



ADMINISTRATION GUIDE | PUBLIC

SAP Adaptive Server Enterprise 16.0 SP03

Document Version: 1.0 – 2020-03-04

HADR System with DR Node Users Guide

Content

- 1 Overview 4**
- 2 Adding a DR Node to an Existing Business Suite HA System. 6**
 - 2.1 Requirements for Adding a DR Node to an HA Environment. 6
 - 2.2 System Resource Requirements. 7
 - 2.3 DR Node Limitations 8
 - 2.4 Adding the Disaster Recovery Node. 8
 - Installing the Business Suite Application. 8
 - Installing the Data Movement Component. 9
 - Run setuphadr to Add the DR Node to the HADR System. 11
 - Sample setup_hadr.rs Response File for Business Suite. 13
 - 2.5 Post-Installation Tasks for Primary and Companion Servers. 18
- 3 Managing the HADR with DR Node System. 19**
 - 3.1 Determining the Mode and State of the DR Node. 19
 - 3.2 Monitoring Replication to the Disaster Recovery Node. 21
 - Monitoring Paths. 21
 - Monitoring Replication Backlog and Resource Usage. 27
 - Monitoring Replication Path Subcomponents. 29
 - 3.3 Managing Failover in the HADR Cluster with the DR Node. 32
 - Performing Planned Failover to the Standby Node. 33
 - Performing Unplanned Failover to the Standby Node. 39
 - Performing Failback to the Former Primary Node. 46
 - 3.4 Resynchronizing the DR Node. 48
 - Resynchronizing a Database on the DR Node. 48
 - Resynchronizing the DR Node. 50
 - 3.5 Resynchronizing the Standby Node. 53
 - Resynchronizing a Database on the Standby Node. 53
 - Resynchronizing the Standby Node. 55
 - 3.6 Recovering the HADR Cluster from the DR Node. 60
 - Recovering a Database in the HADR Cluster. 60
 - Recovering the HADR Cluster. 63
 - 3.7 Disabling Replication to the DR Node. 68
 - 3.8 Removing the DR Node from the HADR System. 69
 - 3.9 Performing a Rolling Upgrade on a DR Node System. 70
 - 3.10 Recovering Replication to DR Node When the Standby SAP Replication Server is Down. 75
 - Recovering Replication Before Primary Server is Full 77

	Recovering Replication After Primary Server is Full.	78
	Recovering Replication When You Cannot Restore the Standby Server.	80
3.11	SQL Statement Replication.	81
4	Troubleshooting.	82
4.1	Performing a Teardown.	82
	Performing a Teardown Using sap_teardown.	83
	Manually Removing the Replication in HADR System with DR Node Environment.	84
5	HADR Reference.	87
5.1	RMA Command Reference.	87
	sap_disable_replication	87
	sap_enable_replication	88
	sap_failover.	89
	sap_failover_drain_to_dr.	93
	sap_host_available.	96
	sap_set_host.	98
	sap_status path.	99
	sap_status resource.	105
	sap_status route.	108
	sap_update_replication {add remove}.	113
	sap_update_replication {move_out move_in}.	114
5.2	sp_hadr_admin.	115

1 Overview

This guide introduces the SAP Adaptive Server Enterprise (ASE) high availability and disaster recovery (HADR) with disaster recovery (DR) node system. The system and its functions are used for the Business Suite implementation. The SAP ASE HADR with DR node system (HADR with DR node for short) adds a standalone third node as the DR server to an existing HADR cluster.

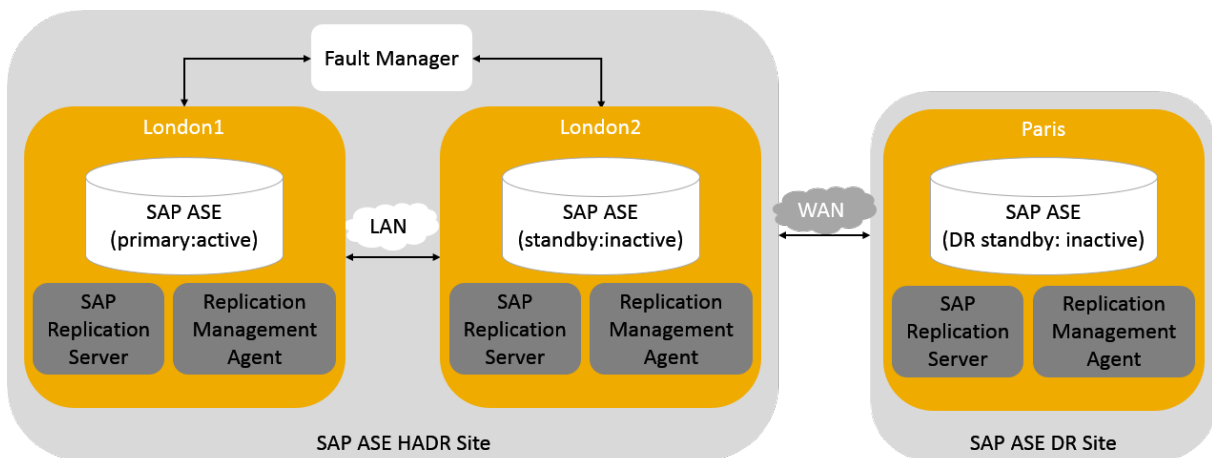
The HADR with DR node system consists of three SAP ASE servers:

- One designated as the primary, on which all transaction processing takes place.
- A standby node, which acts as a warm standby for the primary server, and contains copies of designated databases from the primary server.
- The DR node, which backs up designated databases from the primary on a geographically distant server.

i Note

The HADR with DR node system are supported on the Linux and Windows operating systems.

A HADR with DR node system looks similar to this, where `London1` is the primary node, `London2` is the standby node, and `Paris` is the DR node:



The HADR with DR node system includes an embedded SAP Replication Server, which synchronizes the databases among three servers. The system uses Replication Management Agent (RMA) to perform the initial setup and HADR operations.

The HADR with DR node system supports synchronous replication between the primary and standby servers for high availability so the two servers can synchronize with zero data loss. This requires a network link between the primary and standby servers fast enough for synchronous replication to keep up with the primary server's workload. Generally, this means that the network latency is at approximately the same speed as the local disk IO speed and fewer than 10 milliseconds. Anything longer than a few milliseconds may result in a slower response to write operations at the primary.

The HADR with DR node system supports asynchronous replication between the HADR cluster and the DR server for disaster recovery. The primary and DR servers (the standby and DR servers) using asynchronous replication can be geographically distant, so they can have a slower network link. With asynchronous

replication, Replication Agent thread for SAP ASE captures the workloads of the primary (or the standby server), which is delivered asynchronously to SAP Replication Server. The SAP Replication Server applies these workload change to the DR server.

The HADR with DR node system provides high availability together with protection in the event of a disaster: if the primary server is lost, the standby server can be used as a replacement, client applications can switch to the standby server, and the standby server can quickly become available for users with zero data loss. You can use the DR server as a replacement for disaster recovery if the HADR cluster is lost. There may be some data loss in this situation.

In an HADR with DR node system, the HADR mode for the DR host is `DR Standby`, and the HADR state is `Inactive`.

The HADR with DR node system installs these components:

- SAP ASE
- SAP Replication Server
- Replication Management Agent (RMA)
- (Optional) Fault Manager
- (Optional) SAP ASE Cockpit
- SAP Host Agent (if Fault Manager is used)

This guide only documents operations and command references specially used when adding DR node to the HADR system or managing the HADR system with DR node. For details about the HADR system itself, see the *HADR Users Guide*.

2 Adding a DR Node to an Existing Business Suite HA System

A DR node requires a functioning high-availability system. Install this HA system before you configure the DR node.

The DR node uses the same environment settings as the HA environments. See *Requirements and Limitations* for Business Suite installation in the *HADR Users Guide*.

2.1 Requirements for Adding a DR Node to an HA Environment

The DR and HA environments use many of the same environment settings.

The following are requirements for adding a DR node to an HA environment:

- The primary and companion servers are running in an HA configuration before for you set up the third node for DR.
- The HA and DR servers have the same:
 - Version of SAP ASE.
 - Configurations.
 - Operation system settings.
 - SAP ASE login name and password, default language, character set, sort order, and so on.
 - Supported hardware and operating system versions, and patch levels that you installed according to the installation guides for SAP ASE and Replication Server.
- The DR node requires:
 - The same default language, character set, and sort order, as well as logins and passwords as those of the primary and companion servers.
 - The same operating system version as the primary and companion servers.
- Your infrastructure must provide network speeds that support synchronized transaction propagation so the transaction log can be delivered from the primary to the standby and DR servers with minimum latency.
- Each server has its own dedicated storage to store SAP ASE software, transaction logs, and replication data.

The system described in this guide uses the following configuration:

Configuration Item	Primary server	Companion server	DR node
Host name	SFMACHINE1	SJMACHINE2	LAMACHINE3
SAP ASE server port number	4901	4901	4901

Configuration Item	Primary server	Companion server	DR node
Backup Server port number	4902	4902	4902
Replication Server port number	4905	4905	4905
Replication Server RSSD port number	4906	4906	4906
Installation directory	/sybase/<SID>	/sybase/<SID>	/sybase/<SID>
SAP ASE device files	/sybase/<SID>/data	/sybase/<SID>/data	/sybase/<SID>/data
RMA port number	4909	4909	4909
RMA RMI port	7000	7000	7000
<div style="background-color: #e0e0e0; padding: 5px;"> <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>			
DR_admin password	<password>	<password>	<password>
SAP ASE sa password	<password>	<password>	<password>
Site name for primary companion	SFSAP1	SJSAP2	LASAP3

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 DR Node Limitations

The DR node has these limitations:

- You cannot move a server from the primary or standby mode to a DR standby mode, or vice versa, unless you tear down and rebuild the HADR environment.
- The DR node does not participate in split-brain decisions.
- You cannot use the Fault Manager to monitor the DR node.
- DR node does not support automatic client failover or zero-downtime upgrades.

2.4 Adding the Disaster Recovery Node

Adding the Disaster Recovery node into the HADR system includes installing the Data Management software and running the `setuphadr` utility.

2.4.1 Installing the Business Suite Application

Install the Business Suite application using the software provisioning manager (SWPM).

Context

The installation process varies depending on which installation application (for example, Business Suite or Netweaver) 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 server
 - *Master Password* – is the same as the master password you entered for the primary server
 - *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, this option only applies to SAP ASE installation prior to 16.0
 - *Configure the replication system* – leave blank, this option only applies to SAP ASE installation prior to 16.0
 - *Materialize the secondary database* – leave blank, this option only applies to SAP ASE installation prior to 16.0
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 and entering the path in the box provided, or selecting Browse to explore the system.

2.4.2 Installing the Data Movement Component

Use a response file to install the Data Movement component.

Procedure

1. Log on to the host as user `syb<SID>`, where `<SID>` is your system ID.
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" as "E:\sybase".

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

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

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

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

i Note

Choose one of the following to agree to the SAP License Agreement when installing in silent mode:

- 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`.

5. (UNIX) set the environment variables by sourcing the `SYBASE.csh` or `SYBASE.sh` files:

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

6. Log on to the host as user `syb<SID>`, where `<SID>` is your system ID.
7. Connect to SAP ASE as user `sapsso`. For example:

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

8. Unlock the user `sa`:

```
sp_locklogin sa, unlock
```

2.4.3 Run `setuphadr` to Add the DR Node to the HADR System

Use the `setuphadr` utility to add the DR node to the HADR system.

Procedure

1. Copy the `setup_hadr.rs` file from the primary or the standby server to the DR machine.
2. Log on to the DR machine as `syb<SID>`.
3. Change these properties in the DR version of the `setup_hadr.rs` response file:
 - `setup_site=site1` – change to `site3` on DR:

```
setup_site=DR
```

- `is_secondary_site_setup=false` – change to "true":

```
is_secondary_site_setup=true
```

See [Sample setup_hadr.rs Response File for Business Suite \[page 13\]](#) 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 the following, which means you have successfully added the DR node to the HADR system:

```
./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 "site2" site: Add host "Huge_Machine1.corp" to allow
dump and load...
  Backup server on "site1" site: Add host "Huge_Machine2.corp" to allow
dump and load...
```

```

Setup Backup server allow hosts...Success
Setup RMA
    Set SAP ID to "AS1"...
    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:5005"...
    Set site name "SJHADR2" with SAP ASE host:port to "Huge_Machine2.corp:
4901" and Replication Server host:port to "Huge_Machine2.corp:5005"...
    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 "/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
'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 "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:4901' and data server 'AS1_SFHAD1'..
        Start the Replication Agent for database 'master' on host
'SFMACHINE1.BIG.corp:4901' and data server 'AS1_SFHAD1'..
        Stop the Replication Agent for database 'AS1' on host
'SFMACHINE1.BIG.corp:4901' and data server 'AS1_SFHAD1'..
        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 'SFHADR1' to target 'SJHADR2'...

```

2.4.4 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=sybase123
# 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=sybase123
# 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=sybase123
# 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=sybase123
# 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, dr
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, dr
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_rmi_port=7000
# RMA RMI occupies five consecutive ports, with the configured port occupying
# the highest number.
COMP.rma_tds_port=4909
#
# 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=20000
#####
# Site "DR" on host host3 with dr role
#####
# Host name where SAP ASE run
#
# Enter fully qualified domain name (FQDN)
# if your sites are on different subnet.
DR.ase_host_name=Huge_Machine3.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.
DR.rma_host_name=Huge_Machine3.corp
# Site name
#
# Enter value that identifies this site,
# like a geographical location.
# Value must be unique.
DR.site_name=Site3
# Site role
#
# Enter the role of this site.
# Valid values: primary, companion, dr
DR.site_role=dr
# directory where SAP ASE installed
DR.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).
DR.ase_user_data_directory=
DR.ase_server_name=NW7
DR.ase_server_port=4901
DR.backup_server_name=NW7_BS
DR.backup_server_port=4902
# Directory to store database dumps
# in materialization
#
# Backup server must able to access this directory
DR.backup_server_dump_directory=E:\sybase\NW7\data
# Port numbers for Replication Server and Replication Management Agent on host3
#
# In remote topology, these are the DR 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
DR.rma_rmi_port=7000
# RMA RMI occupies five consecutive ports, with the configured port occupying
# the highest number.
DR.rma_tds_port=4909
#
# Starting port number to use when setup Replication Server.
# Make sure next two ports (+1 and +2) are also available for use.
DR.srs_port=4905
# Device buffer for Replication Server on host3
# 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
DR.device_buffer_dir=E:\\sybase\\NW7\\data
DR.device_buffer_size=20000
# Persistent queue directory for Replication Server running on host3
#
# For synchronous replication (synchronization_mode=async),
# 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
DR.simple_persistent_queue_dir=E:\\sybase\\NW7\\data
DR.simple_persistent_queue_size=20000
```

2.5 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.

See the *Post-Installation Tasks for Primary and Companion Servers* chapter in *HADR Users Guide* for more details.

3 Managing the HADR with DR Node System

You can manage an HADR with DR node by performing activities such as monitoring replication to the DR node, performing failover, resynchronizing the DR node, disabling replication to the DR node, and recovering the HADR cluster from the DR node.

The examples in this chapter use the following values:

- Logical host name for primary node: PR
- Logical host name for standby node: HA
- Logical host name for DR node: DR
- Databases: master, db1, PI2

3.1 Determining the Mode and State of the DR Node

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 DR node's mode. The return values are 5 for `<@@hadr_mode>` and DR Standby for `hadr_mode`:

```
select hadr_mode()
-----
DR Standby
select @@hadr_mode
-----
5
```

- Use the `hadr_state` function and the `<@@hadr_state>` global variable to determine the member state. The return values are 2 for `<@@hadr_state>` and Inactive for `hadr_state`:

```
select hadr_state()
-----
Inactive
select @@hadr_state
-----
2
```

- You can include a return value (-1, 0, 1, 2, 3, or 5) 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(5)
-----
DR Standby
```

- Issuing `hadr_mode` and `hadr_state` functions without arguments returns the mode and state of the server, respectively:

```
select hadr_mode(), hadr_state()
```

```
-----
DR Standby                               Inactive
(1 row affected)
```

- Issue the `HADR mode` configuration parameter to determine the current mode of the server (this server is in non-HADR mode):

```
sp_configure 'HADR mode'
Parameter Name  Default      Memory      Used      Config Value Run
Value Unit Type
-----
HADR mode      -1           0           5         5         not
applicable 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)', 'DR Standby (5)',
'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 the *HADR Users Guide > sap_failover*.

3.2 Monitoring Replication to the Disaster Recovery Node

After you configure the disaster recovery (DR) node, use the Replication Management Agent (RMA) commands to manage primary, standby, and DR nodes in your HADR environment, and monitor performance:

Use RMA to monitor the following parameters:

- Replication performance of the system – the ability of Replication Server to synchronize within the SAP ASE HA Cluster and with the DR node. Performance is measured in terms of latency (in milliseconds) and throughput (in GB/hour).
- Failover performance – the system's ability to allocate extra resources and move operations to the standby systems in an event of disaster, such as when the primary SAP ASE Server is unavailable, the primary SAP ASE host is down, or the HADR system is isolated from the network.

The RMA commands provide information such as device usage, queue backlog, and SAP ASE transaction log size.

The key RMA commands used for monitoring are `sap_status path`, `sap_status resource`, `sap_status route`, `sap_status active_path`, `sap_status synchronization`, and `sap_send_trace`. These commands help you monitor the state, latency time, commit time, distribution path, and the drain status of the participating databases (`master`, `db1`, `PI2`) on the PR, HA, and DR nodes.

i Note

Use the RMA log to get information on specific tasks executed as a part of failover or fallback operations.

3.2.1 Monitoring Paths

The replication paths are at the database level from the primary logical host to the replicate logical host. As a DR administrator, you monitor the state of the replication path from the SAP ASE HA Cluster to the DR node.

Depending on a replication path's current state, you can categorize it in one of the following states:

- Active: The replication path is functional and the data replication is in progress.
- Suspended: Either the Replication Server or the replication agent thread (RAT) is nonfunctional.
- Down: Configured servers (Replication Server or SAP ASE) are unavailable.
- Defined: Replication path is ready for materialization.

The following table summarizes replication, participating databases, and their discovered properties:

Replication	Databases	Properties
<ul style="list-style-type: none"> From PR to HA From PR to DR From HA to DR From HA to PR 	master, PI2, and db1	<p>The following properties are discovered for each database:</p> <ul style="list-style-type: none"> State Latency time Latency Commit time Distribution path Drain status

To monitor the state of a replication path, log in to the RMA and execute the following command:

```
sap_status path
```

The following sample output displays the replication paths from the primary logical host to the standby logical host and the DR host.

In some cases, you may see latency displayed as "Unknown." To refresh its value, run the `sap_send_trace PR` command, and then re-execute the `sap_status path` command.

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.	Hostname	site2	Logical host
DR name.	HADR Status	DR Standby : Inactive	Identify the
DR primary and standby sites.	Synchronization Mode	Asynchronous	The configured
DR Synchronization Mode value.	Synchronization State	Inactive	Synchronization
DR Mode in which replication is currently operating.	Distribution Mode	Local	Configured value
DR for the distribution_mode replication model property.	Replication Server Status	Active	The status of
DR Replication Server.	Hostname	site1	Logical host
HA name.	HADR Status	Standby : Inactive	Identify the
HA primary and standby sites.	Synchronization Mode	Synchronous	The configured
HA Synchronization Mode value.	Synchronization State	Inactive	Synchronization
HA Mode in which replication is currently operating.	Distribution Mode	Remote	Configured value
HA for the distribution_mode replication model property.	Replication Server Status	Active	The status of
HA Replication Server.	Hostname	site0	Logical host
PR name.	HADR Status	Primary : Active	Identify the
PR 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	Unknown	No latency information for database 'db1'.
HA.PR.db1	Latency	Unknown	No latency information for database 'db1'.
HA.PR.db1	Commit Time	Unknown	No last commit time for the database 'db1'.
HA.PR.db1	Distribution Path	PR	The path of Replication Server through which transactions travel.
HA.PR.db1	Drain Status	Unknown	The drain status of the transaction logs of the primary database server.

HA.PR.master State (Replication Agent Thread). Transactions are not being replicated.	Suspended	Path is suspended
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 of the transaction logs of the primary database server.	Unknown	The drain status
PR.DR.PI2 State and replication can occur.	Active	Path is active
PR.DR.PI2 Latency Time calculated	2016-04-27 22:33:00.840	Time latency last
PR.DR.PI2 Latency (ms)	707	Latency
PR.DR.PI2 Commit Time replicated	2016-04-27 22:33:00.840	Time last commit
PR.DR.PI2 Distribution Path Replication Server through which transactions travel.	HA	The path of
PR.DR.PI2 Drain Status of the transaction logs of the primary database server.	Not Drained	The drain status
PR.DR.db1 State and replication can occur.	Active	Path is active
PR.DR.db1 Latency Time calculated	2016-04-27 22:33:00.832	Time latency last
PR.DR.db1 Latency (ms)	703	Latency
PR.DR.db1 Commit Time replicated	2016-04-27 22:33:00.832	Time last commit
PR.DR.db1 Distribution Path Replication Server through which transactions travel.	HA	The path of
PR.DR.db1 Drain Status of the transaction logs of the primary database server.	Not Drained	The drain status
PR.DR.master State and replication can occur.	Active	Path is active
PR.DR.master Latency Time calculated	2016-04-27 22:33:00.832	Time latency last
PR.DR.master Latency (ms)	703	Latency
PR.DR.master Commit Time replicated	2016-04-27 22:33:00.832	Time last commit
PR.DR.master Distribution Path Replication Server through which transactions travel.	HA	The path of
PR.DR.master Drain Status of the transaction logs of the primary database server.	Not Drained	The drain status
PR.HA.PI2 State and replication can occur.	Active	Path is active
PR.HA.PI2 Latency Time calculated	2016-04-27 22:33:00.640	Time latency last
PR.HA.PI2 Latency (ms)	607	Latency
PR.HA.PI2 Commit Time replicated	2016-04-27 22:33:00.646	Time last commit
PR.HA.PI2 Distribution Path Replication Server through which transactions travel.	HA	The path of
PR.HA.PI2 Drain Status of the transaction logs of the primary database server.	Not Drained	The drain status
PR.HA.db1 State and replication can occur.	Active	Path is active
PR.HA.db1 Latency Time calculated	2016-04-27 22:33:00.646	Time latency last
PR.HA.db1 Latency (ms)	610	Latency
PR.HA.db1 Commit Time replicated	2016-04-27 22:33:00.646	Time last commit

PR.HA.db1	Distribution Path	HA	The path of
PR.HA.db1	Drain Status	Not Drained	The drain status
PR.HA.master	State	Active	Path is active
PR.HA.master	Latency Time	2016-04-27 22:33:00.632	Time latency last
PR.HA.master	Latency (ms)	603	Latency
PR.HA.master	Commit Time	2016-04-27 22:33:00.632	Time last commit
PR.HA.master	Distribution Path	HA	The path of
PR.HA.master	Drain Status	Not Drained	The drain status
(92 rows affected)			

To monitor the active replication paths, execute the following command:

```
sap_status active_path
```

PATH	NAME	VALUE	
DR	Start Time	2016-04-27 22:33:58.842	Time command
DR	Elapsed Time	00:00:04	Command execution
DR	Hostname	site2	Logical host
DR	HADR Status	DR Standby : Inactive	Identify the
DR	Synchronization Mode	Asynchronous	The configured
DR	Synchronization State	Inactive	Synchronization
DR	Distribution Mode	Local	Configured value
DR	Replication Server Status	Active	The status of
HA	Hostname	site1	Logical host
HA	HADR Status	Standby : Inactive	Identify the
HA	Synchronization Mode	Synchronous	The configured
HA	Synchronization State	Inactive	Synchronization
HA	Distribution Mode	Remote	Configured value
HA	Replication Server Status	Active	The status of
PR	Hostname	site0	Logical host
PR	HADR Status	Primary : Active	Identify the
PR	Synchronization Mode	Synchronous	The configured
PR	Synchronization State	Synchronous	Synchronization
PR	Distribution Mode	Remote	Configured value
PR	Replication Server Status	Active	The status of

PR.DR.PI2	State	Active	Path is active
PR.DR.PI2	Latency Time	2016-04-27 22:33:00.834	Time latency last calculated
PR.DR.PI2	Latency (ms)	704	Latency
PR.DR.PI2	Commit Time	2016-04-27 22:33:00.834	Time last commit replicated
PR.DR.PI2	Distribution Path	HA	The path of Replication Server through which transactions travel.
PR.DR.PI2	Drain Status	Not Drained	The drain status of the transaction logs of the primary database server.
PR.DR.db1	State	Active	Path is active
PR.DR.db1	Latency Time	2016-04-27 22:33:00.846	Time latency last calculated
PR.DR.db1	Latency (ms)	710	Latency
PR.DR.db1	Commit Time	2016-04-27 22:33:00.846	Time last commit replicated
PR.DR.db1	Distribution Path	HA	The path of Replication Server through which transactions travel.
PR.DR.db1	Drain Status	Not Drained	The drain status of the transaction logs of the primary database server.
PR.DR.master	State	Active	Path is active
PR.DR.master	Latency Time	2016-04-27 22:33:00.840	Time latency last calculated
PR.DR.master	Latency (ms)	707	Latency
PR.DR.master	Commit Time	2016-04-27 22:33:00.840	Time last commit replicated
PR.DR.master	Distribution Path	HA	The path of Replication Server through which transactions travel.
PR.DR.master	Drain Status	Not Drained	The drain status of the transaction logs of the primary database server.
PR.HA.PI2	State	Active	Path is active
PR.HA.PI2	Latency Time	2016-04-27 22:33:00.646	Time latency last calculated
PR.HA.PI2	Latency (ms)	610	Latency
PR.HA.PI2	Commit Time	2016-04-27 22:33:00.652	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
PR.HA.db1	Latency Time	2016-04-27 22:33:00.652	Time latency last calculated
PR.HA.db1	Latency (ms)	613	Latency
PR.HA.db1	Commit Time	2016-04-27 22:33:00.652	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
PR.HA.master	Latency Time	2016-04-27 22:33:00.632	Time latency last calculated
PR.HA.master	Latency (ms)	603	Latency
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.		
(56 rows affected)		

☰ Output Code

3.2.2 Monitoring Replication Backlog and Resource Usage

To maintain consistency in data replication and provide business continuity, you should monitor the available storage space, resource utilization, and backlog for various components.

Use the `sap_status resource` command to view the information on device size, simple persistent queue size, transactional log space, inbound and outbound queue space, replication queue space, replication truncation backlog (for inbound queue and outbound queue), replication route queue truncation backlog, and other parameters.

Log in to the RMA and execute the following command:

```
sap_status resource
```

NAME VALUE	TYPE	

22:38:20.891	Start Time	2016-06-02
00:00:02	Elapsed Time	
-1	Estimated Failover Time	
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)

```

3.2.3 Monitoring Replication Path Subcomponents

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 simple persistent queue). You can use the `sap_status route` command to monitor the sequence of queues and threads.

Log in to the RMA and execute the following command:

```
sap_status route
```

PATH STATE	SEQUENCE	NAME BACKLOG	TYPE	QID	SPID	SITE
PR.DR.master 1	1	ASE	S	NULL	58312	site0
Active		0				
PR.DR.master 2	2	RAT	T	NULL	63	site0
Active		NULL				
PR.DR.master 3	3	RATCI	T	NULL	NULL	site1
Active (Active)		NULL				
PR.DR.master 4	4	SPQ	Q	106	NULL	site1
NULL		0				
PR.DR.master 5	5	CAP	T	NULL	53	site1
Active (Awaiting Command)		NULL				
PR.DR.master 6	6	SQM	T	NULL	22	site1
Active (Awaiting Message)		NULL				
PR.DR.master 7	7	IBQ	Q	106	NULL	site1
NULL		0				
PR.DR.master 8	8	SQT	T	NULL	73	site1
Active (Awaiting Wakeup)		NULL				
PR.DR.master 9	9	DIST	T	NULL	41	site1
Active (Awaiting Wakeup)		NULL				
PR.DR.master 10	10	SQM	T	NULL	95	site1
Active (Awaiting Message)		NULL				
PR.DR.master 11	11	RouteQ	Q	16777319	NULL	site1
NULL		0				
PR.DR.master 12	12	RSI	T	NULL	96	site1
Active (Awaiting Wakeup)		NULL				
PR.DR.master 13	13	SQM	T	NULL	96	site2
Active (Awaiting Message)		NULL				
PR.DR.master 14	14	OBQ	Q	116	NULL	site2
NULL		0				
PR.DR.master 15	15	DSI	T	NULL	171	site2
Active (Awaiting Message)		NULL				
PR.DR.master 16	16	ASE	S	NULL	31638	site2
Active		NULL				
PR.DR.PI2 1	1	ASE	S	NULL	58312	site0
Active		0				
PR.DR.PI2 2	2	RAT	T	NULL	64	site0
Active		NULL				
PR.DR.PI2 3	3	RATCI	T	NULL	NULL	site1
Active (Active)		NULL				
PR.DR.PI2 4	4	SPQ	Q	110	NULL	site1
NULL		0				
PR.DR.PI2 5	5	CAP	T	NULL	59	site1
Active (Awaiting Command)		NULL				
PR.DR.PI2 6	6	SQM	T	NULL	26	site1
Active (Awaiting Message)		NULL				
PR.DR.PI2 7	7	IBQ	Q	110	NULL	site1
NULL		0				
PR.DR.PI2 8	8	SQT	T	NULL	71	site1
Active (Awaiting Wakeup)		NULL				
PR.DR.PI2 9	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
PR.HA.PI2      1      ASE      S      NULL      58312   site0
Active
PR.HA.PI2      2      RAT      T      NULL      64      site0
Active
PR.HA.PI2      3      RATCI   T      NULL      NULL     site1
Active (Active)
PR.HA.PI2      4      SPQ      Q      110      NULL     site1
NULL
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
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
PR.HA.PI2      12     DSI      T      NULL      34      site1
Active (Awaiting Message) NULL
PR.HA.PI2      13     ASE      S      NULL      55597   site1
Active
PR.HA.db1      1      ASE      S      NULL      58312   site0
Active
PR.HA.db1      2      RAT      T      NULL      65      site0
Active
PR.HA.db1      3      RATCI   T      NULL      NULL     site1
Active (Active)
PR.HA.db1      4      SPQ      Q      114      NULL     site1
NULL
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
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
PR.HA.db1      12     DSI      T      NULL      36      site1
Active (Awaiting Message) NULL
PR.HA.db1      13     ASE      S      NULL      55597   site1
Active
(87 rows affected)

```

You can use the `sap_status synchronization` command to monitor the configured replication modes and the current replication states in the HADR environment.

Log in to the RMA and execute the following command:

```
sap_status synchronization
```

PATH INFO	NAME	VALUE	
PR name.	Hostname	site0	Logical host
PR standby sites.	HADR Status	Primary : Active	Identify the primary and
PR Mode value.	Synchronization Mode	Synchronous	The configured Synchronization
PR replication is currently operating.	Synchronization State	Synchronous	Synchronization Mode in which
PR.master replication is currently operating.	Synchronization State	Synchronous	Synchronization Mode in which
PR.PI2 replication is currently operating.	Synchronization State	Synchronous	Synchronization Mode in which
PR.db1 replication is currently operating.	Synchronization State	Synchronous	Synchronization Mode in which
(7 rows affected)			

3.3 Managing Failover in the HADR Cluster with the DR Node

The HADR cluster with the DR node supports failover from the primary node to the standby node.

The two types of failover are planned and unplanned. Planned failover allows you to use the `sap_failover` command to perform regular maintenance work and system upgrades on the primary node by switching application connections to the standby node.

Unplanned failovers are required when the primary server is down or unavailable. If the Fault Manager is configured, the unplanned failover happens automatically. Otherwise, use the `unplanned` option of the `sap_failover` command to perform the unplanned failover.

After a failover, the standby server becomes the new primary server, at which point you need to clean and disable the old replication paths and activate the new replication paths by using the `sap_failover_drain_to_dr` and `sap_host_available` commands.

Fail back to the former primary node when the maintenance work or system upgrades has finished, or the former primary node is back online.

i Note

Failover to the DR node is not supported.


```

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_dr <number_of_seconds_before_timeout>' to complete the
failover for all hosts. 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
15

Failover Total Number of Tasks
15

Failover Task Start           Fri Mar 18 03:26:39 EDT
2016

Failover Task End            Fri Mar 18 03:26:44 EDT
2016

Failover Hostname            site0

(12 rows affected)

```

When `sap_failover` has finished, the SAP ASE server on the HA logical host becomes the new primary server. It is only activated after all transaction backlogs on the PR logical host are drained to the HA logical host.

3. Run the following command to check whether all transaction backlogs are drained from the HADR cluster to the DR node. The example uses a timeout of 120 seconds:

```

sap_failover_drain_to_dr 120
go

```

If this command has successfully finished, all backlogs are drained to the DR node. Use the `sap_status` command to check the progress (see the bolded text), and that all backlogs are drained to the DR node:

```

TASKNAME          TYPE
VALUE
-----
-----
Status          Start Time          Fri Mar 18 03:28:19 EDT
2016
Status          Elapsed Time
00:00:31

```

```

FailoverRemaining Task Name          Failover drain to
DR.

FailoverRemaining Task State
Completed

FailoverRemaining Short Description   Failover drain to DR deactivate old
replication path and activate new replication path for all hosts.
FailoverRemaining Long Description   Started task 'Failover drain to DR.'
asynchronously.
FailoverRemaining Additional Info     Please execute command 'sap_status
task' to determine when task 'Failover drain to DR.' is complete.
FailoverRemaining Current Task Number
7

FailoverRemaining Total Number of Tasks
7

FailoverRemaining Task Start          Fri Mar 18 03:28:19 EDT
2016
FailoverRemaining Task End            Fri Mar 18 03:28:50 EDT
2016
FailoverRemaining Hostname            site0

(12 rows affected)

```

You can now perform planned maintenance work or a system upgrade on the PR logical host or the DR node. Establish replication from the HA logical host to the PR logical host and the DR node separately after the maintenance or upgrade is done, in your preferred order. This example resumes the replication from the HA logical host to the DR node first

4. Run the following command to establish replication to the DR node when the DR node is ready to rejoin the replication system:

```
sap_host_available DR
```

The system displays the following information when the command has finished successfully:

```

TASKNAME      TYPE
VALUE
-----
-----
HostAvailable Start Time          Fri Mar 18 03:32:52 EDT
2016
HostAvailable Elapsed Time
00:00:22

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 'DR' receiving replication from logical host 'HA'.
HostAvailable Current Task Number
9

```

```

HostAvailable Total Number of Tasks
9

HostAvailable Task Start          Fri Mar 18 03:32:52 EDT
2016
HostAvailable Task End            Fri Mar 18 03:33:14 EDT
2016
HostAvailable Hostname
site0

(11 rows affected)

```

Replication from the HA logical host to the DR node is established. You can continue to perform additional maintenance or upgrade tasks on the PR logical host before adding it to the replication system.

Note

If you perform `sap_host_available` for only one node, the partition size of SAP Replication Server on that node must be large enough to save all the data that is waiting to be replicated to another node, or else the SAP Replication Server queue becomes full and affects the data replication for all replication paths.

- Run the following command to establish the replication to the PR logical host:

```

sap_host_available PR
go

```

The system displays the following information when the command has finished successfully:

```

TASKNAME      TYPE
VALUE
-----
-----
HostAvailable Start Time          Fri Mar 18 04:29:32 EDT
2016
HostAvailable Elapsed Time
00:02:05

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
11

HostAvailable Total Number of Tasks
10

HostAvailable Task Start          Fri Mar 18 04:29:32 EDT
2016
HostAvailable Task End            Fri Mar 18 04:31:37 EDT
2016
HostAvailable Hostname
site0

(11 rows affected)

```

6. Check the replication path status by logging in to the RMA, and running the `sap_status path` command:

```
sap_status path

PATH          NAME          VALUE
INFO
-----
-----
-----
      Start Time          2016-03-18 03:36:49.595 Time command
started executing.
      Elapsed Time          00:00:00          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
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
Replication Server.
HA           Hostname          site1          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.
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.DR.PI2 State Active Path is
active and replication can occur.
HA.DR.PI2 Latency Time          2016-03-18 03:22:13.980 Time latency
last calculated
HA.DR.PI2 Latency          414          Latency
(ms)
HA.DR.PI2 Commit Time          2016-03-18 03:19:59.440 Time last
commit replicated
HA.DR.PI2 Distribution Path          PR          The path of
Replication Server through which transactions travel.
```

```

HA.DR.PI2      Drain Status          Not Drained          The drain
status of the transaction logs of the primary database server.
HA.DR.db1      State                Active              Path is
active and replication can occur.
HA.DR.db1      Latency Time                2016-03-18 03:22:14.560 Time latency
last calculated
HA.DR.db1      Latency                          704                  Latency
(ms)
HA.DR.db1      Commit Time                  2016-03-18 03:17:59.152 Time last
commit replicated
HA.DR.db1      Distribution Path              PR                    The path of
Replication Server through which transactions travel.
HA.DR.db1      Drain Status                    Not Drained          The drain
status of the transaction logs of the primary database server.
HA.DR.master State                Active              Path is
active and replication can occur.
HA.DR.master Latency Time                2016-03-18 03:22:13.980 Time latency
last calculated
HA.DR.master Latency                          414                  Latency
(ms)
HA.DR.master Commit Time                  2016-03-18 03:16:13.172 Time last
commit replicated
HA.DR.master Distribution Path              PR                    The path of
Replication Server through which transactions travel.
HA.DR.master Drain Status                    Not Drained          The drain
status of the transaction logs of the primary database server.
HA.PR.PI2      State                Active              Path is
active and replication can occur.
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                  2016-03-18 02:28:01.086 Time last
commit replicated
HA.PR.PI2      Distribution Path              PR                    The path of
Replication Server through which transactions travel.
HA.PR.PI2      Drain Status                    Not Drained          The drain
status of the transaction logs of the primary database server.
HA.PR.db1      State                Active              Path is
active and replication can occur.
HA.PR.db1      Latency Time                Unknown              No latency
information for database 'db1'.
HA.PR.db1      Latency                          Unknown              No latency
information for database 'db1'.
HA.PR.db1      Commit Time                  2016-03-18 02:25:54.744 Time last
commit replicated
HA.PR.db1      Distribution Path              PR                    The path of
Replication Server through which transactions travel.
HA.PR.db1      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                Unknown              No latency
information for database 'master'.
HA.PR.master Latency                          Unknown              No latency
information for database 'master'.
HA.PR.master Commit Time                  2016-03-18 02:23:49.526 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.
. . .
(128 rows affected)

```

The replication path is redirected from the HA logical host to the PR logical host and the DR node.

You can also run `sap_status active_path` to view a summary of all active replication paths.

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, and then re-run the `sap_status path` command.

For more information about monitoring replication path status, see [Monitoring Paths \[page 21\]](#).

7. (If the Fault Manager is configured) Start the Fault Manager:

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

Results

The planned failover from the primary node to the standby node has finished. Client applications are now connected to the SAP ASE server on the HA logical host, and new replication paths are established from the HA logical host to the PR logical host and to the DR node.

3.3.2 Performing Unplanned Failover to the Standby Node

When the primary SAP ASE server is down or lost, perform an unplanned failover so that client applications can continue to work on the SAP ASE server configured on the standby node.

Context

Use the `sap_failover` command with the `unplanned` option to perform an unplanned failover from the primary node (PR logical host in this example) to the standby node (HA logical host).

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. (Skip this step if the Fault Manager is configured) To guarantee zero data loss, run the following command to check the synchronization state of the primary node and make sure the state is synchronous before you proceed with the next step:

```
sap_status path
```

PATH INFO	NAME	VALUE	

. . .			
PR	Hostname	site0	Logical host name.
PR	HADR Status	Unknown	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	Unknown	The status of Replication Server.

2. (Skip this step if the Fault Manager is configured) Enter the following command to initiate the unplanned failover. The example uses a deactivation timeout of 120 seconds:

```
sap_failover PR, HA, 120, unplanned
```

Use the `sap_status` to check the progress, then proceed only after the `sap_failover` task has finished:

```
sap_status
```

TASKNAME VALUE	TYPE	

Status	Start Time	Fri Mar 18 04:16:45 EDT 2016
Status	Elapsed Time	00:00:05
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_dr <number_of_seconds_before_timeout>' to complete the
failover for all hosts. 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
15

Failover    Total Number of Tasks
15

Failover    Task Start            Fri Mar 18 04:16:45 EDT
2016

Failover    Task End              Fri Mar 18 04:16:50 EDT
2016

Failover    Hostname
site0

(12 rows affected)

```

When `sap_failover` finishes, the SAP ASE server on the HA logical host becomes the new primary server. It is activated only after all transaction backlogs on the PR logical host are drained to the HA logical host.

3. Run the following command to check whether all transaction backlogs are drained from the HADR cluster to the DR node. This example uses a timeout of 120 seconds:

```
sap_failover_drain_to_dr 120
```

If this command has successfully finished, all backlogs are drained to the DR node. Use the `sap_status` command to check the progress, and that all backlogs are drained to the DR node:

```
sap_status
```

```

TASKNAME          TYPE
VALUE
-----
-----
-----
Status           Start Time           Fri Mar 18 04:18:19 EDT
2016
Status           Elapsed Time
00:00:31

FailoverRemaining Task Name           Failover drain to
DR.

FailoverRemaining Task State
Completed

FailoverRemaining Short Description   Failover drain to DR deactivate old
replication path and activate new replication path for all hosts.
FailoverRemaining Long Description   Started task 'Failover drain to DR.'
asynchronously.
FailoverRemaining Additional Info     Please execute command 'sap_status
task' to determine when task 'Failover drain to DR.' is complete.
FailoverRemaining Current Task Number
7

FailoverRemaining Total Number of Tasks
7

FailoverRemaining Task Start           Fri Mar 18 04:18:19 EDT
2016
FailoverRemaining Task End             Fri Mar 18 04:18:50 EDT
2016
FailoverRemaining Hostname
site0

(12 rows affected)

```

You are now ready to establish replication from the HA logical host to the DR node.

4. Establish replication from the HA logical host (current primary node) to the DR node:

```
sap_host_available DR
```

The system displays the following information when the command has successfully finished:

```

TASKNAME          TYPE
VALUE
-----
-----
-----
HostAvailable Start Time           Fri Mar 18 05:32:52 EDT
2016
HostAvailable Elapsed Time
00:00:22

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 'DR' receiving replication from logical host 'HA'.
HostAvailable Current Task Number
9

HostAvailable Total Number of Tasks
9

HostAvailable Task Start              Fri Mar 18 05:32:52 EDT
2016
HostAvailable Task End                Fri Mar 18 05:33:14 EDT
2016
HostAvailable Hostname
site0

(11 rows affected)

```

- Restore the SAP ASE server on the PR logical host:

```

source /opt/sap/SYBASE.sh
cd /opt/sap/data/ASE-16_0/install
./RUN_SYBASE &

```

- Establish replication from the HA logical host to the primary node:

```
sap_host_available PR
```

The system displays the following when the command has successfully finished:

```

TASKNAME      TYPE
VALUE
-----
-----
HostAvailable Start Time              Fri Mar 18 05:29:32 EDT
2016
HostAvailable Elapsed Time
00:02:05

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
11

HostAvailable Total Number of Tasks
10

HostAvailable Task Start              Fri Mar 18 05:29:32 EDT
2016
HostAvailable Task End                Fri Mar 18 05:31:37 EDT
2016
HostAvailable Hostname
site0

(11 rows affected)

```

i Note

It does not matter in which order the replication paths to the PR logical host and the DR node are established. If you perform `sap_host_available` for only one node, the partition size of SAP Replication Server on that node must be large enough to save all the data waiting to be replicated to another node, or else the SAP Replication Server queue becomes full and affects the data replication for all replication paths.

- Run the following command to confirm that the replication from the HA logical host (current primary node) to the PR logical host and the DR node is active:

```
sap_status path
```

PATH INFO	NAME	VALUE	
-----	-----	-----	-----
started executing.	Start Time	2016-03-18 03:36:49.595	Time command
execution time.	Elapsed Time	00:00:00	Command
DR name.	DR Hostname	site2	Logical host
primary and standby sites.	HADR Status	DR Standby : Inactive	Identify the
configured Synchronization Mode value.	DR Synchronization Mode	Asynchronous	The
Synchronization Mode in which replication is currently operating.	DR Synchronization State	Inactive	
value for the distribution_mode replication model property.	DR Distribution Mode	Local	Configured
Replication Server.	DR Replication Server Status	Active	The status of
HA name.	HA Hostname	site1	Logical host
primary and standby sites.	HA HADR Status	Primary : Active	Identify the
configured Synchronization Mode value.	HA Synchronization Mode	Synchronous	The
Synchronization Mode in which replication is currently operating.	HA Synchronization State	Synchronous	
value for the distribution_mode replication model property.	HA Distribution Mode	Remote	Configured
Replication Server.	HA Replication Server Status	Active	The status of
PR name.	PR Hostname	site0	Logical host
primary and standby sites.	PR HADR Status	Standby : Inactive	Identify the
configured Synchronization Mode value.	PR Synchronization Mode	Synchronous	The
Synchronization Mode in which replication is currently operating.	PR Synchronization State	Inactive	
value for the distribution_mode replication model property.	PR Distribution Mode	Remote	Configured

```

PR          Replication Server Status Active          The status of
Replication Server.
. . .
HA.DR.PI2 State Active Path is
active and replication can occur.
HA.DR.PI2 Latency Time 2016-03-18 03:22:13.980 Time latency
last calculated
HA.DR.PI2 Latency 414 Latency
(ms)
HA.DR.PI2 Commit Time 2016-03-18 03:19:59.440 Time last
commit replicated
HA.DR.PI2 Distribution Path PR The path of
Replication Server through which transactions travel.
HA.DR.PI2 Drain Status Not Drained The drain
status of the transaction logs of the primary database server.
HA.DR.db1 State Active Path is
active and replication can occur.
HA.DR.db1 Latency Time 2016-03-18 03:22:14.560 Time latency
last calculated
HA.DR.db1 Latency 704 Latency
(ms)
HA.DR.db1 Commit Time 2016-03-18 03:17:59.152 Time last
commit replicated
HA.DR.db1 Distribution Path PR The path of
Replication Server through which transactions travel.
HA.DR.db1 Drain Status Not Drained The drain
status of the transaction logs of the primary database server.
HA.DR.master State Active Path is
active and replication can occur.
HA.DR.master Latency Time 2016-03-18 03:22:13.980 Time latency
last calculated
HA.DR.master Latency 414 Latency
(ms)
HA.DR.master Commit Time 2016-03-18 03:16:13.172 Time last
commit replicated
HA.DR.master Distribution Path PR The path of
Replication Server through which transactions travel.
HA.DR.master Drain Status Not Drained The drain
status of the transaction logs of the primary database server.
HA.PR.PI2 State Active Path is
active and replication can occur.
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 2016-03-18 02:28:01.086 Time last
commit replicated
HA.PR.PI2 Distribution Path PR The path of
Replication Server through which transactions travel.
HA.PR.PI2 Drain Status Not Drained The drain
status of the transaction logs of the primary database server.
HA.PR.db1 State Active Path is
active and replication can occur.
HA.PR.db1 Latency Time Unknown No latency
information for database 'db1'.
HA.PR.db1 Latency Unknown No latency
information for database 'db1'.
HA.PR.db1 Commit Time 2016-03-18 02:25:54.744 Time last
commit replicated
HA.PR.db1 Distribution Path PR The path of
Replication Server through which transactions travel.
HA.PR.db1 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 Unknown No latency
information for database 'master'.

```

```

HA.PR.master Latency                Unknown                No latency
information for database 'master'.
HA.PR.master Commit Time            2016-03-18 02:23:49.526 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.
. . .
(128 rows affected)

```

For more information about monitoring replication path status, see [Monitoring Paths \[page 21\]](#).

Note

In some cases, the HADR cluster can lose synchronization after an unplanned failover. If the synchronization state of the new primary node (HA logical host) is not *Synchronous* in the output of the `sap_status_path` command, rematerialize all databases on the new standby node (PR logical host). See [Resynchronizing the Standby Node \[page 55\]](#).

Results

The unplanned failover from the primary node to the standby node has finished. Client applications are now connected to the SAP ASE server on the HA logical host. New replication paths are established from the HA logical host to the PR logical host and the DR node.

3.3.3 Performing Failback to the Former Primary Node

Fail back to the former primary node when the planned maintenance or upgrade is done, or the former primary server is available.

Context

This procedure describes a general failback workflow, since the procedure of failback to the former primary node is similar to that of a failover. See [Performing Planned Failover to the Standby Node \[page 33\]](#) for more details and examples.

Procedure

1. (If the Fault Manager is configured) Stop the Fault Manager:

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

2. Start failover to the former primary node (PR logical host). This example uses a deactivation timeout of 120 seconds:

```
sap_failover HA, PR, 120
```

Use the `sap_status` to check the progress, and proceed only when the `sap_failover` has finished.

```
sap_status
```

3. Run the following command to check whether all transaction backlogs are drained from the HADR cluster to the DR node. This example uses a timeout of 120 seconds:

```
sap_failover_drain_to_dr 120
```

If this command has successfully finished, all backlogs are drained to the DR node. Use the `sap_status` to check the progress, and that all backlogs are drained to the DR node:

```
sap_status
```

4. Establish replication from the PR logical host to the DR node:

```
sap_host_available DR
```

5. Establish replication from the PR logical host to the HA logical host:

```
sap_host_available HA
```

i Note

It does not matter which order the replication paths to the PR logical host and the DR node are established. If you perform `sap_host_available` for only one node, the partition size of SAP Replication Server must be large enough to save all the data waiting to be replicated to another node, or else the SAP Replication Server queue becomes full and affects the data replication for all replication paths.

6. Verify that the status of replication paths is active from the primary node (PR logical host) to both the standby node (HA logical host) and the DR node:

```
sap_status path
```

7. (If the Fault Manager is configured) Start the Fault Manager:

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

Results

The failback from the HA logical host to the PR logical host has finished. Client applications are now reconnected to the SAP ASE server on the PR logical host. Replication is resumed from the PR logical host to the HA logical host and the DR node.

3.4 Resynchronizing the DR Node

When the DR node is out of sync with the HADR cluster, you can resynchronize a particular database on the DR node or the whole DR node.

This section describes how to resynchronize the DR node. See [Troubleshooting \[page 82\]](#) and the *HADR User Guide* to find out why the DR node is out of sync and how to fix it.

3.4.1 Resynchronizing a Database on the DR Node

When a database is out of sync with the source database in the HADR cluster, resynchronize it on the DR node.

Procedure

1. Log in to the RMA on the primary node.
2. Disable replication from the HADR cluster to the database (db1 in this example) on the DR node to prevent queue accumulation in the HADR cluster:

```
sap_disable_replication PR, DR, db1
```

3. Re-enable replication to the db1 database:

```
sap_enable_replication PR, DR, db1
```

4. Run the following command to check the state of the replication path to the db1 database:

```
sap_status path
```

```
PATH          NAME          VALUE
INFO
-----
. . .
PR.DR.db1    State          Defined    Path is defined and ready for
materialization.
. . .
```

The state of the PR.DR.db1 replication path is now `Defined`, and is ready for materialization.

5. Materialize the db1 database using either the automatic or manual method. This example uses the automatic method:

```
sap_materialize auto PR, DR, db1
```

Use the `sap_status` command to check whether the materialization has finished. This example shows completed materialization:

```
sap_status
```



```

TASKNAME      TYPE
VALUE
-----
-----
Status        Start Time      Thu Mar 19 06:25:28 EDT
2016
Status        Elapsed Time
00:01:43

Materialize Task Name
Materialize

Materialize Task State
Completed

Materialize Short Description Materialize
database
Materialize Long Description Completed automatic materialization of
database 'db1' from source 'PR' to target 'DR'.
Materialize Task Start      Thu Mar 19 06:25:28 EDT
2016
Materialize Task End        Thu Mar 19 06:27:11 EDT
2016
Materialize Hostname
site0

(9 rows affected)

```

- When the materialization has finished, use the `sap_status path` command to check the state of the PR.DR.db1 replication path.

```

sap_status path

PATH          NAME          VALUE
INFO
-----
-----
. . .
PR.DR.db1 State Active Path is active and
replication can occur.
PR.DR.db1 Latency Time 2016-03-19 06:30:14.566 Time latency last
calculated
PR.DR.db1 Latency 707 Latency
(ms)
PR.DR.db1 Commit Time 2016-03-19 06:30:14.572 Time last commit
replicated
PR.DR.db1 Distribution Path HA The path of
Replication Server through which transactions travel.
PR.DR.db1 Drain Status Not Drained The drain status of
the transaction logs of the primary database server.

. . .

```

The state of the PR.DR.db1 replication path is active now.

Results

The db1 database on the DR node is resynchronized.

3.4.2 Resynchronizing the DR Node

Resynchronizing the DR node involves disabling then re-enabling replication, then materializing the databases on the node.

Context

The DR node used in this example contains three databases: `master`, `db1`, and `PI2`. Although you can materialize databases for resynchronization manually, the steps described here use the automatic method.

Procedure

1. Log in to the RMA on the primary node.
2. Disable replication from the HA cluster to the DR node:

```
sap_disable_replication PR, DR
```

3. Re-enable replication to the DR node:

```
sap_enable_replication PR, DR
```

4. Run the `sap_status path` command to check the state of the replication path to all the databases on the DR node:

```
sap_status path
```

PATH INFO	NAME	VALUE	

. . .			
PR.DR.PI2	State	Defined	Path is
defined and ready for materialization.			
PR.DR.PI2	Latency Time	Unknown	No latency
information for database 'PI2'.			
PR.DR.PI2	Latency	Unknown	No latency
information for database 'PI2'.			
PR.DR.PI2	Commit Time	Unknown	No last
commit time for the database 'PI2'.			
PR.DR.PI2	Distribution Path	HA	The path of
Replication Server through which transactions travel.			
PR.DR.PI2	Drain Status	Not Drained	The drain
status of the transaction logs of the primary database server.			
PR.DR.db1	State	Defined	Path is
defined and ready for materialization.			
PR.DR.db1	Latency Time	Unknown	No latency
information for database 'db1'.			
PR.DR.db1	Latency	Unknown	No latency
information for database 'db1'.			
PR.DR.db1	Commit Time	Unknown	No last
commit time for the database 'db1'.			

```

PR.DR.db1      Distribution Path      HA      The path of
Replication Server through which transactions travel.
PR.DR.db1      Drain Status          Not Drained      The drain
status of the transaction logs of the primary database server.
PR.DR.master State      Defined      Path is
defined and ready for materialization.
PR.DR.master Latency Time      Unknown      No latency
information for database 'master'.
PR.DR.master Latency      Unknown      No latency
information for database 'master'.
PR.DR.master Commit Time      Unknown      No last
commit time for the database 'master'.
PR.DR.master Distribution Path      HA      The path of
Replication Server through which transactions travel.
PR.DR.master Drain Status      Not Drained      The drain
status of the transaction logs of the primary database server.
. . .

```

The state of all replication paths is now `Defined`, and both the replication paths are ready for materialization.

5. Materialize the `master` database using the automatic method:

```
sap_materialize auto PR, DR, master
```

6. Use `sap_status` to verify that the materialization has finished:

```

sap_status

TASKNAME      TYPE
VALUE
-----
-----
Status      Start Time      Thu Mar 20 06:21:17 EDT
2016
Status      Elapsed Time
00:00:30

Materialize Task Name
Materialize

Materialize Task State
Completed

Materialize Short Description Materialize
database
Materialize Long Description      Completed automatic materialization of
database 'master' from source 'PR' to target 'DR'.
Materialize Task Start      Thu Mar 20 06:21:17 EDT
2016
Materialize Task End      Thu Mar 20 06:21:47 EDT
2016
Materialize Hostname
site0

(9 rows affected)

```

7. Materialize the `db1` database:

```
sap_materialize auto PR, DR, db1
go
```

8. Materialize the PI2 database:

```
sap_materialize auto PR, DR, PI2
go
```

9. Use sap_status path to check the state of all replication paths when the materialization has finished

```
sap_status path
```

PATH	NAME	VALUE	

. . . .			
PR.DR.PI2	State	Active	Path is
active and replication can occur.			
PR.DR.PI2	Latency Time	2016-03-20 07:22:13.986	Time latency
last calculated			
PR.DR.PI2	Latency	417	Latency
(ms)			
PR.DR.PI2	Commit Time	2016-03-20 07:22:13.992	Time last
commit replicated			
PR.DR.PI2	Distribution Path	HA	The path of
Replication Server through which transactions travel.			
PR.DR.PI2	Drain Status	Not Drained	The drain
status of the transaction logs of the primary database server.			
PR.DR.db1	State	Active	Path is
active and replication can occur.			
PR.DR.db1	Latency Time	2016-03-20 07:22:14.566	Time latency
last calculated			
PR.DR.db1	Latency	707	Latency
(ms)			
PR.DR.db1	Commit Time	2016-03-20 07:22:14.572	Time last
commit replicated			
PR.DR.db1	Distribution Path	HA	The path of
Replication Server through which transactions travel.			
PR.DR.db1	Drain Status	Not Drained	The drain
status of the transaction logs of the primary database server.			
PR.DR.master	State	Active	Path is
active and replication can occur.			
PR.DR.master	Latency Time	2016-03-20 07:22:13.980	Time latency
last calculated			
PR.DR.master	Latency	414	Latency
(ms)			
PR.DR.master	Commit Time	2016-03-20 07:22:13.980	Time last
commit replicated			
PR.DR.master	Distribution Path	HA	The path of
Replication Server through which transactions travel.			
PR.DR.master	Drain Status	Not Drained	The drain
status of the transaction logs of the primary database server.			
. . . .			

The state of all replication paths is now active.

Results

The DR node is resynchronized with the HADR cluster.

3.5 Resynchronizing the Standby Node

When the standby node is out of sync with the primary node, you can resynchronize a particular database on the standby node or the whole standby node.

This section describes how to resynchronize the standby node. See [Troubleshooting \[page 82\]](#) and the *HADR User Guide* to find out why the standby node is out of sync and how to fix it.

3.5.1 Resynchronizing a Database on the Standby Node

When you resynchronize a database on the standby node, the corresponding database on the DR node loses synchronization with the HADR cluster, requiring you to resynchronize this database on the DR node.

Procedure

1. Log in to the RMA on the primary node.
2. Disable replication to the database from the primary node to the standby node. This example uses a database called `db1`:

```
sap_disable_replication PR, HA, db1
```

3. Re-enable replication to the `db1` database:

```
sap_enable_replication PR, HA, db1
```

4. Check the state of the replication path to the `db1` database:

```
sap_status path
```

```
PATH          NAME          VALUE
INFO
-----
. . .
PR.HA.db1    State          Defined      Path is defined and ready for
materialization.
. . .
```

The state of the `PR.HA.db1` replication path is now `Defined`, and is ready for materialization.

5. Materialize the `db1` database using either the automatic or manual method. This example uses the automatic method:

```
sap_materialize auto PR, HA, db1
```

Use the `sap_status` command to check whether the materialization has finished:

```
sap_status
```

```

TASKNAME      TYPE
VALUE
-----
-----
Status        Start Time          Thu Mar 20 06:25:28 EDT
2016
Status        Elapsed Time
00:01:43

Materialize Task Name
Materialize

Materialize Task State
Completed

Materialize Short Description Materialize
database
Materialize Long Description Completed automatic materialization of
database 'db1' from source 'PR' to target 'HA'.
Materialize Task Start          Thu Mar 20 06:25:28 EDT
2016
Materialize Task End            Thu Mar 20 06:27:11 EDT
2016
Materialize Hostname
site0

(9 rows affected)

```

6. Resynchronize the db1 database on the DR node by using the following commands:

```

sap_disable_replication PR, DR, db1
go
sap_enable_replication PR, DR, db1
go
sap_status path
go
sap_materialize auto, PR, DR, db1
go
sap_status
go

```

See [Resynchronizing a Database on the DR Node \[page 48\]](#) for detailed information.

7. When the resynchronization on the DR node has finished, use the `sap_status path` command to check the state of both replication paths to the db1 database:

```

sap_status path

```

PATH	NAME	VALUE
PR.DR.db1	State	Active
PR.DR.db1	Latency Time	2016-03-20 06:31:54.765
PR.DR.db1	Latency (ms)	707
PR.DR.db1	Commit Time replicated	2016-03-20 06:31:04.392

```

PR.DR.db1      Distribution Path  HA          The path of
Replication Server through which transactions travel.
PR.DR.db1      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      2016-03-20 06:30:14.566 Time latency last
calculated
PR.HA.db1      Latency          707        Latency
(ms)
PR.HA.db1      Commit Time      2016-03-20 06:30:14.572 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.
. . .

```

The state of both replication paths to the db1 database is now active.

Results

The db1 database on the standby node is resynchronized.

3.5.2 Resynchronizing the Standby Node

Whenever you resynchronize the standby node, resynchronize the DR node as well to prevent the DR node from losing synchronization with the HADR cluster.

Context

You can either resynchronize the standby and DR nodes separately or, as shown in this example, resynchronize them at the same time.

The standby and the DR nodes used in this example contain three participating databases: `master`, `db1`, and `PI2`. Although you can materialize databases for resynchronization manually, the steps described here use the automatic method.

Procedure

1. Log in to the RMA on the primary node.
2. Disable replication from the primary node to the standby node and the DR node:

```
sap_disable_replication PR
```

3. Re-enable replication to the standby node and the DR node:

```
sap_enable_replication PR
```

Note

This and the previous commands disable and enable the replication from the primary node to the standby node and the DR node at the same time. You can also re-enable the replication separately. For example:

```
sap_disable_replication PR HA
go
sap_enable_replication PR HA
go
sap_disable_replication PR DR
go
sap_enable_replication PR DR
go
```

4. Check the state of the replication paths to all databases on the standby node and the DR node:

```
sap_status path
```

PATH	NAME	VALUE	
INFO			

. . .			
PR.DR.PI2	State	Defined	Path is
defined and ready for materialization.			
PR.DR.PI2	Latency Time	Unknown	No latency
information for database 'PI2'.			
PR.DR.PI2	Latency	Unknown	No latency
information for database 'PI2'.			
PR.DR.PI2	Commit Time	Unknown	No last
commit time for the database 'PI2'.			
PR.DR.PI2	Distribution Path	HA	The path of
Replication Server through which transactions travel.			
PR.DR.PI2	Drain Status	Not Drained	The drain
status of the transaction logs of the primary database server.			
PR.DR.db1	State	Defined	Path is
defined and ready for materialization.			
PR.DR.db1	Latency Time	Unknown	No latency
information for database 'db1'.			
PR.DR.db1	Latency	Unknown	No latency
information for database 'db1'.			
PR.DR.db1	Commit Time	Unknown	No last
commit time for the database 'db1'.			
PR.DR.db1	Distribution Path	HA	The path of
Replication Server through which transactions travel.			
PR.DR.db1	Drain Status	Not Drained	The drain
status of the transaction logs of the primary database server.			
PR.DR.master	State	Defined	Path is
defined and ready for materialization.			
PR.DR.master	Latency Time	Unknown	No latency
information for database 'master'.			
PR.DR.master	Latency	Unknown	No latency
information for database 'master'.			
PR.DR.master	Commit Time	Unknown	No last
commit time for the database 'master'.			


```

PR.DR.master Distribution Path          HA          The path of
Replication Server through which transactions travel.
PR.DR.master Drain Status              Not Drained The drain
status of the transaction logs of the primary database server.
PR.HA.PI2 State Defined Path is
defined and ready for materialization.
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                  Unknown     No last
commit time for the database 'PI2'.
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 Defined Path is
defined and ready for materialization.
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                  Unknown     No last
commit time for the database 'db1'.
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 Defined Path is
defined and ready for materialization.
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               Unknown     No last
commit time for the database 'master'.
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.
. . .

```

The state of all replication paths is now `Defined`, and the replication paths are ready for materialization.

5. Materialize the `master` database on the standby node using either the automatic or manual method. This example uses the automatic method:

```
sap_materialize auto PR, HA, master
```

6. Use `sap_status` to verify that the materialization has finished:

```
sap_status

TASKNAME      TYPE
VALUE
-----
-----
Status        Start Time          Thu Mar 21 06:21:17 EDT
2016
Status        Elapsed Time
00:00:30

```

```

Materialize Task Name
Materialize

Materialize Task State
Completed

Materialize Short Description Materialize
database
Materialize Long Description Completed automatic materialization of
database 'master' from source 'PR' to target 'HA'.
Materialize Task Start Thu Mar 21 06:21:17 EDT
2016
Materialize Task End Thu Mar 21 06:25:47 EDT
2016
Materialize Hostname
site0

(9 rows affected)

```

7. Materialize the db1 database, then verify that materialization has finished:

```

sap_materialize auto PR, HA, db1
go
sap_status
go

```

8. Materialize the PI2 database, then verify that materialization has finished:

```

sap_materialize auto PR, HA, PI2
go
sap_status
go

```

9. Materialize the master database on the DR node then verify that the materialization has finished. You can use either the automatic or manual method. This example uses the automatic method:

```

sap_materialize auto PR, DR, master
go
sap_status
go

```

10. Materialize the db1 database, then verify that the materialization has finished:

```

sap_materialize auto PR, DR, db1
go
sap_status
go

```

11. Materialize the PI2 database, then verify that the materialization has finished:

```

sap_materialize auto PR, DR, PI2
go
sap_status
go

```

12. When the materialization has finished, use the sap_status path command to check the state of all replication paths:

```

sap_status path
go

```

PATH	NAME	VALUE	
INFO			

. . .			
PR.DR.PI2	State	Active	Path is
active and replication can occur.			
PR.DR.PI2	Latency Time	2016-03-21 07:22:17.986	Time latency
last calculated			
PR.DR.PI2	Latency	417	Latency
(ms)			
PR.DR.PI2	Commit Time	2016-03-21 07:22:17.992	Time last
commit replicated			
PR.DR.PI2	Distribution Path	HA	The path of
Replication Server through which transactions travel.			
PR.DR.PI2	Drain Status	Not Drained	The drain
status of the transaction logs of the primary database server.			
PR.DR.db1	State	Active	Path is
active and replication can occur.			
PR.DR.db1	Latency Time	2016-03-21 07:22:16.966	Time latency
last calculated			
PR.DR.db1	Latency	707	Latency
(ms)			
PR.DR.db1	Commit Time	2016-03-21 07:22:16.872	Time last
commit replicated			
PR.DR.db1	Distribution Path	HA	The path of
Replication Server through which transactions travel.			
PR.DR.db1	Drain Status	Not Drained	The drain
status of the transaction logs of the primary database server.			
PR.DR.master	State	Active	Path is
active and replication can occur.			
PR.DR.master	Latency Time	2016-03-21 07:22:15.380	Time latency
last calculated			
PR.DR.master	Latency	414	Latency
(ms)			
PR.DR.master	Commit Time	2016-03-21 07:22:15.280	Time last
commit replicated			
PR.DR.master	Distribution Path	HA	The path of
Replication Server through which transactions travel.			
PR.DR.master	Drain Status	Not Drained	The drain
status of the transaction logs of the primary database server.			
PR.HA.PI2	State	Active	Path is
active and replication can occur.			
PR.HA.PI2	Latency Time	2016-03-21 07:22:14.986	Time latency
last calculated			
PR.HA.PI2	Latency	417	Latency
(ms)			
PR.HA.PI2	Commit Time	2016-03-21 07:22:14.992	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	2016-03-21 07:22:14.566	Time latency
last calculated			
PR.HA.db1	Latency	707	Latency
(ms)			
PR.HA.db1	Commit Time	2016-03-21 07:22:14.572	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 last calculated	2016-03-21 07:22:13.881	Time latency
PR.HA.master Latency (ms)	414	Latency
PR.HA.master Commit Time commit replicated	2016-03-21 07:22:13.780	Time last
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.		
. . .		

The state of all replication paths is now active.

Results

The standby node is resynchronized with the primary node.

3.6 Recovering the HADR Cluster from the DR Node

The databases in the HADR cluster may lose synchronization after an unplanned failover, or the entire cluster may be damaged after a disaster. You can recover a specific database or the HADR cluster from databases on the DR node.

3.6.1 Recovering a Database in the HADR Cluster

When a database in the HADR cluster loses synchronization, you can recover the database from the corresponding database on the DR node.

Context

To recover a database in the HADR cluster from the DR node, dump the database from the DR node and load the dump into the primary node. Then rematerialize the corresponding databases on the standby node and the DR node. This example shows you how to recover the `db1` database in the HADR cluster.

Procedure

1. Log in to the RMA on the primary node.
2. Disable replication to the db1 database on the DR node:

```
sap_disable_replication PR, DR, db1
```

3. Disable replication to the db1 database on the standby node:

```
sap_disable_replication PR, HA, db1
```

4. Log in to the SAP ASE running on the DR node and dump the db1 database:

```
dump database db1 to "/SAP_backup/db1.dmp"
```

5. Log in to the SAP ASE running on the primary node and load the db1.dmp database.

```
load database db1 from "/SAP_backup/db1.dmp"
```

6. Bring the database online by running the following command on the primary node:

```
online database db1
```

7. In the primary RMA, enable replication to the db1 database on the standby node:

```
sap_enable_replication PR, HA, db1
```

8. Materialize the db1 database on the standby node using either the automatic or manual method. This example uses the automatic method:

```
sap_materialize auto PR, HA, db1
```

Use the `sap_status` command to check whether the materialization is completed has finished:

```
sap_status
```

```
TASKNAME      TYPE
VALUE
-----
-----
Status        Start Time      Thu Mar 22 06:25:28 EDT
2016
Status        Elapsed Time    00:01:43

Materialize Task Name
Materialize

Materialize Task State
Completed

Materialize Short Description Materialize
database
Materialize Long Description Completed automatic materialization of
database 'db1' from source 'PR' to target 'HA'.
Materialize Task Start      Thu Mar 22 06:25:28 EDT
2016
```

```

Materialize Task End          Thu Mar 22 06:27:11 EDT
2016
Materialize Hostname
site0

(9 rows affected)

```

9. Enable replication to the db1 database on the DR node:

```
sap_enable_replication PR, DR, db1
```

10. Materialize the db1 database on the DR node, then verify that the materialization has finished:

```
sap_materialize auto PR, DR, db1
go
sap_status
go
```

11. When the materialization has finished, use the `sap_status path` command to check the state of both replication paths to the db1 database:

```
sap_status path
```

```

PATH          NAME          VALUE
INFO
-----
. . .
PR.DR.db1 State Active Path is active and
replication can occur.
PR.DR.db1 Latency Time 2016-03-22 06:31:54.365 Time latency last
calculated
PR.DR.db1 Latency 707 Latency
(ms)
PR.DR.db1 Commit Time 2016-03-22 06:31:04.392 Time last commit
replicated
PR.DR.db1 Distribution Path HA The path of
Replication Server through which transactions travel.
PR.DR.db1 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 2016-03-22 06:30:14.566 Time latency last
calculated
PR.HA.db1 Latency 707 Latency
(ms)
PR.HA.db1 Commit Time 2016-03-22 06:30:14.572 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.
. . .

```

The state of both replication paths to the db1 database is now active.

Results

The `db1` database in the HADR cluster is recovered from the DR node.

3.6.2 Recovering the HADR Cluster

If the HADR cluster is out of sync due to some fatal errors or disasters, you can recover the HADR cluster with the databases on the DR node.

Context

To recover the HADR cluster, load each database from the DR node to the primary node, then rematerialize each database on the standby node and the DR node.

The primary node used in the following procedure contains three participating databases: `master`, `db1`, and `PI2`.

Procedure

1. Log in to the RMA on the primary node.
2. Disable replication from the primary node to the standby node and the DR node:

```
sap_disable_replication PR
```

3. Log into the SAP ASE running on the DR node and dump the `master`, `db1`, and `PI2` databases:

```
dump database master to "/SAP_backup/master.dmp"  
dump database db1 to "/SAP_backup/db1.dmp"  
dump database PI2 to "/SAP_backup/PI2.dmp"
```

4. Log into the SAP ASE running on the primary node and load the `master`, `db1`, and `PI2` databases:

```
load database master from "/SAP_backup/master.dmp"  
load database db1 from "/SAP_backup/db1.dmp"  
load database PI2 from "/SAP_backup/PI2.dmp"
```

5. Bring the database online by running the following command on the primary node:

```
online database master  
go  
online database db1  
go  
online database PI2  
go
```

- In the primary RMA, enable replication from the primary node to the standby node and the DR node.

```
sap_enable_replication PR
```

- Check the state of the replication paths to all databases on the standby node and the DR node:

```
sap_status path
```

PATH INFO	NAME	VALUE	

. . .			
PR.DR.PI2	State	Defined	Path is
defined and ready for materialization.			
PR.DR.PI2	Latency Time	Unknown	No latency
information for database 'PI2'.			
PR.DR.PI2	Latency	Unknown	No latency
information for database 'PI2'.			
PR.DR.PI2	Commit Time	Unknown	No last
commit time for the database 'PI2'.			
PR.DR.PI2	Distribution Path	HA	The path of
Replication Server through which transactions travel.			
PR.DR.PI2	Drain Status	Not Drained	The drain
status of the transaction logs of the primary database server.			
PR.DR.db1	State	Defined	Path is
defined and ready for materialization.			
PR.DR.db1	Latency Time	Unknown	No latency
information for database 'db1'.			
PR.DR.db1	Latency	Unknown	No latency
information for database 'db1'.			
PR.DR.db1	Commit Time	Unknown	No last
commit time for the database 'db1'.			
PR.DR.db1	Distribution Path	HA	The path of
Replication Server through which transactions travel.			
PR.DR.db1	Drain Status	Not Drained	The drain
status of the transaction logs of the primary database server.			
PR.DR.master	State	Defined	Path is
defined and ready for materialization.			
PR.DR.master	Latency Time	Unknown	No latency
information for database 'master'.			
PR.DR.master	Latency	Unknown	No latency
information for database 'master'.			
PR.DR.master	Commit Time	Unknown	No last
commit time for the database 'master'.			
PR.DR.master	Distribution Path	HA	The path of
Replication Server through which transactions travel.			
PR.DR.master	Drain Status	Not Drained	The drain
status of the transaction logs of the primary database server.			
PR.HA.PI2	State	Defined	Path is
defined and ready for materialization.			
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	Unknown	No last
commit time for the database 'PI2'.			
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	Defined	Path is
defined and ready for materialization.			


```

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              Unknown      No last
commit time for the database 'db1'.
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      Defined      Path is
defined and ready for materialization.
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              Unknown      No last
commit time for the database 'master'.
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.
. . .

```

The state of all replication paths is `Defined` now, and the replication paths are ready for materialization.

8. Materialize the `master` database first on the standby node using either the automatic or manual method. This example uses the automatic method:

```
sap_materialize auto PR, HA, master
```

Use the `sap_status` command to check whether the materialization has finished:

```

sap_status

TASKNAME      TYPE
VALUE
-----
-----
Status        Start Time           Thu Mar 23 06:25:28 EDT
2016
Status        Elapsed Time         00:01:43

Materialize Task Name
Materialize

Materialize Task State
Completed

Materialize Short Description Materialize
database
Materialize Long Description Completed automatic materialization of
database 'master' from source 'PR' to target 'HA'.
Materialize Task Start           Thu Mar 23 06:25:28 EDT
2016
Materialize Task End             Thu Mar 23 06:27:11 EDT
2016
Materialize Hostname
site0

(9 rows affected)

```

9. Materialize the db1 database, then verify that the materialization has finished:

```
sap_materialize auto PR, HA, db1
go
sap_status
go
```

10. Materialize the PI2 database, then verify that the materialization has finished:

```
sap_materialize auto PR, HA, PI2
go
sap_status
go
```

11. Materialize the master database on the DR node then verify that the materialization has finished. You can use either the automatic or manual method. This example uses the automatic method:

```
sap_materialize auto PR, DR, master
go
sap_status
go
```

12. Materialize the db1 database, then verify that the materialization has finished:

```
sap_materialize auto PR, DR, db1
go
sap_status
go
```

13. Materialize the PI2 database, then verify that the materialization has finished:

```
sap_materialize auto PR, DR, PI2
go
sap_status
go
```

14. When the materialization has finished, use the sap_status path command to check the state of all replication paths:

```
sap_status path
```

PATH	NAME	VALUE
PR.DR.PI2	State	Active
PR.DR.PI2	Latency Time last calculated	2016-03-23 07:22:17.986
PR.DR.PI2	Latency (ms)	417
PR.DR.PI2	Commit Time last committed	2016-03-23 07:22:17.992
PR.DR.PI2	Distribution Path	HA
PR.DR.PI2	Drain Status	Not Drained
PR.DR.db1	State	Active

```

PR.DR.db1      Latency Time                2016-03-23 07:22:16.966 Time latency
last calculated
PR.DR.db1      Latency                      707                      Latency
(ms)
PR.DR.db1      Commit Time                 2016-03-23 07:22:16.872 Time last
commit replicated
PR.DR.db1      Distribution Path             HA                       The path of
Replication Server through which transactions travel.
PR.DR.db1      Drain Status                   Not Drained              The drain
status of the transaction logs of the primary database server.
PR.DR.master State Active Path is
active and replication can occur.
PR.DR.master Latency Time                2016-03-23 07:22:15.380 Time latency
last calculated
PR.DR.master Latency                      414                      Latency
(ms)
PR.DR.master Commit Time                 2016-03-23 07:22:15.280 Time last
commit replicated
PR.DR.master Distribution Path             HA                       The path of
Replication Server through which transactions travel.
PR.DR.master Drain Status                   Not Drained              The drain
status of the transaction logs of the primary database server.
PR.HA.PI2 State Active Path is
active and replication can occur.
PR.DR.PI2      Latency Time                2016-03-23 07:22:14.986 Time latency
last calculated
PR.DR.PI2      Latency                      417                      Latency
(ms)
PR.DR.PI2      Commit Time                 2016-03-23 07:22:14.992 Time last
commit replicated
PR.DR.PI2      Distribution Path             HA                       The path of
Replication Server through which transactions travel.
PR.DR.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.DR.db1      Latency Time                2016-03-23 07:22:14.566 Time latency
last calculated
PR.DR.db1      Latency                      707                      Latency
(ms)
PR.DR.db1      Commit Time                 2016-03-23 07:22:14.572 Time last
commit replicated
PR.DR.db1      Distribution Path             HA                       The path of
Replication Server through which transactions travel.
PR.DR.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.DR.master Latency Time                2016-03-23 07:22:13.881 Time latency
last calculated
PR.DR.master Latency                      414                      Latency
(ms)
PR.DR.master Commit Time                 2016-03-23 07:22:13.780 Time last
commit replicated
PR.DR.master Distribution Path             HA                       The path of
Replication Server through which transactions travel.
PR.DR.master Drain Status                   Not Drained              The drain
status of the transaction logs of the primary database server.
. . .

```

The state of all replication paths is now active.

Results

The HADR cluster is recovered from the DR node.

3.7 Disabling Replication to the DR Node

In some situations, you might consider disabling replication to the DR node, which drops the subscription at the SAP ASE server on the DR node.

Context

Consider disabling replication to the DR node when:

- The DR node is lost.
- Network connectivity to the DR node is lost and cannot be restored within a reasonable amount of time.
- The amount of backlog remaining to be replicated to the DR node is very large, making it impossible to drain it to the DR node within a reasonable amount of time.
- The SAP ASE on the DR node is nonfunctional (for example, the database log is full or the database has crashed).

Procedure

Issue this command to disable replication:

```
sap_disable_replication <logical_host_name_for_primary>  
<logical_host_name_for_dr>
```

See [sap_disable_replication \[page 87\]](#) in [RMA Command Reference \[page 87\]](#) for more details.

To enable the replication to the DR node again, use `sap_enable_replication` and then run the materialization.

3.8 Removing the DR Node from the HADR System

Removing the DR node from the HADR system removes the replication system that is used for replicating data from the HADR system to the DR node. This action drops subscriptions, replication definitions, and connections, and also removes the SAP Replication Server instances from the DR node.

Context

i Note

Do not use `sap_teardown` or the `removehadr` utility to remove the DR node from the HADR system, because `sap_teardown` tears down the whole HADR system, and `removehadr` invokes `sap_teardown` when executing.

Procedure

1. Remove DR node:

```
sap_update_replication remove <DR_logical_host>
go
```

2. Clean up replication definition on DR host:

```
sap_drop_host <DR_logical_host>
go
```

3. Drop SAP 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 procedure rs_update_lastcommit
go
drop procedure rs_update_threads
```

```

go
drop procedure rs_get_thread_seq
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

```

4. Drop temporary tables in the tempdb database from SAP ASE on the DR host:

```

use tempdb
go
drop table <temp_table_name>
go

```

5. Remove DR_admin and DR_maint from SAP ASE on the DR host.
 - a. Execute following commands on all participated databases:

```

sp_dropalias <SID>_maint
go

```

- b. Execute following commands on the master database:

```

use master
go
drop login <SID>_maint
go
drop role sap_maint_user_role with override
go
use master
go
sp_droplogin DR_admin
go

```

3.9 Performing a Rolling Upgrade on a DR Node System

You can update applications in a DR node system with zero downtime.

Prerequisites

Generally, the steps for performing a rolling upgrade are:

- Upgrade the DR node Replication Server and SAP ASE.
- Upgrade the primary node Replication Server.
- Perform a failover from the primary node to the standby node.
- Upgrade the standby node Replication Server

- Upgrade the primary node SAP ASE.
- Perform a fail back from the standby node to the primary node.
- Upgrade the standby node SAP ASE.

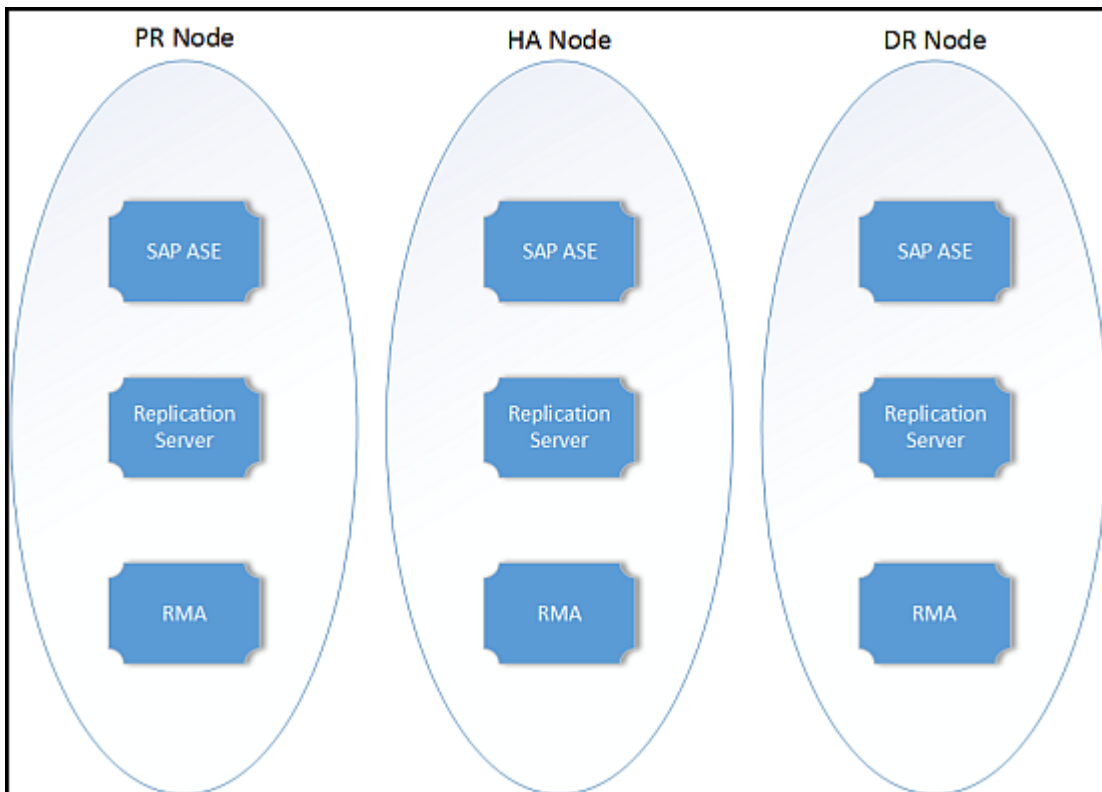
i Note

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.

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

- Primary logical hostname – PR
- Standby logical hostname – HA
- DR Standby logical hostname – DR
- Old version – V1
- New version – V2



Procedure

1. Upgrade the DR node's Replication Server:

- a. Log into the DR node's RMA as the user `DR_admin` using `isql`, and issue this to start the upgrade process:

```
sap_upgrade_server SRS, start, DR
```

- b. Shut down the DR node's RMA. From the RMA `isql` prompt, issue:

```
shutdown
```

- c. Install the new version of Replication Server using the `setup.bin` utility with the `-DDR=TRUE` parameter:

```
<installation_directory>/setup.bin -DDR=TRUE
```

- d. Start RMA on the DR node by issuing this from the command line:

```
$SYBASE/DM/SCC-3_2/bin/scc.sh
```

- e. Log into RMA on the DR node as the `DR_admin` user and issue this to complete the upgrade for Replication Server and RMA:

```
sap_upgrade_server SRS, finish, DR
```

2. Upgrade the DR node's SAP ASE:

- a. Temporarily unlock the `sa` user on the DR node for upgrade. Log in with `isql` as a user with `sapss0` permission on the DR node's SAP ASE and issue:

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

- b. Log into the DR node's RMA as the `DR_admin` user and issue `sap_upgrade_server` to start the upgrade for SAP ASE:

```
sap_upgrade_server ASE, start, DR
```

- c. Shut down the DR node's SAP ASE and Backup Server. Log into the DR node's SAP ASE with `isql` and issue:

```
shutdown SYB_BACKUP
go
shutdown
go
```

- d. Upgrade the DR node's SAP ASE. Start the SAP installer from the installation directory:

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

- e. Log into DR node's RMA as the `DR_admin` user and issue `sap_upgrade_server` to finish the upgrade for SAP ASE (the process resumes the DSI connections from the DR node's Replication Server to the DR node's SAP ASE):

```
sap_upgrade_server ASE, finish, DR
```


- f. Execute the installation script for DR:

```
installmaster instmsgs.ebf installjsdb installdbextend installcommit  
installmodel
```

- g. Log into the DR node's SAP ASE as the user with `sapsso` permissions and issue this once the upgrade is complete:

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

3. Upgrade the PR node's Replication Server:

- a. Log into the PR node's RMA as `DR_admin`, and issue (this shuts down the PR node's Replication Server and RMA):

```
sap_upgrade_server SRS, start, PR
```

- b. Shut down the PR node's RMA. From the `isql` prompt, issue:

```
shutdown
```

- c. Overwrite the current version of Replication Server using the `setup.bin` utility with the `-DDR=TRUE` parameter:

```
<installation_directory>/setup.bin -DDR=TRUE
```

- d. Start the PR node's RMA by issuing:

```
$SYBASE/DM/SCC-3_2/bin/scc.sh
```

- e. Log into the PR node's RMA as the `DR_admin` user and issue `sap_upgrade_server` to finish the upgrade for the Replication Server:

```
sap_upgrade_server SRS, finish, PR
```

4. Fail over from the PR node to the HA node. Log into the PR node's RMA as `DR_admin` and issue:

```
sap_failover PR, HA, 30, force  
sap_failover_drain_to_dr 30  
sap_host_available PR  
sap_host_available DR
```

5. Upgrade the HA node's Replication Server:

- a. Log into the HA node's RMA as `DR_admin`, and issue (this shuts down the HA node's Replication Server and RMA):

```
sap_upgrade_server SRS, start, HA
```

- b. Shut down the HA node's RMA. From the `isql` prompt, issue:

```
shutdown
```

- c. Overwrite the current version of Replication Server using the `setup.bin` utility with the `-DDR=TRUE` parameter:

```
<installation_directory>/setup.bin -DDR=TRUE
```

- d. Start the HA node's RMA by issuing:

```
$SYBASE/DM/SCC-3_2/bin/scc.sh
```

- e. Log into the HA node's RMA as the `DR_admin` user and issue `sap_upgrade_server` to finish the Replication Server upgrade:

```
sap_upgrade_server SRS, finish, HA
```

6. Upgrade the PR node's SAP ASE:

- a. Temporarily unlock the `sa` user on the PR node for upgrade. Log in as the user with `sapssso` permission on the PR node's SAP ASE and issue:

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

- b. Log into PR node's RMA as the `DR_admin` user and issue `sap_upgrade_server` to start the SAP ASE upgrade:

```
sap_upgrade_server ASE, start, PR
```

- c. Shut down the PR node's SAP ASE and Backup Server. Log into the PR node's SAP ASE and issue:

```
shutdown SYB_BACKUP
go
shutdown
go
```

- d. Upgrade the PR node's SAP ASE. Start the SAP installer from the installation directory:

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

- e. Log into the PR node's RMA as the `DR_admin` user and issue `sap_upgrade_server` to finish the upgrade for SAP ASE (the process resumes the DSI connections from the PR node's Replication Server to the PR node's SAP ASE):

```
sap_upgrade_server ASE, finish, PR
```

- f. Execute the installation script for PR:

```
installmaster instmsgs.ebf installjsdb installdbextend installcommit
installmodel
```

- g. Log in to PR node's SAP ASE as the user with `sapssso` permissions and issue this once the upgrade is complete:

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

7. Fail back from the HA node to the PR node. Log into the HA node's RMA as `DR_admin` and issue:

```
sap_failover HA, PR, 30, force
sap_failover drain_to_dr 30
sap_host_available HA, suspend
sap_host_available DR
```

8. Upgrade the HA node's SAP ASE:

- a. Temporarily unlock the `sa` user on the HA node. Log in as the user with `sapssso` permission on the HA node's SAP ASE and issue:

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

- b. Log into the HA node's RMA as the `DR_admin` user and issue `sap_upgrade_server` to start the upgrade for SAP ASE:

```
sap_upgrade_server ASE, start, HA, suspend
```

- c. Shut down the HA node's SAP ASE and Backup Server. Log into the HA node's SAP ASE server and issue:

```
shutdown SYB_BACKUP
go
shutdown
go
```

- d. Upgrade the HA node's SAP ASE. Start the SAP installer from the installation directory.

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

- e. Log into the HA node's RMA as the `DR_admin` user and issue `sap_upgrade_server` to finish the upgrade for SAP ASE (the process resumes the DSI connections from the PR node's Replication Server to the PR node's SAP ASE):

```
sap_upgrade_server ASE, finish, HA
```

- f. Execute the installation script for DR:

```
installmaster instmsgs.ebf installjsdb installdbextend installcommit
installmodel
```

- g. Log in to HA node's SAP ASE as the user with `sapss0` permissions and issue this once the upgrade is complete:

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

3.10 Recovering Replication to DR Node When the Standby SAP Replication Server is Down

When the standby SAP Replication Server is down in an HADR with DR node system, a temporary solution is to configure the primary SAP ASE server to replicate data to the DR node directly, then, after the standby SAP Replication Server is restored, configure the primary SAP ASE server to connect to it once again.

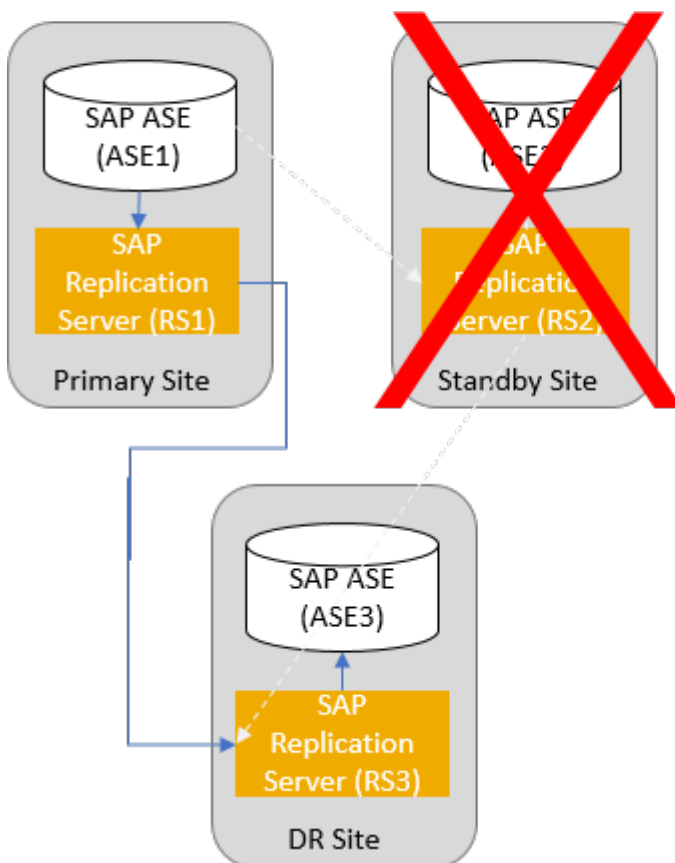
In an HADR with DR node system, the SAP Replication Server instance on the standby host receives data from the primary SAP ASE server, then replicates data to the SAP Replication Server instance on the DR node. Replication to the DR node stops when the standby SAP Replication Server goes down. To resume data replication, configure the primary SAP ASE server to bypass the standby SAP Replication Server, and replicate data to the DR SAP Replication Server directly. After you restore the standby host, reconfigure the primary SAP ASE server to connect to the standby SAP Replication Server.

The examples in this section use the following values:

Value	Example
Primary SAP ASE	ASE1

Value	Example
Primary SAP Replication Server	RS1
Standby SAP ASE	ASE2
Standby SAP Replication Server	RS2
SAP ASE on the DR node	ASE3
SAP Replication Server on the DR node	RS3

The following diagram shows the data flow after you configure the primary SAP ASE server to connect to the SAP Replication Server instance on the DR node:



In normal replication:

1. ASE1 sends data to RS2 remotely.
2. RS2 applies data to ASE2 while propagating data to RS3.

3. RS3 applies data to ASE3.

If RS2 is down, you can continue replicating data to the DR node by switching the replication path so that ASE1 connects to RS1 locally, RS1 propagates data to RS3, and RS3 applies data to ASE3.

When RS2 is restored, switch the replication path back, so that ASE1 connects to RS2 remotely to provide high availability functionality again.

Recover replication to the DR node based on your restoration status of RS2:

- You can restore RS2 before RS1 is filled up. See [Recovering Replication Before Primary Server is Full \[page 77\]](#).
- You cannot restore RS2 before RS1 is filled up. See [Recovering Replication After Primary Server is Full \[page 78\]](#).
- You cannot restore RS2. See [Recovering Replication When You Cannot Restore the Standby Server \[page 80\]](#).

3.10.1 Recovering Replication Before Primary Server is Full

Use these steps to recover replication when you can restore the standby SAP Replication Server before the primary SAP Replication Server has filled up.

Context

See [Recovering Replication to DR Node When the Standby SAP Replication Server is Down \[page 75\]](#) for the list of server names used in the examples.

After the standby SAP Replication Server is down, if you switch the replication path to recover the replication to the DR node, the data replicated to DR is also cached in the route queue from the primary to the standby SAP Replication Server.

If you can restore the standby SAP Replication Server before the route queue is filled up, data can be temporarily cached in the route queue, and you need not disable the replication path from the primary SAP ASE to the standby SAP Replication Server. Data to be replicated to the standby SAP Replication Server is not lost and you have no need to rematerialize the standby SAP ASE after the standby SAP Replication Server is restored.

Procedure

1. Switch the replication topology from remote (ASE1 → RS2 → RS3 → ASE3) to local (ASE1 → RS1 → RS3 → ASE3):

```
sap_update_replication move_out, <standby_logical_host_name>
```

After executing the command, the replication from ASE1 to ASE3 works immediately. Data is replicated from ASE1 to RS1 locally and then to RS3 on the DR node. The same data to be replicated to ASE2 is cached in the route queue from RS1 to RS2.

2. Re-create replication from the primary host to the standby host after the standby host is restored, while the replication topology remains as local:

```
sap_update_replication move_in, <standby_logical_host_name>, start
```

The replication topology is still local; that is, replicate data from ASE1 to RS1 locally, RS1 propagates data to both RS2 and RS3.

3. Switch the replication topology from local to remote to enable high availability functionality:

```
sap_update_replication move_in, <standby_logical_host_name>, finish [,  
<data_drain_time>]
```

<data_drain_time> is the maximum time allowed to drain the pending data to ASE2 and ASE3 before switching the replication topology from local to remote. If the draining time exceeds the value of <data_drain_time>, the command returns an error; you can execute the command again to drain the remaining data.

After you execute the command, the replication topology switches to remote; that is, replicate data from ASE1 to RS2 remotely, RS2 propagates data to RS3. Using remote topology enables high availability functionality of the HADR cluster.

i Note

Before switching the topology from local to remote, the system stops the RepAgent on ASE1 and waits until all replicating data is drained to ASE2 and ASE3. If RepAgent on ASE1 stops for too long, the log file on ASE1 fills up and ASE1 cannot work normally. To avoid this situation, use `sap_status path` to check the replication latency before switching the topology to remote. If the replication latency from ASE1 to ASE2 is similar to that of ASE1 to ASE3, the accumulated data to be replicated to ASE2 while RS2 is down is almost drained, so that you can then switch the topology to remote without waiting for the data to be drained for a long time.

3.10.2 Recovering Replication After Primary Server is Full

Use these steps to recover replication when you cannot restore the standby SAP Replication Server before the primary SAP Replication Server has filled up.

Context

See [Recovering Replication to DR Node When the Standby SAP Replication Server is Down \[page 75\]](#) for the list of server names used in the examples.

After the standby SAP Replication Server is down, if you switch the replication path to recover the replication to the DR node, the data replicated to DR is also cached in the route queue from the primary to the standby SAP Replication Server.

If you cannot restore the standby SAP Replication Server before the route queue is filled up, the data uses up disk space and blocks the replication from the primary SAP ASE. When this happens, you need to disable the replication path from the primary SAP ASE to the standby SAP Replication Server, then rematerialize the standby SAP ASE after the standby SAP Replication Server is restored.

Procedure

1. Switch the replication topology from remote (ASE1 → RS2 → RS3 → ASE3) to local (ASE1 → RS1 → RS3 → ASE3):

```
sap_update_replication move_out, <standby_logical_host_name>
```

After executing the command, the replication from ASE1 to ASE3 works immediately. Data is replicated from ASE1 to RS1 locally and then to RS3 on the DR node. The same data to be replicated to ASE2 is cached in the route queue from RS1 to RS2.

2. Disable replication to the standby host to avoid more data from accumulating in the route queue:

```
sap_disable_replication <primary_logical_host_name>,  
<standby_logical_host_name>
```

3. Re-create replication from the primary host to the standby host after the standby host is restored, while the replication topology remains as local:

```
sap_update_replication move_in, <standby_logical_host_name>, start
```

The replication topology is still local; that is, replicate data from ASE1 to RS1 locally, RS1 propagates data to both RS2 and RS3.

4. Enable replication to the standby host:

```
sap_enable_replication <primary_logical_host_name>,  
<standby_logical_host_name>
```

5. Rematerialize the standby databases. See [Resynchronizing the Standby Node \[page 53\]](#).
6. Switch the replication topology from local to remote to enable high availability functionality:

```
sap_update_replication move_in, <standby_logical_host_name>, finish [,  
<data_drain_time>]
```

`<data_drain_time>` is the maximum time allowed to drain the pending data to ASE2 and ASE3 before switching the replication topology from local to remote. If the draining time exceeds the value of `<data_drain_time>`, the command returns an error; you can execute the command again to drain the remaining data.

After you execute the command, the replication topology switches to remote; that is, replicate data from ASE1 to RS2 remotely, RS2 propagates data to RS3. Using remote topology enables high availability functionality of the HADR cluster.

Note

Before switching the topology from local to remote, the system stops the RepAgent on ASE1 and waits until all replicating data is drained to ASE2 and ASE3. If RepAgent on ASE1 stops for too long, the log file on ASE1 fills up and ASE1 cannot work normally. To avoid this situation, use `sap_status path to`

check the replication latency before switching the topology to remote. If the replication latency from ASE1 to ASE2 is similar to that of ASE1 to ASE3, the accumulated data to be replicated to ASE2 while RS2 is down is almost drained, so that you can then switch the topology to remote without waiting for the data to be drained for a long time.

3.10.3 Recovering Replication When You Cannot Restore the Standby Server

If you cannot restore the SAP Replication Server instance on the standby host, remove it permanently and recover the replication path from the primary SAP ASE to the SAP ASE on the DR node only.

Context

See [Recovering Replication to DR Node When the Standby SAP Replication Server is Down \[page 75\]](#) for the list of server names used in the examples.

i Note

High availability functionality does not work in this solution. If you need high availability, tear down the current environment and set up a new HADR with DR node system.

Procedure

1. Switch the replication topology from remote (ASE1 → RS2 → RS3 → ASE3) to local (ASE1 → RS1 → RS3 → ASE3):

```
sap_update_replication move_out, <standby_logical_host_name>
```

After executing the command, the replication from ASE1 to ASE3 works immediately. Data is replicated from ASE1 to RS1 locally and then to RS3 on the DR node. The same data to be replicated to ASE2 is cached in the route queue from RS1 to RS2.

2. Disable replication to the standby host to avoid more data from accumulating in the route queue:

```
sap_disable_replication <primary_logical_host_name>,  
<standby_logical_host_name>
```


3.10.3.1 Resetting Environment When You Cannot Restore Standby Server

To enable high availability when you cannot restore the SAP Replication Server instance on the standby host, tear down the current environment and set up a new HADR with DR node system.

Context

When you cannot restore an SAP Replication Server instance on the standby host, removing it and recovering the replication path from the primary SAP ASE to the SAP ASE on the DR node only enables disaster recovery. To enable high availability, you need to tear down the current environment and set up a new HADR with DR node system.

Procedure

1. Tear down the environment by executing `sap_teardown`. See [Performing a Teardown Using `sap_teardown` \[page 83\]](#).
2. Clean up the environment using the `removehadr` utility. See *Removing an HADR Environment Using the `removehadr` Utility* in *HADR Users Guide*.
3. Set up the HADR with DR node environment. See [Adding a DR Node to an Existing Business Suite HA System \[page 6\]](#).

3.11 SQL Statement Replication

HADR with DR node system supports the same SQL statement replication functionality that is supported in the HADR system. See *SQL Statement Replication* in the *HADR Users Guide* for more details.

4 Troubleshooting

Find out solutions or workarounds for issues you may encounter with the HADR cluster with the DR node.

Replication is suspended

- To find out why replication is not working for a particular path, run the `sap_status route` command to view all sub-threads status and replication queue backlogs.

DSI on DR node shuts down

The Data Server Interface (DSI) of SAP Replication Server on the Disaster Recovery (DR) node shuts down if you execute `sp_config_rep_agent` to change the RepAgent configuration on the primary SAP ASE while the RepAgent is running.

Starting from version 16.0 SP03 PL04, the RepAgent on the DR SAP ASE is disabled. If the RepAgent on the primary SAP ASE is not stopped, configuration commands to the RepAgent on the primary SAP ASE are replicated to the RepAgent on the DR SAP ASE, which causes the shutdown of the SAP Replication Server DSI on the DR node with the following error:

```
Message from server: Message: 18374, State 1, Severity 16 --
'Database 'DB1' is not configured to use Replication Agent.
Run sp_config_rep_agent without parameters to see a list of databases that use
Replication Agent.
Use the ENABLE option of sp_config_rep_agent to configure.
```

Workaround:

1. Run `set replication off` before you run the `sp_config_rep_agent` command to change the RepAgent configuration.
2. Run `set replication on` after you change the RepAgent configuration.

See the *HADR Users Guide* for more information about troubleshooting the HADR cluster.

4.1 Performing a Teardown

Performing a teardown requires tasks on the primary, standby, and DR 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 all the three SAP ASE servers.
- Drops the HADR group from all the three servers.
- Disables HADR on all the three 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`.

4.1.1 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 with DR node 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 and DR 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 and DR database (the data remains current as of when replication was last active).
- Deletes these directories:
 - Instance directories and their contents on the primary, standby, and DR Replication Servers:
 - On the primary server – `SYBASE/DM/<CID_REP_primary_logical_host_name>`
 - On the standby server – `SYBASE/DM/<CID_REP_standby_logical_host_name>`
 - On the DR server – `SYBASE/DM/<CID_REP_DR_logical_host_name>`
 - Any device files created by Replication Server.

The primary, standby, and DR 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. The user is responsible for maintaining these directories.

The primary, standby, and DR 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
```

4.1.2 Manually Removing the Replication in HADR System with DR Node Environment

Use the `sap_teardown` command to perform a teardown.

Perform the following to manually remove HADR replication:

1. Log into the primary, standby, and DR 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
drop table hadrStatusActivePath
go
drop table hadrStatusSynchronization
go
```

2. Log into the primary, standby, and DR 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 procedure rs_update_lastcommit
go
drop procedure rs_update_threads
go
drop procedure rs_get_thread_seq
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_DR_logical_host_name>'
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_DR_logical_host_name>'
go
sp_hadr_admin dropgroup, '<group_name>'
go
sp_configure 'HADR mode',-1
go

```

5. Log into the DR 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_DR_logical_host_name>'
go
sp_hadr_admin dropgroup, '<group_name>'
go
sp_configure 'HADR mode',-1
go

```

6. Shut down Replication Server on the primary, standby, and DR sites.
7. Remove the Replication Server instance directory from the primary site at `$$SYBASE/DM/CID_REP_primary_logical_host_name`.
8. Remove the Replication Server instance directory from the standby site at `$$SYBASE/DM/CID_REP_standby_logical_host_name`.
9. Remove the Replication Server instance directory from the DR site at `$$SYBASE/DM/CID_REP_DR_logical_host_name`.
10. Shut down RMA on the primary, standby, and DR sites.
11. 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/*
```

12. Restart RMA.

5 HADR Reference

This chapter lists the commands, system procedures, and proxy tables that can be used to administer the HADR system with DR node.

5.1 RMA Command Reference

Use the Replication Management Agent (RMA) commands to administer, monitor, and modify the properties of your replication environment.

5.1.1 sap_disable_replication

Disables replication from an SAP ASE HA cluster to an SAP ASE DR node.

Syntax

```
sap_disable_replication <logical_host_name_for_primary> [,  
<logical_host_name_for_ha> | <logical_host_name_for_dr>][, <database_name>]
```

Parameters

<logical_host_name_for_primary>

The name of the logical host that identifies the primary site.

<logical_host_name_for_ha>

The name of the logical host that identifies the HA site.

<logical_host_name_for_dr>

The name of the logical host that identifies the DR site.

<database_name>

The name of the database.

Examples

Example 1

Disables replication from primary host PR (SAP ASE HA cluster) to DR (SAP ASE DR) node for database 'db1':

```
sap_disable_replication PR,DR,db1
go
```

5.1.2 sap_enable_replication

Enables replication from an SAP ASE HA cluster to an SAP ASE DR node.

Syntax

```
sap_enable_replication <logical_host_name_for_primary>,
<logical_host_name_for_dr> [, <database_name>]
```

Parameters

<logical_host_name_for_primary>

The name of the logical host that identifies the primary site.

<logical_host_name_for_dr>

The name of the logical host that identifies the DR site.

<database_name>

The name of the database.

Examples

Example 1

Enables replication from primary host PR (SAP ASE HA cluster) to DR (SAP ASE DR) node for database db1:

```
sap_enable_replication PR,DR,db1
go
```


5.1.3 sap_failover

Use the `sap_failover` command to perform planned and unplanned failovers.

Syntax

Variable declaration:

```
sap_failover <logical_host_name_for_primary>, <logical_host_name_for_standby>,  
<time_out> [force | ,unplanned]
```

Parameters

<logical_host_name_for_primary>

The name of the logical host that identifies the primary site.

<logical_host_name_for_standby>

The name of the logical host that identifies the standby site.

<time_out>

Specifies the number of seconds the process waits 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).

unplanned

(Optional) Specifies an unplanned failover.

Examples

Example 1

Performs planned failover to designate standby site as the new primary site:

```
sap_failover PR, HA, 120  
go
```

Checks the status:

```
sap_status task  
go
```



```
Failover      Current Task Number
15

Failover      Total Number of Tasks
15

Failover      Task Start                Fri Mar 18 03:26:39 EDT
2016

Failover      Task End                  Fri Mar 18 03:26:44 EDT
2016

Failover      Hostname
site0

(12 rows affected)
```

Usage

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 finish before you issue a second.

See [Managing Failover in the HADR Cluster with the DR Node \[page 32\]](#)

5.1.4 sap_failover_drain_to_dr

Use the `sap_failover_drain_to_dr` command while performing failover inside an HADR with DR node environment. The `sap_failover_drain_to_dr` command makes sure that the incremental backlog from the HA cluster is drained to the DR node.

Syntax

```
sap_failover_drain_to_dr <time_out> | skip
```

Parameters

<time_out>

Specifies the number of seconds the command waits for the remaining backlog to be fully applied to the DR node. If this timeout is reached and draining backlog to DR node is not finished, the `sap_failover_drain_to_dr` command reports an error, retry this command with a higher `<time_out>` value.

skip

Causes the failover process to continue without applying the remaining backlog to the DR node. The `skip` keyword drops the subscription to the DR node, and the DR node must be rematerialized after the primary host is available.

Examples

Example 1

Transfers the transaction backlog to the DR node:

```
sap_failover_drain_to_dr 120
go
```

Checks the status:

```
1> sap_status
2> go
TASKNAME          TYPE
VALUE
-----
-----
Status           Start Time          Fri Mar 18 03:28:19 EDT
2016
```

```

Status                               Elapsed Time
00:00:31

FailoverRemaining Task Name           Failover
remaining.

FailoverRemaining Task State
Completed

FailoverRemaining Short Description   Failover remaining deactivate old
replication path and activate new replication path for all hosts.
FailoverRemaining Long Description   Verifying databases are not in data
loss mode.
FailoverRemaining Additional Info     Please run command
'sap_host_available PR' to complete disabling replication from the old
source.
FailoverRemaining Current Task Number
7

FailoverRemaining Total Number of Tasks
7

FailoverRemaining Task Start           Fri Mar 18 03:28:19 EDT
2016
FailoverRemaining Task End             Fri Mar 18 03:28:50 EDT
2016
FailoverRemaining Hostname
site0

(12 rows affected)

```

Example 2

Skips the transfer of the transaction backlog to the DR node, and you need not execute the `sap_host_available DR` command:

```

1> sap_failover_drain_to_dr skip
2> go

```

The returned result is:

```

TASKNAME                               TYPE
VALUE
-----
-----
Failover drain to DR. Start Time         Mon May 23 05:12:31 EDT
2016

Failover drain to DR. Elapsed Time
00:00:01

DRExecutorImpl Task Name           Failover drain to
DR.

DRExecutorImpl Task State
Completed

DRExecutorImpl Short Description   Failover drain to DR deactivate old
replication path and activate new replication path for all
hosts.
DRExecutorImpl Long Description   Started task 'Failover drain to DR.'
asynchronously.

```

```

DRExecutorImpl      Additional Info  Please execute command 'sap_status
task' to determine when task 'Failover drain to DR.' is
complete.
  FailoverRemaining  Task Name          Failover drain to
DR.

  FailoverRemaining  Task State
Completed

  FailoverRemaining  Short Description  Failover drain to DR deactivate old
replication path and activate new replication path for all
hosts.
  FailoverRemaining  Long Description  The DR host 'DR' will be skipped
from the failover processing, please re-materialize it after
sap_host_available 'PR' is done.
  FailoverRemaining  Task Start        Mon May 23 05:12:31 EDT
2016

  FailoverRemaining  Task End          Mon May 23 05:12:32 EDT
2016

  FailoverRemaining  Hostname
site0

(14 rows affected)

```

For a new primary host, execute the `sap_host_available PR` command. Re-materialize the DR node after you execute the `sap_host_available PR`:

```

sap_host_available PR
go

```

The returned result is:

```

1> sap_host_available PR
2> go
  TASKNAME          TYPE
VALUE
-----
-----
  HostAvailable Start Time          Mon May 23 05:29:32 EDT
2016
  HostAvailable Elapsed Time
00:02:05

  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
11

  HostAvailable Total Number of Tasks
10

  HostAvailable Task Start          Mon May 23 05:29:32 EDT
2016

```

```

HostAvailable Task End      Mon May 23 05:31:37 EDT
2016
HostAvailable Hostname
site0

(11 rows affected)

```

5.1.5 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.

Syntax

```
sap_host_available <logical_host_name>
```

Parameters

<logical_host_name>

Specifies the logical host name of the DR node.

Examples

Example 1

Re-establishes the replication path from the new primary SAP ASE to the new standby SAP ASE inside the HA cluster:

```
sap_host_available PR
go
```

The returned result is:

```

1> sap_host_available PR
2> go
TASKNAME      TYPE
VALUE
-----
HostAvailable Start Time      Fri Mar 18 03:29:32 EDT
2016

```



```

HostAvailable Elapsed Time
00:02:05

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
11

HostAvailable Total Number of Tasks
10

HostAvailable Task Start              Fri Mar 18 03:29:32 EDT
2016
HostAvailable Task End                Fri Mar 18 03:31:37 EDT
2016
HostAvailable Hostname
site0

(11 rows affected)

```

Example 2

Re-establishes the replication path from the HA cluster to the DR node:

```

sap_host_available DR
go

```

The returned result is:

```

1> sap_host_available DR
2> go
TASKNAME      TYPE
VALUE
-----
-----
HostAvailable Start Time          Fri Mar 18 03:32:52 EDT
2016
HostAvailable Elapsed Time
00:00:22

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 'DR' receiving replication from logical host 'HA'.
HostAvailable Current Task Number
9

HostAvailable Total Number of Tasks
9

```

```

HostAvailable Task Start      Fri Mar 18 03:32:52 EDT
2016
HostAvailable Task End      Fri Mar 18 03:33:14 EDT
2016
HostAvailable Hostname
site0
(11 rows affected)

```

Usage

After the completion of the `sap_failover` and the `sap_failover_drain_to_dr` commands, use the `sap_host_available` command to:

- Re-establish the replication path from the new primary SAP ASE to the new standby SAP ASE inside the HA cluster.
- Re-establish the replication path from the HA cluster to the DR node.

Note

The order in which the two replication paths are re-established is not relevant.

5.1.6 sap_set_host

Use the `sap_set_host` command to register the SAP ASE DR logical host. The logical host consists of an SAP ASE server, a Replication Server, and an RMA.

Syntax

```

sap_set_host <logical_host_name_for_dr>, <ASEDR_dataserver_host_name>,
<ASEDR_dataserver_server_port_number>,
<ASE_DR_repserver_host_name>, <ASE_DR_repserver_port_number>, <ASE_DR_rma_port_number>

```

Parameters

<logical_host_name_for_dr>

Specifies the logical host name for the DR.

<ASEDR_dataserver_host_name>

Specifies the host name for the SAP ASE DR data server.

<ASEDR_dataserver_server_port_number>

Specifies the server port number for the SAP ASE DR data server.

<ASE_DR_repserver_host_name>

Specifies the host name for the SAP ASE DR Replication Server.

<ASE_DR_repserver_port_number>

Specifies the server port number for the SAP ASE DR Replication Server.

<ASE_DR_rma_port_number>

Specifies the port number for the SAP ASE DR RMA.

Examples

Example 1

Registers DR node running on site0 with logical host name DR, SAP ASE host name site0 and listening on port 5000, Replication Server running on site0 and listening on port 5005, and RMA listening on port 7001:

```
sap_set_host DR, site0, 5000, site0, 5005, 7001
```

Usage

Execute this command on the RMA running inside the HA cluster.

5.1.7 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
DR name.	DR HADR Status	DR Standby : Inactive	Identify the
primary and standby sites.	DR Synchronization Mode	Asynchronous	The configured
DR Synchronization Mode value.	DR Synchronization State	Inactive	
DR Synchronization Mode in which replication is currently operating.	DR Distribution Mode	Local	Configured
DR 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
HA 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
HA value for the distribution_mode replication model property.	HA Replication Server Status	Active	The status of
HA Replication Server.	PR Hostname	site0	Logical host
PR name.	PR HADR Status	Primary : Active	Identify the
PR primary and standby sites.	PR Synchronization Mode	Synchronous	The configured
PR Synchronization Mode value.	PR Synchronization State	Synchronous	
PR Synchronization Mode in which replication is currently operating.	PR Distribution Mode	Remote	Configured
PR value for the distribution_mode replication model property.	PR Replication Server Status	Active	The status of
PR Replication Server.	HA.DR.PI2 State	Suspended	Path is
HA.DR.PI2 suspended (Replication Agent Thread). Transactions are not being replicated.	HA.DR.PI2 Latency Time	Unknown	No latency
HA.DR.PI2 information for database 'PI2'.			

HA.DR.PI2	Latency information for database 'PI2'.	Unknown	No latency
HA.DR.PI2	Commit Time time for the database 'PI2'.	Unknown	No last commit
HA.DR.PI2	Distribution Path Replication Server through which transactions travel.	PR	The path of
HA.DR.PI2	Drain Status status of the transaction logs of the primary database server.	Unknown	The drain
HA.DR.db1	State suspended (Replication Agent Thread). Transactions are not being replicated.	Suspended	Path is
HA.DR.db1	Latency Time information for database 'db1'.	Unknown	No latency
HA.DR.db1	Latency information for database 'db1'.	Unknown	No latency
HA.DR.db1	Commit Time time for the database 'db1'.	Unknown	No last commit
HA.DR.db1	Distribution Path Replication Server through which transactions travel.	PR	The path of
HA.DR.db1	Drain Status status of the transaction logs of the primary database server.	Unknown	The drain
HA.DR.master	State suspended (Replication Agent Thread). Transactions are not being replicated.	Suspended	Path is
HA.DR.master	Latency Time information for database 'master'.	Unknown	No latency
HA.DR.master	Latency information for database 'master'.	Unknown	No latency
HA.DR.master	Commit Time time for the database 'master'.	Unknown	No last commit
HA.DR.master	Distribution Path Replication Server through which transactions travel.	PR	The path of
HA.DR.master	Drain Status status of the transaction logs of the primary database server.	Unknown	The drain
HA.PR.PI2	State suspended (Replication Agent Thread). Transactions are not being replicated.	Suspended	Path is
HA.PR.PI2	Latency Time information for database 'PI2'.	Unknown	No latency
HA.PR.PI2	Latency information for database 'PI2'.	Unknown	No latency
HA.PR.PI2	Commit Time time for the database 'PI2'.	Unknown	No last commit
HA.PR.PI2	Distribution Path Replication Server through which transactions travel.	PR	The path of
HA.PR.PI2	Drain Status status of the transaction logs of the primary database server.	Unknown	The drain
HA.PR.db1	State suspended (Replication Agent Thread). Transactions are not being replicated.	Suspended	Path is
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	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	HA	The path of
PR.DR.PI2	Drain Status	Not Drained	The drain
PR.DR.db1	State	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	HA	The path of
PR.DR.db1	Drain Status	Not Drained	The drain
PR.DR.master	State	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	HA	The path of
PR.DR.master	Drain Status	Not Drained	The drain
PR.HA.PI2	State	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	HA	The path of
PR.HA.PI2	Drain Status	Not Drained	The drain
PR.HA.db1	State	Active	Path is active
PR.HA.db1	Latency Time last calculated	2016-04-27 22:33:00.646	Time latency
PR.HA.db1	Latency (ms)	610	Latency
PR.HA.db1	Commit Time commit replicated	2016-04-27 22:33:00.646	Time last
PR.HA.db1	Distribution Path	HA	The path of
PR.HA.db1	Drain Status	Not Drained	The drain
PR.HA.master	State	Active	Path is active
PR.HA.master	Latency Time last calculated	2016-04-27 22:33:00.632	Time latency
PR.HA.master	Latency (ms)	603	Latency
PR.HA.master	Commit Time commit replicated	2016-04-27 22:33:00.632	Time last

```

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 1: 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
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.
<div style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <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	One of two replication distribution modes you have configured between a database and the Replication Server: <ul style="list-style-type: none"> • Local • Remote
Replication Server Status	The status of the Replication Server, which can be one of: <ul style="list-style-type: none"> • Active • Down • Unknown

Information	Description
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	The timestamp of the most recent <code>trace</code> command that was applied to the target database and used to calculate latency.
Latency	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.
<div style="background-color: #f0f0f0; padding: 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	The local timestamp of a command applied to the target database.
Distribution Path	The logical host name of the distribution target server.
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.

5.1.8 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
-----
22:38:20.891  Start Time                2016-06-02
00:00:02      Elapsed Time
-1            Estimated Failover Time
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 2: 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.
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: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> <p>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.

Information	Description	Value
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.

5.1.9 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    0
PR.DR.master 2      RAT      T          NULL     63
site0        Active    NULL
PR.DR.master 3      RATCI    T          NULL     NULL
site1        Active (Active)  NULL
```

PR.DR.master 4		SPQ	Q	106	NULL
site1	NULL		0		
PR.DR.master 5		CAP	T	NULL	53
site1	Active (Awaiting Command)		NULL		
PR.DR.master 6		SQM	T	NULL	22
site1	Active (Awaiting Message)		NULL		
PR.DR.master 7		IBQ	Q	106	NULL
site1	NULL		0		
PR.DR.master 8		SQT	T	NULL	73
site1	Active (Awaiting Wakeup)		NULL		
PR.DR.master 9		DIST	T	NULL	41
site1	Active (Awaiting Wakeup)		NULL		
PR.DR.master 10		SQM	T	NULL	95
site1	Active (Awaiting Message)		NULL		
PR.DR.master 11		RouteQ	Q	16777319	NULL
site1	NULL		0		
PR.DR.master 12		RSI	T	NULL	96
site1	Active (Awaiting Wakeup)		NULL		
PR.DR.master 13		SQM	T	NULL	96
site2	Active (Awaiting Message)		NULL		
PR.DR.master 14		OBQ	Q	116	NULL
site2	NULL		0		
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 3: 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.

Information	Description
State	<p>The status of the thread and server.</p> <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; margin-top: 10px;"> <p>i Note</p> <p>Threads also have some other specific states.</p> </div>
Backlog	<p>The accumulated logs to be processed. Displays 0 when there are no logs to be processed. Displays NULL when the information cannot be queried.</p> <div style="background-color: #f0f0f0; padding: 5px; margin-top: 10px;"> <p>i Note</p> <p>Backlogs are only available for queues and the primary ASE, so NULL is displayed for threads and the standby ASE.</p> </div>

Table 4: 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.

5.1.10 sap_update_replication {add | remove}

Use the `sap_update_replication` command to add or remove a DR host.

Syntax

```
sap_update_replication {add | remove}, <logical_host_name_for_dr>
```

Parameters

<logical_host_name_for_dr>

Specifies the name of the logical host to be added or removed.

add

Adds the DR node.

remove

Removes the DR node.

Examples

Example 1

Adds a DR node with the logical host name DR:

```
sap_update_replication add, DR
```

Example 2

Removes a DR node with the logical host name DR:

```
sap_update_replication remove DR
```

5.1.11 sap_update_replication {move_out | move_in}

Use the `sap_update_replication` command to move a standby host into or out from the HADR with DR node system.

Syntax

```
sap_update_replication {move_out, <standby_logical_host_name> | move_in, <standby_logical_host_name>, {start | finish [<data_drain_time>] }}
```

Parameters

`move_out, <standby_logical_host_name>`

Moves the specified standby host out from the replication system and switches the replication topology from remote (primary SAP ASE → standby SAP Replication Server → DR SAP Replication Server → DR SAP ASE) to local (primary SAP ASE → primary SAP Replication Server → DR SAP Replication Server → DR SAP ASE).

`move_in, <standby_logical_host_name>, start`

Creates the replication path from the primary host to the standby host using the local topology. That is, replicate data from the primary SAP ASE to the primary SAP Replication Server and then to SAP Replication Server on both the standby site and the DR site.

`move_in, <standby_logical_host_name>, finish [<data_drain_time>]`

Indicates the HADR with DR node system to switch the replication topology from local to remote. That is, replicate data from the primary SAP ASE to the standby SAP Replication Server and then to SAP Replication Server on the DR site.

`<data_drain_time>` indicates the maximum time (in seconds) allowed to drain the pending data to the standby SAP ASE and DR SAP ASE before switching the replication topology from local to remote. The default value is 1200 seconds. If the draining time exceeds the value of `<data_drain_time>`, the command returns an error and you can execute the command again to drain the remaining data.

Examples

Example 1

On a standby SAP Replication Server that is not running, this example removes the standby host from the HADR with DR node system and switches the replication topology from remote to local:

```
sap_update_replication move_out, site2
```

Example 2

On a standby SAP Replication Server that has been restored, this example creates the replication from the primary host to the standby host in the local topology:

```
sap_update_replication move_in, site2, start
```

Example 3

After `sap_update_replication start` is executed successfully, this example switches the replication topology from local to remote:

```
sap_update_replication move_in, site2, finish, 1000
```

5.2 sp_hadr_admin

`sp_hadr_admin` allows you to add a DR node.

Syntax

- Adds a DR node to the HADR member list:

```
sp_hadr_admin addserver_dr, '<HADR_server_name>' [, [<pname>] [, 'nopropagate' ]]
```

Parameters

`addserver_dr`

adds a DR server to the HADR system and member list.

Examples

Example 1

Adds a server named ROME to the member list:

```
sp_hadr_admin addserver_dr, ROME
(return status = 0)
Adding server 'ROMEDR', physical name 'ROME'
Server added.
Command 'addserver_dr' successful.
(1 row affected)
```



Usage

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.