YYYY-MM-DD HH:MM:SS'

Examples

Displays ticket information on the HA host for ERP database since '2010-01-01 15:30:00':

```

select * from master..hadrGetTicketHistory where _logicalHost='HA' and
_DBName='ERP' and _startDate='2010-01-01 15:30:00'

```

12.3.3 `hadrStatusActivePath`

The `hadrStatusActivePath` displays information about the active connection status between local and remote HADR hosts.

Syntax

```
select * from master..hadrStatusActivePath
```

Information	Description
<code>LogicalHost</code>	The logical host name of the server in HADR system.
<code>Hostname</code>	The host name of the server.
<code>hadrMode</code>	An external mode that is visible to and known by other HADR members, including "Primary", "Standby", "Disabled", "Unreachable" and "Starting".
<code>hadrState</code>	An internal state that is known only by the member, including "Active", "Inactive" and "Deactivating".
<code>SynchronizationMode</code>	The replication synchronization mode you have configured between a database and the SAP Replication Server, which can be one of: <ul style="list-style-type: none">• Synchronous• Asynchronous
<code>SynchronizationState</code>	Synchronization Mode in which replication is currently operating.
<code>DistributionMode</code>	The replication distribution mode you have configured between a database and the Replication Server, which can be one of: <ul style="list-style-type: none">• Local• Remote
<code>DistributionPath</code>	The target distribution server.
<code>Pathname</code>	The path of Replication Server through which transactions travel
<code>State</code>	The state of the replication path, including "Active", "Suspended", and so on. Displays only those paths where <code>State</code> is <code>Active</code> in <code>hadrStatusActivePath</code> .
<code>LatencyTime</code>	The time at which the latency was last calculated.
<code>Latency</code>	The latency in milliseconds.
<code>CommitTime</code>	The time at which the last commit was replicated.

Information	Description
DrainStatus	<p>The status of draining the primary database server's transaction logs. Values are:</p> <ul style="list-style-type: none"> • Drained: The primary database server's transaction logs are completely transferred to Replication Server. • Not Drained: The primary database server's transaction logs are only partially transferred to Replication Server. • Unknown: The status cannot be queried.

Examples

Displays connection information between the PR and HA hosts.

```
select * from master..hadrStatusActivePath
go
```

```
LogicalHost Hostname  hadrMode  hadrState  SynchronizationMode
SynchronizationState DistributionMode DistributionPath Pathname
State      LatencyTime          Latency    CommitTime
DrainStatus
-----
-----
-----
PR          site0      Primary   Active     Synchronous
Synchronous NULL      Remote    NULL      NULL      NULL
NULL       NULL
HA          site1     Standby   Inactive   Synchronous
Inactive   NULL      Remote    NULL      NULL      NULL
NULL       NULL
NULL       NULL      NULL      NULL      NULL
NULL       NULL      HA        NULL      PR.HA.PI2
Active     Unknown   Unknown   2016-01-21 01:20:28.952 Not
Drained
NULL       NULL      NULL      NULL      NULL
NULL       NULL      NULL      HA        PR.HA.db1
Active     Unknown   Unknown   Unknown   Unknown   Not
Drained
NULL       NULL      NULL      NULL      NULL
NULL       NULL      NULL      HA        PR.HA.master
Active     2016-01-21 01:06:33.866 884      2016-01-21 01:12:20.260 Not
Drained
(5 rows affected)
```

12.3.4 hadrStatusPath

Displays information about the connection status between local and remote HADR hosts.

Syntax

Use this syntax to retrieve information from `hadrStatusPath`:

```
select * from master..hadrStatusPath
```

Columns

The columns for `hadrStatusPath` are:

Column Name	Description
<code>LogicalHost</code>	Logical host name of the server in the HADR system.
<code>Hostname</code>	Host name of the server.
<code>hadrMode</code>	An external mode that is visible to and known by other HADR members. One of: "Primary", "Standby", "Disabled", "Unreachable", and "Starting".
<code>hadrState</code>	Internal state that is known only by the member. One of: "Active", "Inactive", and "Deactivating".
<code>SynchronizationMode</code>	The replication synchronization mode configured between a database and Replication Server. One of: <ul style="list-style-type: none">• Synchronous• Asynchronous
<code>SynchronizationState</code>	Synchronization Mode in which replication is currently operating
<code>DistributionMode</code>	The replication distribution mode configured between a database and Replication Server. One of: <ul style="list-style-type: none">• Local• Remote
<code>DistributionPath</code>	Target distribution server.
<code>Pathname</code>	Path of Replication Server through which transactions travel
<code>State</code>	State of the replication path. Set to "Active", "Suspended", and so on.
<code>LatencyTime</code>	Time that latency was last calculated.
<code>Latency</code>	Amount of latency, in milliseconds.
<code>CommitTime</code>	Time the last commit was replicated.

Column Name	Description
DrainStatus	<p>The status of draining the primary database server's transaction logs. Values are:</p> <ul style="list-style-type: none"> • Drained: The primary database server's transaction logs are completely transferred to Replication Server. • Not Drained: The primary database server's transaction logs are only partially transferred to Replication Server. • Unknown: The status cannot be queried.

Examples

```
select * from master..hadrStatusPath
go
```

```
LogicalHost Hostname  hadrMode  hadrState  SynchronizationMode
SynchronizationState DistributionMode DistributionPath Pathname
State        LatencyTime                Latency    CommitTime
DrainStatus
-----
-----
-----
PR          site0      Primary   Active     Synchronous
Synchronous Remote      NULL      NULL      NULL
NULL       NULL
NULL
HA          site1      Standby   Inactive   Synchronous
Inactive   Remote      NULL      NULL      NULL
NULL       NULL
NULL       NULL      NULL      NULL      NULL
Suspended  Unknown   Unknown   Unknown   Unknown   HA.PR.master
Unknown
NULL       NULL      NULL      NULL      NULL
Suspended  Unknown   Unknown   Unknown   Unknown   HA.PR.db1
Unknown
NULL       NULL      NULL      NULL      NULL
NULL       NULL      NULL      HA        NULL      PR.HA.PI2
Active     Unknown   Unknown   Unknown   2016-01-21 01:20:28.952 Not
Drained
NULL       NULL      NULL      NULL      NULL
NULL       NULL      NULL      HA        NULL      PR.HA.db1
Active     Unknown   Unknown   Unknown   Unknown   Unknown   Not
Drained
NULL       NULL      NULL      NULL      NULL
NULL       NULL      NULL      HA        NULL      PR.HA.master
Active     2016-01-21 01:06:33.866 884      2016-01-21 01:12:20.260 Not
Drained
NULL       NULL      NULL      NULL      NULL
NULL       NULL      NULL      PR        NULL      HA.PR.PI2
Suspended  Unknown   Unknown   Unknown   Unknown
Unknown
(8 rows affected)
```

12.3.5 hadrStatusResource

The `hadrStatusResource` command monitors the resources in the Replication Server and ASE that are known to the DR Agent.

Use the `hadrStatusResource` command to see:

- Replication device size and usage
- Primary ASE transaction log size and backlog
- Replication simple persistent queue backlog
- Replication inbound queue backlog
- Replication outbound queue backlog
- Replication route queue backlog
- Replication truncation backlog (inbound queue - IBQ, outbound queue - OBQ, and route queue - RQ)

Syntax

```
select * from master..hadrStatusResource
```

Table 19: Column Information for `hadrStatusResource`

Information	Description
LogicalHost	Logical host name of the server in the HADR system.
DBname	Database name.
RSDeviceSize	Replication Server device size.
RSDeviceUsed	Replication Server device usage.
ASELogSize	Primary SAP ASE transaction log size.
ASELogBacklog	Primary SAP ASE transaction log backlog.
RSQueueBacklog	Summary of the replication queues of Replication Server.
EstimatedFailoverTime	Estimated amount of time left for the primary server to failover to the standby server.
RSSPQBacklog	Replication Server simple persistent queue backlog.
RSIBQBacklog	Replication Server inbound queue backlog.
RSOBQBacklog	Replication Server outbound queue backlog.
RSRouteQBacklog	Replication Server route queue backlog.
TargetLogicalHost	Logical name of the target distribution server.
RSIBQTruncationBacklog	Replication Server inbound queue truncation backlog.
RSOBQTruncationBacklog	Replication Server outbound queue truncation backlog.

Information

Description

RSRouteQTruncationBacklog
g
Replication Server route queue truncation backlog

Example

```
select * from master..hadrStatusResource
go
```

```
LogicalHost DBName RSDeviceSize RSDeviceUsed ASELogSize
ASELogBacklog RSQueueBacklog EstimatedFailoverTime RSPQBacklog
RSIBQBacklog RSOBQBacklog RSRouteQBacklog TargetLogicalHost
RSIBQTruncationBacklog RSOBQTruncationBacklog
RSRouteQTruncationBacklog
```

```
-----
```

LogicalHost	DBName	RSDeviceSize	RSDeviceUsed	ASELogSize	ASELogBacklog	RSQueueBacklog	EstimatedFailoverTime	RSPQBacklog	RSIBQBacklog	RSOBQBacklog	RSRouteQBacklog	TargetLogicalHost	RSIBQTruncationBacklog	RSOBQTruncationBacklog	RSRouteQTruncationBacklog
PR	NULL	256	208	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL
NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL
NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL
DR	NULL	512	64	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL
NULL	NULL	NULL	-1	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL
NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL
HA	NULL	256	208	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL
NULL	NULL	NULL	-1	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL
NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL
PR	master	NULL	NULL	100	0	0	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL
NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL
PR	PI2	NULL	NULL	10	0	0	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL
NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL
PR	db1	NULL	NULL	20	0	0	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL
NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL
DR	master	NULL	NULL	NULL	NULL	0	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL
NULL	0	NULL	NULL	NULL	NULL	0	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL
DR	PI2	NULL	NULL	NULL	NULL	0	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL
NULL	0	NULL	NULL	NULL	NULL	0	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL
DR	db1	NULL	NULL	NULL	NULL	0	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL
NULL	0	NULL	NULL	NULL	NULL	0	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL
HA	master	NULL	NULL	NULL	0	0	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL
NULL	0	NULL	NULL	0	0	0	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	0
0	0	NULL	NULL	NULL	0	0	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	0
0	0	0	NULL	NULL	0	0	NULL	NULL	NULL	NULL	NULL	NULL	NULL	NULL	0

```

HA          PI2          NULL          NULL          NULL
NULL        0           NULL          NULL          0
0           0           NULL          NULL          NULL
0           0           0           NULL          NULL
HA          db1          NULL          NULL          NULL
NULL        0           NULL          NULL          0
0           0           NULL          NULL          NULL
0           0           0           NULL          NULL
HA          NULL         NULL          NULL          NULL
NULL        NULL         NULL          NULL          NULL
NULL        NULL         0           DR           NULL
NULL        NULL         NULL          0           NULL

(13 rows affected)

```

12.3.6 hadrStatusRoute

Displays the status of the connection routes between the HADR nodes.

Use this syntax to retrieve information from hadrStatusRoute:

```
select * from master..hadrStatusRoute
```

Table 20: Column Information for hadrStatusRoute

Column Name	Description
Path	Replication Server path through which transactions travel.
Sequence	Order of the queues or threads of the replication path
Name	Name of the queues or threads.
Type	Indicates whether it is a queue or a thread
Qid	ID of the queue
Spid	ID of the thread
Site	Replication Server name
State	State of the thread
Backlog	Backlog of the queue.

Example

```
select * from master..hadrStatusRoute
go
```

```

Path      Sequence  Name      Type      Qid      Spid
Site      State     Backlog
-----

```

SFHADR1.SJHADR2.master	1	ASE	S	NULL
16298 SFMACHINE1	Active		0	
SFHADR1.SJHADR2.master	2	RAT	T	NULL
25 SFMACHINE1	Active		NULL	
SFHADR1.SJHADR2.master	3	RATCI	T	NULL
NULL SJMACHINE2	Active		NULL	
SFHADR1.SJHADR2.master	4	SPQ	Q	106
NULL SJMACHINE2	NULL		0	
SFHADR1.SJHADR2.master	5	CAP	T	NULL
293 SJMACHINE2	Awaiting Command		NULL	
SFHADR1.SJHADR2.master	6	SQM	T	NULL
231 SJMACHINE2	Awaiting Message		NULL	
SFHADR1.SJHADR2.master	7	IBQ	Q	106
NULL SJMACHINE2	NULL		0	
SFHADR1.SJHADR2.master	8	SQT	T	NULL
274 SJMACHINE2	Awaiting Wakeup		NULL	
SFHADR1.SJHADR2.master	9	DIST	T	NULL
273 SJMACHINE2	Awaiting Wakeup		NULL	
SFHADR1.SJHADR2.master	10	SQM	T	NULL
270 SJMACHINE2	Awaiting Message		NULL	
SFHADR1.SJHADR2.master	11	OBQ	Q	105
NULL SJMACHINE2	NULL		0	
SFHADR1.SJHADR2.master	12	DSI	T	NULL
337 SJMACHINE2	Awaiting Message		NULL	
SFHADR1.SJHADR2.master	13	ASE	S	NULL
14663 SJMACHINE2	Active		NULL	
SFHADR1.SJHADR2.db1	1	ASE	S	NULL
16298 SFMACHINE1	Active		0	
SFHADR1.SJHADR2.db1	2	RAT	T	NULL
155 SFMACHINE1	Active		NULL	
SFHADR1.SJHADR2.db1	3	RATCI	T	NULL
NULL SJMACHINE2	Active		NULL	
SFHADR1.SJHADR2.db1	4	SPQ	Q	114
NULL SJMACHINE2	NULL		0	
SFHADR1.SJHADR2.db1	5	CAP	T	NULL
305 SJMACHINE2	Awaiting Command		NULL	
SFHADR1.SJHADR2.db1	6	SQM	T	NULL
230 SJMACHINE2	Awaiting Message		NULL	
SFHADR1.SJHADR2.db1	7	IBQ	Q	114
NULL SJMACHINE2	NULL		0	
SFHADR1.SJHADR2.db1	8	SQT	T	NULL
285 SJMACHINE2	Awaiting Wakeup		NULL	
SFHADR1.SJHADR2.db1	9	DIST	T	NULL
283 SJMACHINE2	Awaiting Wakeup		NULL	
SFHADR1.SJHADR2.db1	10	SQM	T	NULL
269 SJMACHINE2	Awaiting Message		NULL	
SFHADR1.SJHADR2.db1	11	OBQ	Q	113
NULL SJMACHINE2	NULL		0	
SFHADR1.SJHADR2.db1	12	DSI	T	NULL
341 SJMACHINE2	Awaiting Message		NULL	
SFHADR1.SJHADR2.db1	13	ASE	S	NULL
14663 SJMACHINE2	Active		NULL	
SFHADR1.SJHADR2.PI2	1	ASE	S	NULL
16298 SFMACHINE1	Active		0	
SFHADR1.SJHADR2.PI2	2	RAT	T	NULL
144 SFMACHINE1	Active		NULL	
SFHADR1.SJHADR2.PI2	3	RATCI	T	NULL
NULL SJMACHINE2	Active		NULL	
SFHADR1.SJHADR2.PI2	4	SPQ	Q	110
NULL SJMACHINE2	NULL		0	
SFHADR1.SJHADR2.PI2	5	CAP	T	NULL
299 SJMACHINE2	Awaiting Command		NULL	
SFHADR1.SJHADR2.PI2	6	SQM	T	NULL
229 SJMACHINE2	Awaiting Message		NULL	
SFHADR1.SJHADR2.PI2	7	IBQ	Q	110
NULL SJMACHINE2	NULL		0	
SFHADR1.SJHADR2.PI2	8	SQT	T	NULL
282 SJMACHINE2	Awaiting Wakeup		NULL	

```

SFHADR1.SJHADR2.PI2          9          DIST      T          NULL
281      SJMACHINE2      Awaiting Wakeup  NULL
SFHADR1.SJHADR2.PI2          10         SQM        T          NULL
268      SJMACHINE2      Awaiting Message NULL
SFHADR1.SJHADR2.PI2          11         OBQ        Q          109
NULL     SJMACHINE2      NULL
SFHADR1.SJHADR2.PI2          12         DSI        T          NULL
345      SJMACHINE2      Awaiting Message NULL
SFHADR1.SJHADR2.PI2          13         ASE        S          NULL
14663    SJMACHINE2      Active
(39 rows affected)

```

12.3.7 hadrStatusSynchronization

Displays the status of synchronization between the HADR nodes.

Use this syntax to retrieve information from `hadrStatusSynchronization`:

```
select * from master..hadrStatusSynchronization
```

Table 21: Column Information for `hadrStatusSynchronization`

Column Name	Description
LogicalHost	Logical host name of the server in HADR system.
hadrMode	An external mode that is visible to and known by other HADR members. One of <ul style="list-style-type: none"> • Primary • Standby • Disabled • Unreachable • Starting
hadrState	An internal state that is known only by the member. One of: <ul style="list-style-type: none"> • Active • Inactive • Deactivating
Hostname	Host name of the server.
SynchronizationMode	Configured synchronization mode value. One of <ul style="list-style-type: none"> • Synchronous • Asynchronous
DBName	The database name.
SynchronizationState	Synchronization mode in which replication is currently operating.

Example

```
select * from master..hadrStatusSynchronization
```

```
go
```

```
LogicalHost  Hostname  hadrMode  hadrState  SynchronizationMode  DBName
SynchronizationState
-----
SFHADR1     SFMACHINE1  Primary  Active     Synchronous
NULL        Synchronous
SFHADR1     SFMACHINE1  Primary  Active     Synchronous
master      Synchronous
SFHADR1     SFMACHINE1  Primary  Active     Synchronous
db1         Synchronous
SFHADR1     SFMACHINE1  Primary  Active     Synchronous
PI2         Synchronous
```

12.4 Replication Server Commands

Replication Server provides commands and parameters that you use in an HADR system.

12.4.1 Replication Server Configuration Parameters

Replication Server provides parameters for configuring stream replication, simple persistent queue (SPQ), Capture, and external replication.

Table 22: Stream Replication Configuration Parameters

Parameter	Description
<code>ci_pool_size</code>	<p>Specifies the number of packages in a stream replication buffer pool. When a stream replication stream starts, the stream replication buffer pool pre-allocates the number of stream replication packages in the pool as specified by <code>ci_pool_size</code>. The package is fetched from this pool when needed.</p> <p>Value: 10 to 10,000</p> <p>Default: 50</p>
<code>ci_package_size</code>	<p>Specifies the size of a stream replication package. Each stream replication package in the stream replication buffer pool shares the same size specified by <code>ci_package_size</code>.</p> <p>Value: 2,048 (bytes) to 104,857,600 (bytes)</p> <p>Default: 1,048,576 (bytes)</p>

Parameter	Description
<code>ci_max_cmds</code>	Specifies the maximum number of commands that can be put in a stream replication package. Value: 1 to 1,000,000 Default: 80

The total size of the stream replication buffer pool is `ci_pool_size x ci_package_size`. The buffer is allocated when the stream replication stream starts. For example, if you set 50 to `ci_pool_size` and 1 MB to `ci_package_size`, about 50 MB memory is allocated when the stream replication stream starts.

Table 23: Simple Persistent Queue (SPQ) Configuration Parameters

Parameter	Description
<code>spq_file_directory</code>	Specifies the SPQ data file directory. <code>spq_file_directory</code> must point to a valid OS directory. If the directory does not exist or is an empty directory, stream replication creates a new directory accordingly and adds SPQ data files in it. Two SPQs cannot share a common directory. Value: file path Default: the path of the first Replication Server partition
<code>spq_min_size</code>	Specifies the minimum size of an SPQ. When an SPQ is created for the first time, two SPQ data files are created with half the size specified by <code>spq_min_size</code> . Value: 20 MB to 100 GB Default: 2 GB
<code>spq_max_size</code>	Specifies the maximum size of an SPQ. The total size of all SPQ data files cannot exceed the size set by <code>spq_max_size</code> . When data reach the value of <code>spq_max_size</code> , SPQ becomes full, a new SPQ is created. Value: 100 MB to 1 TB Default: 100 GB
<code>spq_data_file_size</code>	Specifies the size of each SPQ data file. If adding the next message makes the data file larger than the specified size, SPQ writes the message to the next SPQ data file. SPQ truncates or removes the data file after all the messages have been read by Capture. Value: 1 MB to 10 GB Default: 1 GB

Parameter	Description
<code>spq_write_flush</code>	<p>Set <code>spq_write_flush</code> to on to bypass writing on OS buffer. Disk writes are committed only when data are actually written on disk, rather than memory buffers.</p> <p>Value: on or off</p> <p>Default: on</p>
<code>spq_cache_size</code>	<p>Specifies the SPQ cache size. SPQ cache is a memory pool for caching en-queue data before they are written on disk and reading the de-queue data beforehand.</p> <p>Value: 1 KB to 100 MB</p> <p>Default: 10 MB</p>

Table 24: Capture-Related Configuration Parameters

Parameter	Description
<code>stream_replication</code>	<p>Specifies the communication protocol between Replication Agent and Replication Server. The parameters must be set on both Replication Agent and Replication Server and agree with each other. If the values are different, the connection is refused.</p> <p>Set <code>stream_replication</code> to true to switch the primary connection to synchronous replication.</p> <p>Value: true or false</p> <p>Default: false</p>
<code>cap_sqm_write_request_limit</code>	<p>Specifies the maximum size (in bytes) of pending requests a Capture can have waiting to be written to the inbound queue (IBQ). If the number of pending requests reaches the value, the Capture stops reading data from the SPQ. The parameter is used for flow controlling of the Capture.</p> <p>Value: 2 MB to 2 GB</p> <p>Default: 8 MB</p>
<code>cap_prs_num_threads</code>	<p>Specifies the number of concurrent parsers in the Capture. In a Capture, multiple parsers can process different stream replication packages concurrently. Increasing the number of parsers improves the throughput of the Capture.</p> <p>Value: 1 to 20</p> <p>Default: 2</p>

Parameter	Description
<code>truncation_mode</code>	<p>Specifies when to move the secondary truncation point (STP) in SAP ASE. When the truncation mode is <code>ci</code>, STP is requested to move when a stream replication package is persisted in the SPQ. When the truncation mode is <code>capture</code>, STP is requested to move when the Capture persists commands in the IBQ.</p> <p>Value: <code>ci</code> or <code>capture</code></p> <p>Default: <code>capture</code></p>

Table 25: External Replication Configuration Parameter

Parameter	Description
<code>cap_filter_dr_admin_name</code>	<p>Changes the DR_admin HADR administrator user to the login name that you used to set up HADR. The external Replication Server then filters out all commands issued under the login name when replicating out of an HADR system.</p> <p>If you change the HADR administrator user at the:</p> <ul style="list-style-type: none"> • Server level – restart all capture components • Connection level – restart the corresponding capture

12.4.2 Monitoring State Information

Replication Server adds new parameters into some commands to monitor states for stream replication, Capture, and SPQ.

12.4.2.1 admin disk_space, mb, spq

Displays disk space usage information of a specific SPQ or SPQs for the Capture path.

Syntax

```
admin disk_space, mb, spq[, <dsname>, <dbname>]
```


Parameters

mb

Prints backlog information for the Capture. The Backlog column displays the value of backlog.

spq

Displays disk usage information about the simple persistent queues for the Capture path.

<dsname>

Specifies the name of the primary data server for the simple persistent queue.

<dbname>

Specifies the name of the primary database for the simple persistent queue.

Examples

Example 1

Displays information for an SPQ:

```
admin disk_space, spq, QG30, tpcc
```

```
Queue #          SPQ Directory
-----
 108 QG30.tpcc  /replinuxb25-2/users/xdai/RS25/spq108

Total Size      Used Size      Packages      State
-----
          100              0              0  ON-LINE
```

Usage

Table 26: Column Descriptions for the output of admin disk_space, mb, spq

Column Name	Description
Queue #	The number of this SPQ.
SPQ Directory	The file directory of this SPQ.
Total Size	
Used Size	Total size in megabytes currently used by the Replication Server.
Packages	The number of packages which are not read out by Capture; the backlog of the SPQ for the Capture path.

Column Name	Description
State	State of this device: <ul style="list-style-type: none"> • ON-LINE – The device is normal • OFF-LINE – The device cannot be found
Backlog	Displays the value of backlog.

12.4.2.2 admin disk_space, mb, spqra

Displays disk space usage and backlog information of a specific SPQ or SPQs for the SPQ Agent path.

Syntax

```
admin disk_space, mb, spqra[, <dsname>, <dbname>]
```

Parameters

mb

Prints backlog information for the SPQ Agent. The Backlog column displays the value of backlog.

spqra

Displays disk usage information about the simple persistent queue for the SPQ Agent.

<dsname>

Specifies the name of the primary data server for the SPQ Agent.

<dbname>

Specifies the name of the primary database for the SPQ Agent.

Examples

Example 1

Displays backlog and other information for an SPQ Agent:

```
admin disk_space, mb, spqra
```

```
Queue #           SPQ Directory
```

```

-----
114 ERP_c0site0_R1.ERP /remote/repeng11/users/gzhong/SP305DR2/spq114
-----
Total Size      Used Size      Packages      State      BACKLOG
-----
                100              0              0  ON-LINE  0
-----

```

Usage

Table 27: Column Descriptions for the output of admin disk_space, mb, spqra

Column Name	Description
Queue #	The number of this SPQ.
SPQ Directory	The file directory of this SPQ.
Total Size	Total size, in megabytes, of the Replication Server.
Used Size	Total size, in megabytes, currently used by the Replication Server.
Packages	The number of packages that are not read out by the SPQ Agent; the backlog of SPQ for the SPQ Agent.
State	State of this device: <ul style="list-style-type: none"> ON-LINE – The device is normal OFF-LINE – The device cannot be found
Backlog	Displays the value of backlog.

12.4.2.3 admin who, spq

Displays information and statistics related to the SPQ writer and readers (Capture and SPQ Agent) for a specific SPQ or all SPQs.

Syntax

```
admin who, spq[, {<ds>, <db>} | <dbid>]
```

Parameters

spq

Displays information for SPQ.

<ds>

Specifies the name of the primary data server for the SPQ.

<db>

Specifies the name of the primary database for the SPQ.

<dbid>

Specifies the database ID.

Examples

Example 1

Displays detailed information, including both SPQ writer and readers for a specific SPQ:

```
admin who, spq, ERP_c0site0_R1, ERP
```

The returned result is:

```
Spid State          Info
-----
      Active          114 ERP_c0site0_R1.ERP
SPQ Reader Statistics Report
=====
ReaderName          ReadState          ReadIndex
SPQP-ReadBytes      SPQP-ReadDiskTime
SPQP-ReadFileNum    SPQP-ReadFileLoc
SPQP-ReadMsgIndex
-----
CAP                  QWait              68
0                    0                  0
0                    0                  0
0                    0                  0
SPQ AGENT           QWait              68
0                    0                  0
0                    0                  0
0                    0                  0
SPQ Writer Statistics Report
=====
WriteState          WriteIndex          WriteDuplicate
SPQP-WriteBytes     SPQP-WriteDiskTime
SPQP-WriteFileNum   SPQP-WriteFileLoc
SPQP-WriteMsgIndex
-----
Dup                 68                  59
813476              1665022
1                   813476
68
```

Example 2

Displays generic information for SPQs:

```
admin who, spq
go
```

The returned result is:

Spid	State	Info	ReadState	WriteState
	Active	114 ERP_c0site0_R1.ERP	QWait	DUP
ReadIndex	WriteIndex	MQSize		
3	3	1826		

Usage

- This command supports the monitoring of multiple SPQ readers, including Capture and SPQ Agent.
- If you run the command:
 - without any parameters, it prints the results of the SPQ writer, and the Capture reader
 - with the `<ds.dbname>` or `<dbid>` parameters, it prints detailed results which include both SPQ writer and SPQ readers statistics.

Table 28: Column Descriptions for `admin who, spq` Output

Column Name	Description
ReadIndex	The current index of SPQ reader.
WriteIndex	The current index of SPQ writer.
Duplicate	The number of duplicate packages received.
MQSize	The size of message queue used in SPQ.
SPQP-ReadBytes	The number of bytes read from SPQ.
SPQP-WriteBytes	The number of bytes written into SPQ.
SPQP-ReadDiskTime	The total time used by reading disk.
SPQP-WriteDiskTime	The total time used by writing disk.
SPQP-ReadFileNum	The number of the file currently being read.
SPQP-ReadFileLoc	The current offset SPQ reader.
SPQP-ReadMsgIndex	The current index of SPQ reader.
SPQP-WriteFileNum	The number of the file currently being written.
SPQP-WriteFileLoc	The current offset SPQ writer.
SPQP-WriteMsgIndex	The current SPQ writer index.

12.4.2.4 admin who, spqra

Displays state and configuration information for the SPQ Agent.

Syntax

```
admin who, spqra
```

Examples

Example 1

Run `admin who, spqra`:

```
admin who, spqra
```

The returned result is:

Spid	State	Info	MemberState
361	Awaiting Command	110(33554432) ERP_site0_R1.ERP	ACTIVE
	Down	109(33554432) ERP_site1.ERP	INACTIVE
	Drained	connect_dataserver	connect_database
NO		mo5_28260	
ERP			
NO		mo5_28260	
ERP			
rs_servername		maint_username	
mo-d558c2551:9885		ERP_ex_maint	
mo-d558c2551:9885		ERP_ex_maint	

12.4.2.5 admin who, cap

Displays information for a specific Capture component or all Capture components.

Syntax

```
admin who, cap[, {<ds>, <db>} | <dbid>]
```

Parameters

cap

Displays information for Capture components.

<ds>

Name of the primary data server for the Capture component.

<db>

Name of the primary database for the Capture component.

<dbid>

Specifies the database ID.

Examples

Example 1

Displays information for the Capture:

```
admin who, cap, QG30, tpcc
```

The returned result is:

```
Spid      State          Info          PkgsRecv
-----
   107  Awaiting Command  108 QG30.tpsc      5498
CmdsRecv  SchemasRecv  BytesRecv    Pending Pkgs
-----
   91183           0      98816036           0
Pending Writes  Pending Bytes  Pending Max  Pending truncs
-----
           0           0      8388608           0

Last OQID received
-----
0x00000000092e24d4003ec0580003003ec0580000000a4f201814e8a0000000000000000

Last OQID moved
-----
0x00000000092e24d4003ec0580003003ec0580000000a4f201814e8a0000000000000000

Last OQID delivered
-----
0x00000000092e24d4003ec0580003003ec0580000000a4f201814e8a0000000000000000
```

Usage

Table 29: Column Descriptions for `admin who, cap` Output

Column Name	Description
PkgsRecv	The number of packages this Capture receives.
CmdsRecv	The number of commands this Capture receives.
SchemasRecv	The number of schemas this Capture receives.
BytesRecv	The number of bytes this Capture receives.
Pending Pkgs	The number of packages which are pending to be parsed by this Capture.
Pending Writes	The number of pending IBQ writes requested by this Capture.
Pending Bytes	The number of bytes pending to write into IBQ requested by this Capture.
Pending Max	The maximum number of pending bytes to write to the IBQ. The output of this option is configured by <code>cap_sqm_write_request_limit</code> .
Pending Truncs	The number of possible truncation points.
Last OQID Received	The OQID of the last command this Capture receives.
Last OQID Moved	The truncation point this Capture requests to move last time.
Last OQID Delivered	The OQID of the latest command which is persisted in IBQ.

12.4.2.6 admin who, ci

Displays information for a specific stream replication stream or all stream replication streams.

Syntax

```
admin who, ci[, {<ds>, <db>} | <dbid>]
```


Parameters

`ci`

Displays information for stream replication streams.

`<ds>`

Specifies the name of the primary data server for the stream replication stream.

`<db>`

Specifies the name of the primary database for the stream replication stream.

`<dbid>`

Specifies the database ID.

Examples

Example 1

Displays information for a stream replication stream:

```
admin who, ci, QG30, tpcc
```

The returned result is:

Spid	State	Info	StreamState	ConsumerState
	Active	108 QG30.tpcc	up	NULL
StreamMode	PkgBufSize	NumPkgBufAlloc	PkgUsedCnt	PkgReleaseCnt
Async	1048576	50	0	80127
WaitForPkgCnt	EncodedCmds	BytesReceived	NumPkgsRead	
48022	1297951	1671582296	80127	
ReceivedWorkload	LoggedACK	ReceivedACK	BytesLogged	
1387200000	77745	79011	1671580896	

Usage

Table 30: Column Descriptions for `admin who, ci` Output

Column Name	Description
<code>StreamMode</code>	The current synchronous state: <ul style="list-style-type: none"> • Sync • Async

Column Name	Description
PkgBufSize	The value configured by <code>ci_package_size</code> .
NumPkgBufAlloc	The value configured by <code>ci_pool_size</code> .
PkgUsedCnt	The number of package buffers being used.
PkgReleaseCnt	The number of free package buffers.
WaitForPkgCnt	The number of times stream replication decoder waits for a free package buffer.
EncodedCmds	The number of commands encoded by this stream replication stream.
BytesReceived	The number of bytes received.
NumPkgsRead	The number of packages read by the Capture.
ReceivedWorkload	The workload of this stream replication stream.
LoggedACK	The number of acknowledgments sent to replication agent thread (RAT).
ReceivedACK	The number of acknowledgments received from RAT.
BytesLogged	The number of bytes this stream replication stream writes into SPQ.

12.4.2.7 admin version, ci

Displays information about the current version of the Component Interface (CI) in use by the SAP Replication Server. If you intend to use stream replication in the Replication Server, CI must be enabled. This command does not accept any additional parameters.

Syntax

```
admin version, ci
```

Parameters

ci

Displays version information of stream replication streams. Replication Server works even without a CI (without the function of stream replication), but under this scenario, running the command does nothing.

Examples

Example 1

Display information about the current version of the CI:

```
admin version, ci
go
```

The returned result is:

```
CI Library
Version

-----
SAP CI-Library/15.7.1/EBF 26750 SP306 rs1571sp306/CI 1.7.1/Linux AMD64/Linux
2.6.18-164.el5 x86_64/1/DEBUG64/Thu Jan 5 22:23:09
2017

-----
Info                Negotiated Version
-----
mo0_13339.tdb1      1.7
mo0_13339.tdb2      1.7
-----
```

12.4.3 Starting and Stopping Capture

Replication Server provides commands to start and stop the Capture.

12.4.3.1 suspend capture

Stops Capture and the SPQ reader.

Syntax

Suspend Capture:

```
suspend capture {<dsname>.<dbname> | all}
```

Suspend Capture and disable replication:

```
suspend capture <dsname.dbname> [disable replication]
```

Parameters

<dsname>

Specifies the name of the primary data server for the Capture.

<dbname>

Specifies the name of the primary database for the Capture.

all

Stops all Captures.

disable replication

Disables replication in an HADR environment.

Usage

- The `suspend capture` command has no impact on the connection between Replication Agent and Replication Server. If the log transfer is not suspended and the Replication Agent is active, it can still send data to SPQ after suspending the Capture.
- When you use `<dsname.dbname>`, the command stops the active path from replicating downstream in the HADR replication path. This, while keeping the replicate data from the SAP ASE Replication Agent to an external system.
- When only `<dbname>` is used: The command stops the SPQ Agent thread in the active Replication Server and disables its truncation point in SPQ such that the SPQ can be truncated properly.
- If you use this command with the `disable replication` option to disable replication in an HADR environment, data loss may occur, and you have to re-materialize when the capture starts again.

12.4.3.2 resume capture

Starts Capture and SPQ reader.

Syntax

```
resume capture {<dsname>.<dbname> | all}
```

Parameters

<dsname>

Specifies the name of the primary data server for the Capture.

<dsname>

Specifies the name of the primary database for the Capture.

all

Starts all Captures.

12.4.4 Managing SPQ

Replication Server provides commands to manage the SPQ.

12.4.4.1 SPQ Data Retention for a Configurable Time

During replication, the Replication Server receives data from the Replication Agent, and writes it to the SPQ file. If the Replication Server encounters any error, it can refer to the SPQ file to retrieve the lost data. However, if the SPQ file data is removed immediately, the Replication Server will not be able to authenticate the validity of data, and it results into erroneous or inadequate information. In order to prevent the immediate removal of data in the SPQ file, you can use the `spq_save_interval` parameter in the `alter connection` command.

The `spq_save_interval` parameter allows you to configure the data retention time (in minutes). This value indicates the amount of time that the system waits before removing or truncating an outdated data or data set.

The default value of the `spq_save_interval` parameter is 0 minute. The max value is 525600 minutes.

An example of the `spq_save_interval` parameter is as follows:

Sample Code

```
alter connection to repl0_52544.tdb1  
  
set spq_save_interval to '10'
```

Where:

- `repl0_52544` is the name of the server.
- `tdb1` is the name of the database.
- `10` is the value of data retention period (in minutes).

12.4.4.2 spq_dump_queue

The `spq_dump_queue` command enhances the diagnostic ability of Replication Server in the HADR environment. It provides the vital information that you can use to investigate the root cause of data loss in your HADR environment. With this information, a remedial action can be taken if necessary.

The `spq_dump_queue` command interprets the raw binary data (stored in the SPQ file) into a more comprehensible, user-consumable form. The command is used with an existing command (`dump_file`) to dump the data into a file for debugging purposes. After you have specified the dump file location using the `dump_file` command, the SPQ dump operation triggered by the `spq_dump_queue` command writes the content to this specified file. You can use this output to identify the lost data, and then proceed with the subsequent troubleshooting steps. Note that if you have not specified the file location before using the `spq_dump_queue` command, the dump file is saved in the SPQ file directory itself.

Before you execute the dump operation, ensure that:

- The stream replication mode is enabled.
- Capture is suspended.
- The log transfer is not suspended.

Typically, the following commands should be executed to dump the SPQ queue:

- Specify the dump file by using the following commands:
 - Execute the `touch` command on the dump file.
 - Thereafter execute the `sysadmin <dump_file>, <"dump_file_path">` command.
- Dump the SPQ queue.

Syntax

```
sysadmin spq_dump_queue, <dsname>, <dbname> [[, <filename> [, <number>]] | [, <begin_time>, <end_time>]]
```

Parameters

<filename>

specifies the SPQ file.

<number>

indicates the number of messages from the SPQ file that you want to include in the data dump. If this value is not specified, the Replication Server selects up to the last message in the SPQ file when the command is executed.

<begin_time>

indicates the begin timestamp. If you use this option, all the SPQ files that are updated after the timestamp, are considered as the dumping target. It must be used with an end timestamp that is specified by the `end_time` parameter. A special value `NULL` can be used to indicate `<anytime>`.

<end_time>

indicates the end timestamp. If you use this option, all the SPQ files that are updated before the timestamp are considered as the dumping target. It must be used with the `begin_time` parameter. Similarly, `NULL` can be used to indicate <anytime>.

Examples

Example 1

Dumps all SPQ files that are not yet truncated:

```
1> sysadmin spq_dump_queue, rqac0_9979, tdb1, 'NULL', 'NULL'
2> go
Dumped SPQ for 'rqac0_9979.tdb1' to file '/ramcache/quasr/s29572318/rep_se0/
sqp_folder/SPQ_1447047299496877063.dump'
```

Example 2

Dumps all the SPQ files that are updated before 11/07/15 08:00:00:

```
1> sysadmin spq_dump_queue, rqac0_9979, tdb1, 'NULL', '11/07/15 08:00:00'
2> go
Dumped SPQ for 'rqac0_9979.tdb1' to file '/ramcache/quasr/s29572318/rep_se0/
sqp_folder/SPQ_1447048790472164917.dump'
```

Example 3

Dumps the data in the specific SPQ file called `SPQ_2.dat`:

```
1> sysadmin spq_dump_queue, rqac0_9979, tdb1, "/ramcache/quasr/s29572318/
rep_se0/sqp_folder/SPQ_2.dat", 10
2> go
Dumped SPQ for 'rqac0_9979.tdb1' to file '/ramcache/quasr/s29572318/rep_se0/
sqp_folder/SPQ_144704895342054283.dump'
```

Example 4

Dumps the first 10 messages in the specific SPQ file called `SPQ_2.dat`:

```
1> sysadmin spq_dump_queue, rqac0_9979, tdb1, "/ramcache/quasr/s29572318/
rep_se0/sqp_folder/SPQ_2.dat", 10
2> go
Dumped SPQ for 'rqac0_9979.tdb1' to file '/ramcache/quasr/s29572318/rep_se0/
sqp_folder/SPQ_1447049022408175347.dump'
```

i Note

In all the examples, since a dump file is not specified, each dump is generated into a single random file in the SPQ file directory.

12.4.4.3 sysadmin issue_ticket

Inserts a ticket to the SPQ. Use this command to flush a replication path when SAP ASE is down.

Syntax

```
sysadmin issue_ticket {, <dbid> | {, <ds>, <db>}}, <q_type>, h1[, h2[, h3[, h4 ]]]
```

Parameters

<dbid>

Specifies the database ID.

<ds>

Specifies the name of the primary data server for the SPQ.

<db>

Specifies the name of the primary database for the SPQ.

<q_type>

Specifies the type of queue. For SPQ, <q_type> must be 2. If <q_type> is 0, 1, or not provided, the ticket is inserted into IBQ or OBQ.

Examples

Example 1

```
sysadmin issue_ticket, QG20, tpcc, 2, 'test1'
```


12.4.4.4 sysadmin spq_force_send_all

Forces the SPQ to send unconfirmed commands (NC-CMDs) to Capture. When SAP ASE is down, by using this command, log transfer from the specified data server and database is suspended automatically, DSI to primary database is suspended, and outbound queue for primary database is purged.

Syntax

```
sysadmin spq_force_send_all, <ds>, <db>
```

Parameters

<ds>

Specifies the name of the primary data server for the SPQ.

<db>

Specifies the name of the primary database for the SPQ.

12.4.5 Managing SPQ Agent

The commands in this section enable you to manage the SPQ Agent, and are issued by the system.

To run any of these commands, simply execute the corresponding commands for Replication Management Agent. When you do this, the system internally triggers SAP Replication Server commands when the corresponding Replication Management Agent commands are executed.

Note

Do not execute these SAP Replication Server commands by themselves.

12.4.5.1 configure spq_agent enable

Configures an SPQ Agent while creating a connection from HADR.

Syntax

```
configure spq_agent <dbname> enable
```

```
<external_rs_host:external_rs_port> <ext_dataservername> <active_dataservername>  
<spq_agent_username> <spq_agent_password> <maintuser_name>
```

Parameters

<dbname>

Specifies the database name of the external connection which the SPQ Agent uses for the `connect source` command.

<external_rs_host:external_rs_port>

Specifies the connection information of the external Replication Server.

<ext_dsname>

Specifies the data server name of the external connection which the SPQ Agent will use for the `connect source` command.

<active_dsname>

Specifies the current data server name used by the Replication Agent thread for SAP ASE. It is retrieved by using the following command:

```
select rep_agent_config (1, 'config', 'connect dataserver')
```

<spq_agent_user>

Specifies the user that SPQ Agent uses to connect to the external Replication Server.

<spq_agent_passwd>

Specifies the password that SPQ Agent uses to connect to the external Replication Server.

<maintuser_name>

The maintenance user used by the external Replication Server. Transactions from the maintenance user will be filtered by the external Replication Server.

Usage

The `configure spq_agent enable` command:

- Is issued in the active Replication Server inside HADR, and for the command to succeed, all Replication Servers in the HADR domain must be up.
- Adds all connections for the database into the `rs_spqratgroup` table.
- Is triggered by the `create connection` and `alter connection` command to configure the SPQ Agent.
- Marks the current active connection as the active member of the SPQ Agent group in the `rs_spqratstate` table.
- Adds all SPQ Agent related configurations into the `rs_spqratcfg` table.

Permissions

Requires `sa` or `manage_spq_agent` permissions.

12.4.5.2 configure spq_agent disable

Disables SPQ Agent.

Syntax

```
configure spq_agent <dbname> disable
```

Parameters

<dbname>

Name of the database of the external connection that the SPQ Agent uses for the `connect source` command.

Usage

- Executing this command deletes all related rows in the tables `rs_spqratgroup`, `rs_spqratstate`, and `rs_spqratcfg`, and SPQ can be truncated without considering external replication.
- The SPQ Agent must be stopped before this command is issued.

Permissions

Requires `sa` or `manage_spq_agent` permissions.

12.4.5.3 configure spq_agent set

Alters the configuration value used by the SPQ Agent. Execute this command in Replication Servers within an HADR system that are either on standby or in an active site.

Syntax

```
configure spq_agent <dbname> set <opname> to '<value>'
```

Parameters

<dbname>

Specifies the name of the database of the external connection that the SPQ Agent uses for the `connect source` command.

<opname>

Specifies the name of configuration. The configurations that can be altered are: `connect_dataserver`, `connect_database`, `rs_servername`, and `maint_username`.

'<value>'

The value to set the configuration to.

Examples

Example 1

Sets the `rs_servername` configuration to the `'mo-d558c2551:9885'` value.

```
configure spq_agent ERP set rs_servername to 'mo-d558c2551:9885'
```

Example 2

Sets the `rs_servername` configuration to the `'mo-d558c2551:9885 ssl="CN=ERS"'` value. `ssl="CN=ERS"` indicates that SSL is enabled for the SAP Replication Server.

```
configure spq_agent ERP set rs_servername to 'mo-d558c2551:9885 ssl="CN=ERS"'
```

Usage

- For a new configuration to take effect, restart the SPQ Agent.

Permissions

Requires `sa OR manage spq_agent` permissions.

12.4.5.4 resume spq_agent

Starts the SPQ Agent thread for external replication. Issue it within any Replication Server inside HADR.

Syntax

```
resume spq_agent <dbname>
```

Parameters

<dbname>

Specifies the name of the database of the external connection that the SPQ Agent uses for the `connect source` command.

Usage

- If the current Replication Server does not control the active SPQ Agent, it uses the same credentials to log in to the controlling Replication Server, and forwards this command. After you execute this command, the SPQ Agent on the active member (specified in the `rs_spqratstate.active_member` table) of the group starts.
- To execute this command, the SPQ Agent must be enabled, but not already started.

Permissions

Requires `sa` or `manage spq_agent` permissions.

12.4.5.5 suspend spq_agent

Stops the SPQ Agent thread for external replication. Issue it within any Replication Server inside HADR.

Syntax

```
suspend spq_agent <dbname> [disable replication]
```

Parameters

<dbname>

Specifies the name of the database of the external connection that the SPQ Agent uses for the `connect source` command.

<disable replication>

This option disables replication to an external replication system, stops the SPQ Agent immediately, and disables the truncation point.

Usage

- If the current Replication Server does not control the active SPQ Agent, it uses the same credentials to log in to the controlling Replication Server, and forwards this command. After you execute this command, the `state` column in the `rs_spqratstate` table is set to 0x03, meaning enabled and suspended.
- If you specify the `<disable replication>` option, the truncation point of SPQ Agent is disabled from its SPQ, so that you can truncate SPQ without considering external replication. Because this may cause you to experience data loss, rematerialize when you resume the SPQ agent.
- A suspended SPQ Agent does not restart automatically after Replication Server restarts.
- This is a synchronized command, which means it will wait and return success after the SPQ Agent thread finally exits, or returns immediately for any error that occurs.
- You do not see a success message from this command until the SPQ Agent thread completely exits. If any errors occur, it displays a return immediately.

Permissions

Requires `sa` or `manage spq_agent` permissions.

12.4.5.6 activate all spq_agent

Starts all local SPQ Agents from the same SAP ASE, and requires all other SPQ Agents for this database to be in an inactive, or a draining state.

Syntax

```
activate all spq_agent from <servername> [with force]
```

Parameters

<servername>

The primary SAP ASE HADR where data is replicating from.

[with force]

Gives you a way to force switch when the previous active Replication Server is down and the data in the previous SPQ can be discarded.

Usage

When you do not specify the `with force` option:

- If the SPQ Agent for the previous HADR active site is in a draining state, the local SPQ Agent is put in a waiting state. It is then put in an active state when the previous agent has drained completely and become inactive. As long as it has not been previously explicitly suspended, the SPQ Agent thread starts after all other members are in an inactive state.
- The command returns with no change if one of local members is already in active or waiting state.
- If the local member is in an inactive state, it will sync up with the other Replication Server and experiences one the following:
 - Enters into a waiting state if the current active member is in a draining state
 - Enters into a active state if all other members are in an inactive state
 - Fails if the current active member is in an active or a waiting state.
- If the local member is in a draining sate, the command returns with no change and tells the administrator to issue the command again after the drain has completed.

When you specify the `with force` option:

- If the local SPQ Agent has not been previously explicitly suspended and if all other Replication Servers are down, the agent starts immediately, and local SPQ Agents are moved into active state. If all other Replication Servers are up, the command checks to see whether other members are in an inactive state before it starts.
- The behavior is the same as if you do not specify `with force` if all the Replication Servers in the HADR domain that are in an active, draining, or waiting state are up.
- The behavior is the same as if you do not specify `with force` if all the Replication Servers in the HADR domain that are in an active, draining, or waiting state are down, with the exception of the server that is down. Such a server is not considered.

Permissions

Requires `sa` or `manage spq_agent` permissions.

12.4.5.7 deactivate all spq_agent

Stops all local SPQ Agents either immediately, or after data is drained, putting the local SPQ Agents in an inactive or drained state depending on the option used. This command plays a vital role in failover scenarios.

Syntax

```
deactivate all spq_agent from <servername>  
[after drain | disable replication]
```

Parameters

<servername>

The current, active HADR SAP ASE server where data is replicating from.

after drain

Table 31: `after drain` parameter usage

SPQ Agent Condition	Action
Already inactive	Do nothing.
In draining state	If you: <ul style="list-style-type: none">• Specify <code>after drain</code> — do nothing.• Do not specify <code>after drain</code> — deactivate immediately; the SPQ Agent changes to inactive state.
If another SPQ Agent in the same group is in a waiting state or inactive state	If you: <ul style="list-style-type: none">• Specify <code>after drain</code> — fail the command.• Do not specify <code>after drain</code> — deactivate immediately; the SPQ Agent changes to inactive state.

SPQ Agent Condition	Action
In active state	If you: <ul style="list-style-type: none"> Specify <code>after drain</code> – set to draining state. After SPQ Agent drains all data to the external system, the state changes to drained. Do not specify <code>after drain</code> – deactivate immediately; the SPQ Agent changes to inactive state.
Is a standby member	The command fails.

disable replication

Immediately stops all local SPQ Agents without draining data. Also disables the truncation point for each SPQ Agent from its SPQ. This means that database re-materialization is needed if the SPQ Agent is re-activated later.

Permissions

Requires `sa` or `manage_spq_agent` permissions.

12.4.6 Managing External Replication Connections

Replication Server provides commands to create a connection to an HADR database in an external Replication Server, and to change the attributes of a database connection.

12.4.6.1 create connection

Creates a connection to an HADR database in an external Replication Server.

Syntax

```
create connection to <data_server>.<database>
set error class [to] <error_class>
set function string class [to] <function_class>
set username [to] <user>
[set password [to] <passwd>]
[set replication server error class [to] <rs_error_class>]
[set database_param [to] '<value>' [set database_param [to] '<value>']...]
```

```
[set security_param [to] '<value>' [set security_param [to] '<value>']...]
[with log transfer on [, dsi_suspended]]
[with primary only]
[connect to hadr [set spq_agent_username [to] '<value>'
set spq_agent_password [to] '<value>']]
```

Parameters

<data_server>

Specifies the data server that holds the database to be added to the replication system.

<database>

Specifies the database to be added to the replication system.

<error_class>

Specifies the error class that is to handle errors for the database.

<function_class>

Specifies the function string class to be used for operations in the database.

set username [to] <user>

Specifies the username that the external Replication Server uses to log onto the active Replication Server, and to HADR-enabled SAP ASE, and which meets these requirements:

- Has valid credentials in HADR ASE with a replication_role
- Has valid credentials in the <dbname> database
- Does not have sa_role and sso_role (to guarantee a login to the active ASE)
- Cannot be a DR_admin nor DR_maint user, that are internally used by the HADR system.
- Exists in the HADR Replication Server as a normal user with the “manage spq_agent” permission.

set password [to] <passwd>

Specifies the password for <user>.

set replication server error class [to] <rs_error_class>

Specifies the error class that handles Replication Server errors for a database. The default error class is <rs_repserver_error_class>.

set database_param [to] '<value>'

Allows you to specify a value that affects database connections from the Replication Server.

set security_param [to] '<value>'

Allows you to specify a value that affects network-based security.

with log transfer on

Indicates that the connection may be a primary data source or the source of replicated functions. When you specify this clause, RepAgent or SPQ Agent creates an inbound queue and is prepared to accept a RepAgent connection. These commands are sent to the primary Replication Server for distribution and replication. RepAgent also

coordinates database log truncation with the Adaptive Server and the primary Replication Server.

dsi_suspended

Starts the connection with the DSI thread suspended. You can resume the DSI later. This option is useful if you are connecting to a non-SAP data server that does not support Replication Server connections.

with primary only

Use this clause when you only intend to replicate out of HADR. When this clause is used, the DSI of this connection is disabled, and the outbound queue, SQM, and DSI threads are not started.

connect to hadr

With this clause, HADR is added in a replication environment as a primary SAP ASE or a target SAP ASE. It creates a primary connection replicating from HADR, or creates a target connection replicating into HADR. When used to add HADR to a primary SAP ASE, the username and password of the SPQ Agent must be specified at the same time. When used to add HADR as a target SAP ASE, it is not necessary to specify the username and password (`spq_agent_username` and `spq_agent_password`) of the SPQ Agent.

set spq_agent_username [to] '<value>'

Specifies the username for the credential SPQ Agent uses to connect to an external Replication Server to replicate data. Both this username and its password must already exist in an external Replication Server with a "connect source" permission.

set spq_agent_password [to] '<value>'

Specifies the password for `<set spq_agent_username>`.

Examples

Example 1

Create a primary-only connection from HADR:

```
create connection to HADR_ASE.ERP
set error class to rs_sqlserver_error_class
set function string class to rs_sqlserver_function_class
set username to ext_maint_user
set password to Sybase123
with primary only
connect to hadr
set spq_agent_username to spq_agent_user
set spq_agent_password to Sybase123
```

Example 2

Create a connection to an HADR:

```
create connection to HADR_ASE.ERP
set error class to rs_sqlserver_error_class
set function string class to rs_sqlserver_function_class
set username to ext_maint_user
set password to Sybase123
```

```
connect to hadr
```

Usage

Executing `create connection` prompts the following actions:

- Connects to the ASE server as the maintenance user, and issues the following commands to obtain the name of active Replication Server and the current data server name, with the results displaying in `<host>:<port>` format:

```
select rep_agent_config(1, 'config', 'rs servername')
select rep_agent_config(1, 'config', 'connect dataserver')
```

- Connects to the active Replication Server as the maintenance user and issues following command to enable SPQ Agent:

```
configure spq_agent <dbname> enable ...
```

- Configures the SPQ Agent by:
 - copying required configurations from the SAP ASE Replication Agent by using the following command:

```
select rep_agent_config(...)
```

- adding appropriate configuration to connect it to an external Replication Server (`host:port`).
- Ensures that the active Replication Server will connect to the other Replication Server to enable and sync up the state of the SPQ Agent
- Starts the SPQ Agent thread by sending the following command to the active Replication Server:

```
resume spq_agent <dbname>
```

- SPQ Agent connects to the external Replication Server and starts replication.
- Introduces a new Replication Server system table to store configuration data of the SPQ Agent. The structure of the table is as under:

Table 32:

SPQ Agent Config	ASE RAT Config	Value
<code>rs_servername</code>	<code>rs servername</code>	host:port of the external Replication Server
<code>rs_username</code>	<code>rs username</code>	<code>spq_agent username</code>
<code>connect_dataserver</code>	<code>connect dataserver</code>	The name of the dataserver known to the external system as part of the server name in the external connection <code><server_name>.<dbname></code>

SPQ Agent Config	ASE RAT Config	Value
connect_database	connect_database	The name of the database known to the external system as part of the dbname in the external connection <server_name>.<dbname>

12.4.6.2 alter connection

Changes the attributes of a database connection.

Syntax

```
alter connection to <data_server>.<database> {
[for replicate table named [<table_owner>.]<table_name>
[set table_param [to] '<value>' ] ] |
set function_string class [to] <function_class> |
set error class [to] <error_class> |
set replication server error class [to] <rs_error_class> |
set password [to] <passwd> |
set log transfer [to] {on | off} [set spq_agent_username [to] '<value>'
set spq_agent_password [to] '<value>'] |
set database_param [to] '<value>' |
set security_param [to] '<value>' |
set security_services [to] <default>]
set dataserver and database name [to] <new_ds>.<new_db> |
set trace [to] '<value>' |
not connect to hadr |
connect to hadr [set spq_agent_username [to] '<value>'
set spq_agent_password [to] '<value>']}]
```

Parameters

<data_server>

Specifies the data server that holds the database to be added to the replication system.

<database>

Specifies the database to be added to the replication system.

for replicate table named [<table_owner>.]<table_name>

Specifies the name and owner of the table at the replicate database. <table_owner> is an optional qualifier for the table name, representing the table owner. <table_name> is the name of the table at the replicate database, and can be up to 200 characters long.

set table_param [to] '<value>'

Specifies the table-level parameter that affects a table you specify with the `for replicate table name` clause.

set function string class [to] <function_class>

Specifies the function string class to be used for operations in the database.

set error class [to] <error_class>

Specifies the error class that is to handle errors for the database.

set replication server error class [to] <rs_error_class>

Specifies the error class that handles Replication Server errors for a database. The default is `rs_repserver_error_class`

set password [to] <passwd>

Use this clause to modify the password (if needed) for the maintenance user, which was specified in the `create connection` command, by using the clauses `set password [to] <passwd>`, and `set username [to] <user>`. Note that `<user>` (the username assigned to the maintenance user) is not a changeable configuration in `alter connection`.

set log transfer [to] {on | off}

Indicates that the connection may be a primary data source or the source of replicated functions. When you specify this clause, RepAgent or SPQ Agent creates an inbound queue and is prepared to accept a RepAgent connection. These commands are sent to the primary Replication Server for distribution and replication. RepAgent also coordinates database log truncation with the Adaptive Server and the primary Replication Server.

set spq_agent_username [to] '<value>'

Specifies the username for the credential SPQ Agent uses to connect to an external Replication Server to replicate data. Both this username and its password must already exist in an external Replication Server with a "connect source" permission.

set spq_agent_password [to] '<value>'

Specifies the password for `<spq_agent_username>`.

set database_param [to] '<value>'

Allows you to specify a value that affects database connections from the Replication Server.

set security_param [to] '<value>'

Allows you to specify a value that affects network-based security for connections. This parameter does not apply to non-ASE connectors.

set security_param [to] '<value>'

Allows you to specify a value that affects network-based security for connections. This parameter does not apply to non-ASE connectors.

set security_services [to] <default>

Resets all network-based security features for the connection to match the global settings of your Replication Server. This parameter does not apply to non-ASE connectors.

set dataserver and database name [to] <new_ds>.<new_db>

`<new_ds>.<new_db>` specify the names of the new database and dataserver.

set trace [to] '<value>'

Allows ExpressConnect tracing at the DSI level. '<value>' represents a character string containing a new value for the option. The syntax for '<value>' is `econn`, `condition`, `[on|off]`, where:

- `<econn>` – This is the valid module type.
- `<condition>` – Specifies the trace condition to set.
- `on | off` – Allows you to enable or disable the condition.

Note

If you do not specify '<value>', the empty string disables ExpressConnect tracing values after the connection, or after Replication Server restarts. For example:

```
alter connection to data_server.database
set trace to ''
```

An empty string disables ExpressConnect tracing values after the connection or when the Replication Server is restarted.

not connect to hadr

Drops HADR from the primary or target SAP ASE in a replication system.

connect to hadr

With this clause, HADR can be added to a primary SAP ASE or a target SAP ASE. It converts a normal primary connection to a connection replicating from HADR, or converts a normal target connection to a connection replicating into HADR. When used to add HADR to a primary SAP ASE, the username and password of the SPQ Agent must be specified at the same time. When used to add HADR to a target SAP ASE, the SPQ Agent username and password are not necessary.

Examples

Example 1

Convert a normal primary connection to a connection replicating from HADR:

```
alter connection to repl1_17149.ERP
connect to hadr
set spq_agent_username to spq_agent_user
set spq_agent_password to Sybase123
```

Example 2

Convert a normal DSI connection to a connection replicating into HADR:

```
alter connection to repl1_17149.ERP
connect to hadr
```

12.5 Replication Agent Thread for SAP ASE Commands

Replication Agent Thread for SAP ASE (RepAgent for short) provides some commands and parameters that you can use in an HADR system.

12.5.1 RepAgent Configuration Parameters

Use `sp_config_rep_agent` configuration parameters to configure and control RepAgent behavior.

This topic lists the RepAgent configuration parameters that apply to stream replication only. For more other RepAgent configuration parameters, refer to *SAP Replication Server Reference Manual*.

RepAgent Configuration Parameter	Description
<code>'buffer pool size', {'<buffer pool size value>'}</code>	<p>Specifies the maximum size of the buffer pool. Specifies the number of buffers (packages) stream replication can allocate on startup.</p> <p>Value: 1 to 2,147,483,647</p> <p>Default: 8</p>
<code>'initial log scan percent', {'<initial log scan percent value>'}</code>	<p>A dynamic configuration parameter, this specifies the percentage of the initial log that RepAgent should scan before it evaluates whether to slow down user tasks. This also applies to every multiple of the percentage until RepAgent scans all the initial log.</p> <p>Value: 1 to 100</p> <p>Default: 100</p>
<code>'max commands per package', {'<max commands per package value>'}</code>	<p>Specifies the maximum number of commands that can be put in a stream replication package.</p> <p>Value: 1 to 1,000,000</p> <p>Default: 80</p>

RepAgent Configuration Parameter

Description

```
'max commit wait', {'<max commit wait value>'}
```

Specifies the maximum amount of time in microseconds that a user task committing a transaction waits for acknowledgment of the commit from SAP Replication Server. When the commit time expires because the user task did not receive acknowledgment, the task makes additional calculations to determine if it needs to request RepAgent to switch to asynchronous replication mode to allow the application to proceed. The additional logic is based on other configuration parameters. See `peak transaction threshold` and `peak transaction timer` for more information.

The `max commit wait` parameter only applies when `stream mode` is set to `sync` or `near_sync` for the RepAgent of the database serving replication.

The `max commit wait` parameter is dynamic and any change to the parameter takes effect without having to restart RepAgent.

Range of valid values: 1 to the value of `MAXINT` which is 2,147,483,647 microseconds.

Default: 10,000,000 microseconds (10 seconds)

Note

If the value is set to zero, all user tasks committing transactions wait indefinitely for acknowledgment.

You must consider two competing factors when you specify a value for `max commit wait`:

- Zero data loss, and
- Application response time

Specifying a:

- High value – increases the possibility all data has been distributed to SAP Replication Server if SAP ASE fails. If all commits have been acknowledged, then no loss of data occurs if a failover to a standby site is performed. However, a high value may decrease the application response time if the acknowledgment from SAP Replication Server is slow.
- Low value – A low value may result in frequent switches to asynchronous mode. In asynchronous mode, the primary and the standby databases are temporally out of synchronization. Therefore, a failover performed in asynchronous mode will result in the loss of data at the standby site.

RepAgent Configuration Parameter	Description
'max stream retry', {'<max stream retry value>'}	<p>However, a low value reduces the maximum database response time an application may experience while waiting for commit processing to complete.</p> <p>Specifies how often RepAgent retries to set up a connection. RepAgent shuts down once the configured value is reached.</p> <p>Value: -1 or 2,147,483,647</p> <p>Default: -1 (means never)</p>
'max user task slowdown', {'<max user task slowdown value>'}	<p>A dynamic configuration parameter, this specifies the maximum amount of time in milliseconds for the slowdown imposed on user tasks for replication to switch to sync mode at RepAgent startup.</p> <p>Value: 0 to 3000 milliseconds</p> <p>Default: 3000 milliseconds</p>
'stream buffer size', {'<stream buffer size value>'}	<p>Specifies the size of a stream replication buffer (package) in bytes. Each stream replication package in the stream replication buffer pool shares the same size specified by <code>stream buffer size</code>.</p> <p>Value: 1 to 2,147,483,647</p> <p>Default: 1,048,576 (bytes)</p>
'stream mode', {'<stream mode value>'}	<p>Specifies the replication synchronization mode between RepAgent and SAP Replication Server. The database must be configured for stream replication for this option to take effect.</p> <p>Value: sync, async, or near sync.</p> <p>Default: sync.</p>
'stream replication', {'<stream replication value>'}	<p>Specifies the communication protocol between RepAgent and SAP Replication Server. The parameters must be set on both Replication Agent and Replication Server and agree with each other. If the values are different, the connection is refused.</p> <p>Set <code>stream_replication</code> to true to use stream replication to distribute replicated data to SAP Replication Server.</p> <p>Value: true or false</p> <p>Default: false</p>

RepAgent Configuration Parameter	Description
<pre>'peak transaction threshold'[, '<peak_transaction_threshold>']</pre>	<p>Specifies the maximum number that a global counter can reach before RepAgent switches from synchronous or near-synchronous mode to asynchronous mode. The global counter increases by one when a task has an average commit wait time that is greater than the configured maximum commit wait time. When the global counter reaches the specified peak transaction threshold, the task requests RepAgent to switch the stream mode.</p> <p>The <code>peak transaction threshold</code> parameter applies when stream replication is set to true.</p> <p>The <code>peak transaction threshold</code> parameter is dynamic and any change to the parameter takes effect without having to restart RepAgent.</p> <p>Valid values: 1 to the value of MAXINT which is 2,147,483,647</p> <p>Default: 5</p>
<pre>'peak transaction timer'[, '<peak_transaction_timer>']</pre>	<p>Sets the amount of time, in seconds, that RepAgent waits before resetting the global counter to zero. The global counter records the number of times when a task commit wait time exceeds the configured maximum commit wait time. The timer restarts again after the global counter is reset to zero. Use the peak transaction timer to avoid a mode switch caused by accumulated spikes in the average commit wait time during a long period of time.</p> <p>The <code>peak transaction timer</code> parameter applies when stream replication is set to true.</p> <p>The <code>peak transaction timer</code> parameter is dynamic and any change to the parameter takes effect without having to restart RepAgent.</p> <p>Valid values: 1 to the value of MAXINT which is 2,147,483,647 seconds</p> <p>Default: 300 seconds</p>

12.5.2 Output from sp_help_rep_agent

When RepAgent is configured for stream replication to support synchronous replication in an HADR system, the output for the `process` parameter in `sp_help_rep_agent` shows additional status information about the Coordinator, Scanner, and Secondary Truncation Point Manager processes. This output differs from the output when the replication mode is through log transfer language (LTL).

However, when RepAgent is configured for stream replication, the output for other parameters in `sp_help_rep_agent` matches the output when the replication mode is through log transfer language (LTL).

Refer to the output for other parameters in `sp_help_rep_agent` in the *SAP Replication Server Reference Manual*.

Table 33: Column Descriptions for Output from `sp_help_rep_agent` with 'process' during Stream Replication

Column	Description
<code><dbname></code>	The name of the database for which you are querying process information.
<code><pathname></code>	The name of the replication path associated with each sender or scanner process if you configure multiple replication paths and scanners (multi-path Replication only).
<code><spid></code>	The system process ID of a process in the dataserver. For a Multithreaded RepAgent – <code>spid</code> identifies the coordinator task if you enable multiple scanners.
<code><scanner_spid></code>	The system process ID of each scanner process in the dataserver.
<code><sender_spid></code>	The system process ID of each sender process in the dataserver (not applicable for Stream Replication).
<code><start marker></code>	Identifies the first log record scanned in current batch.
<code><end marker></code>	Identifies the last log record to be scanned in current batch.
<code><current marker></code>	Identifies the current log record being scanned.
<code><trunc_pts_confirmed></code>	The number of confirmed truncation points. A confirmed locator is received for replicated log operations that were actually written to disk.
<code><trunc_pts_processed></code>	The number of truncation points processed. For example, the number of times the Secondary Truncation Point could be moved in the primary database.
<code><total_processing_time></code>	The total processing time for all truncation points.
<code><longest_process_time></code>	The longest time spent on processing the truncation points.
<code><sleep_status></code> <code><state></code>	See the <code><sleep_status></code> and <code><state></code> columns together for the status of the coordinator, scanner, and secondary truncation point manager. <ul style="list-style-type: none"> • RepAgent Coordinator Process Status During Stream Replication [page 629] • RepAgent Scanner Process Status During Stream Replication [page 629] • RepAgent Secondary Truncation Point Manager Process Status During Stream Replication [page 629]
<code>scanner_type</code>	Indicates the type of scanner. RepAgent supports the following two types of scanners when you enable in-memory row storage (IMRS) on a database. <ul style="list-style-type: none"> • <code>syslogs_scanner</code> • <code>sysimrslogs_scanner</code> – not available for non-IMRS databases.

When stream replication mode is enabled, you need to look at values for both the columns, `<sleep_status>` and `<state>`, to determine the thread and process status of the RepAgent Coordinator, Scanner, and Secondary Truncation Point Manager. Some of the statuses are tagged "rare" because they may appear very

briefly as the process task moves from one common state to another and are therefore extremely unlikely to be seen.

Table 34: RepAgent Coordinator Process Status During Stream Replication

<code><state></code>	<code><sleep status></code>	Description
starting	not sleeping	Coordinator doing its own startup.
spawning	not sleeping	Coordinator is starting scanner or the Secondary Truncation Point Manager task.
spawning	sleep on task spawn	Coordinator is waiting on spawned task to indicate it has started.
running	updating mode	Coordinator is servicing a mode update.
running	not sleeping	Coordinator is otherwise working (rare).
sleeping	sleeping	Coordinator is waiting on alarm or event.
stopping	sleep on task terminate	Coordinator is waiting for scanner and Secondary Truncation Point Manager to stop.
stopping	not sleeping	Coordinator is shutting down (rare).

Table 35: RepAgent Scanner Process Status During Stream Replication

<code><state></code>	<code><sleep status></code>	Description
starting	not sleeping	Scanner is doing startup work (rare).
starting	opening stream	Scanner is sleeping, waiting on stream to open.
running	not sleeping	Scanner is working but not scanning (rare).
scanning	not sleeping	Scanner is actively scanning the log.
sleeping	stream sleep	Scanner is sleeping on a stream flush or package allocation.
sleeping	rewrite	Scanner is sleeping on a rewrite.
sleeping	sleeping	Scanner is sleeping on a dirty log.
sleeping	closing stream	Scanner is closing stream on shutdown.
sleeping	shutdown sync	Scanner is synchronizing with coordinator on shutdown (rare).
stopping	not sleeping	Scanner is shutting down (rare).

Table 36: RepAgent Secondary Truncation Point Manager Process Status During Stream Replication

<code><state></code>	<code><sleep status></code>	Description
starting	not sleeping	Secondary Truncation Point Manager is doing startup work (rare).
process trunc. point	not sleeping	Secondary Truncation Point Manager is processing a new truncation point.
running	not sleeping	Secondary Truncation Point Manager is working but not processing truncation points (rare).

<state>	<sleep status>	Description
sleeping	sleeping	Secondary Truncation Point Manager is waiting on an alarm or wakeup event.
stopping	not sleeping	Secondary Truncation Point Manager is stopping (rare).

12.5.3 sp_configure 'stream rep msg channel timeout'

Specifies the timeout value for the message channel from RepAgent to SAP Replication Server.

Syntax

```
sp_configure 'stream rep msg channel timeout', <stream_rep_msg_channel_timeout>
```

Parameters

<stream_rep_msg_channel_timeout> The maximum time, in seconds, that RepAgent waits for the response from SAP Replication Server.

Examples

Examples 1

This example sets the RepAgent message channel timeout to 100 seconds:

```
sp_configure 'stream rep msg channel timeout', 100
go
00:0002:00000:00041:2018/08/10 10:33:57.05 server Configuration file '/
marlinux5_work6/reneq/asecoronahadr_reneq_vu/release/SYBASE_16_SP03.cfg' has
been written and the previous version has been renamed to '/marlinux5_work6/
reneq/asecoronahadr_reneq_vu/release/SYBASE_16_SP03.000000003'.
00:0002:00000:00041:2018/08/10 10:33:57.07 server The configuration option
'stream rep msg channel timeout' has been changed by 'sa' from '60' to '100'.
Parameter Name          Config Value      Run Value      Unit      Default      Memory Used
-----
stream rep msg channel timeout 100              100           number    60           0
dynamic
```

Usage

- Use `sp_configure 'Rep Agent Thread administration'` to check the current timeout value.
- The default timeout value is 60 seconds. Valid range is 0 - MAXINT. If you set it to 0, it means no timeout.

Permissions

`sp_configure` requires “sa” or “sso” permission to modify configuration parameters.

Anyone can execute `sp_configure` to display information about parameters and their values.

Important Disclaimers and Legal Information

Hyperlinks

Some links are classified by an icon and/or a mouseover text. These links provide additional information.

About the icons:

- Links with the icon : You are entering a Web site that is not hosted by SAP. By using such links, you agree (unless expressly stated otherwise in your agreements with SAP) to this:
 - The content of the linked-to site is not SAP documentation. You may not infer any product claims against SAP based on this information.
 - SAP does not agree or disagree with the content on the linked-to site, nor does SAP warrant the availability and correctness. SAP shall not be liable for any damages caused by the use of such content unless damages have been caused by SAP's gross negligence or willful misconduct.
- Links with the icon : You are leaving the documentation for that particular SAP product or service and are entering a SAP-hosted Web site. By using such links, you agree that (unless expressly stated otherwise in your agreements with SAP) you may not infer any product claims against SAP based on this information.

Beta and Other Experimental Features

Experimental features are not part of the officially delivered scope that SAP guarantees for future releases. This means that experimental features may be changed by SAP at any time for any reason without notice. Experimental features are not for productive use. You may not demonstrate, test, examine, evaluate or otherwise use the experimental features in a live operating environment or with data that has not been sufficiently backed up.

The purpose of experimental features is to get feedback early on, allowing customers and partners to influence the future product accordingly. By providing your feedback (e.g. in the SAP Community), you accept that intellectual property rights of the contributions or derivative works shall remain the exclusive property of SAP.

Example Code

Any software coding and/or code snippets are examples. They are not for productive use. The example code is only intended to better explain and visualize the syntax and phrasing rules. SAP does not warrant the correctness and completeness of the example code. SAP shall not be liable for errors or damages caused by the use of example code unless damages have been caused by SAP's gross negligence or willful misconduct.

Gender-Related Language

We try not to use gender-specific word forms and formulations. As appropriate for context and readability, SAP may use masculine word forms to refer to all genders.

Videos Hosted on External Platforms

Some videos may point to third-party video hosting platforms. SAP cannot guarantee the future availability of videos stored on these platforms. Furthermore, any advertisements or other content hosted on these platforms (for example, suggested videos or by navigating to other videos hosted on the same site), are not within the control or responsibility of SAP.

© 2020 SAP SE or an SAP affiliate company. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP SE or an SAP affiliate company. The information contained herein may be changed without prior notice.

Some software products marketed by SAP SE and its distributors contain proprietary software components of other software vendors. National product specifications may vary.

These materials are provided by SAP SE or an SAP affiliate company for informational purposes only, without representation or warranty of any kind, and SAP or its affiliated companies shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP or SAP affiliate company products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP SE (or an SAP affiliate company) in Germany and other countries. All other product and service names mentioned are the trademarks of their respective companies.

Please see <https://www.sap.com/about/legal/trademark.html> for additional trademark information and notices.