WHAT'S NEW | PUBLIC
SAP Adaptive Server Enterprise 16.0 SP03
Document Version: 1.0 – 2020-03-04

What's New

THE BEST RUN SAP

Content

1 Version Update - 16.0 SP03 PL08

SAP Adaptive Server Enterprise (SAP ASE) version 16.0 SP03 PL08 introduces several enhancements.

Table 1: 16.0 SP03 PL08 Enhancements

Feature	Description
`plldegree` added in the output of `sp_help` and `sp_helpindex`	The `plldegree` displays the maximum number of threads the query optimizer can use for a table or each index in the table. See sp_help and sp_helpindex.
`monProcELC`	`monProcELC` provides statistics for the Engine Local Cache (ELC). See monProcELC for more details.
Support for Solaris Sparc in Hardware Security Module (HSM) external keystore	Encrypting keys using an external keystore is supported on Solaris Sparc starting from 16.0 SP03 PL08 release. See Key Protection Using External Keystore for more details about HSM.
The `optdiag` utility supports use of the secure user store	The `optdiag` utility adds the `-k` and `-H` options to support connections to the SAP ASE server using a user-generated key that retrieves connection information (user name, password, host, and port number) from the secure store. See Securely Automate Database Administration Tasks Using Secure User Store and optdiag for more details.
Enable FIPS in the backup server	You can set the `enable fips` configuration parameter to enable FIPS login password encryption in the backup server. See enable fips for more details.

Table 2: HADR New Features

Feature	Description
Support for replication of `installjava` and Java columns	The replication of `installjava` and Java columns is supported in stream replication.

Feature	Description
Disabling the checks for reference constraints when applying transactions	To improve performance, you can choose to bypass the checks for reference constraints when replicating data in an HADR environment by setting the `disable_referential_constraints` entry in the `setuphadr` response file from `false` to `true`. See Sample setup_hadr.rs Response File and Sample setup_hadr.rs Response File for Business Suite and rs_referential_integrity for more details.
Purging outdated information in the `rs_ticket_history` table	you can use the `sap_purge_trace` command to purge `rs_ticket` information inserted in the `rs_ticket_history` table. See sap_purge_trace for more details.
Device space management	You can add and drop device spaces, and also view device information using the enhanced `sap_add_device` command and the newly added `sap_drop_device` and `sap_list_device` commands. See sap_add_device and sap_drop_device and sap_list_device for more details.

2 Version Update - 16.0 SP03 PL07

SAP Adaptive Server Enterprise (SAP ASE) version 16.0 SP03 PL07 introduces several enhancements.

Table 3: 16.0 SP03 PL07 Enhancements

Feature	Description
View SSL certificate expiration date.	You can use the `monSSLCertInfo` monitoring table to view the expiration date of the SSL certificate on the SAP ASE server. See Viewing a Certificate Expiration Date for more details.
Audit the truncation of audit tables.	In addition to audit event 64, SAP ASE now also uses audit event 58 to track the truncation of audit tables. See Audit Event Values for more details.
`create_table` parameter in the `sp_dump_history` command	Use this parameter to create a proxy table to allow access to the dump history file. See sp_dump_history and Creating a Proxy Table.
Enable latch-free indexes at the database level	The `latch_free_index` option on the `create database` or `alter database` command enables or disables latch-free indexes database-wide. See create database, alter database, Using Latch-Free Indexes.
`set retaindays <number_days>` Backup Server configuration parameter	Use this parameter to specify the number of days to retain each database or transaction dump on the backup server. See backupserver and Prevent Dump Files from Being Overwritten.
`HCB index auto tuning` enhancement	`HCB index auto tuning` re-enables hash caching if it is previously automatically disabled. See HCB index auto tuning.
`set statistics io` enhanced to display usage for hash-cached Btree (HCB)-enabled indexes	The `set statistics io` command reports the number of hash cache lookups and hits for HCB-enabled indexes on data row cached (DRC) tables. See Statistics for In-Memory Row Storage.

Feature	Description
Dumping and loading in-memory databases enhancement	The dump and load of an in-memory database is supported when using single-signed-on authentication mechanisms, such as Kerberos. See Dumping and Loading In-Memory Databases.
The `wlacliutil` Workload Analyzer command line utility is available	Use the `wlacliutil` utility to capture, analyze, and replay workloads. See Workload Analyzer Command Line Utility.
The `infmsg` attribute in the `SYS_SESSION` system application context.	You can use the `infmsg` attribute in the `SYS_SESSION` system application context to store the informational messages in a string. See Using infmsg to Get the Informational Messages for more details.
OData Server is upgraded to OData 17.	OData 16 is end of life and is replaced with OData 17.
Limiting the number of elements in the `in` clause.	You can set the `max number of IN elements` configuration parameter to limit the number of elements in the largest `in` clause in a query. See Limiting the Number of Elements in the in Clause for more details.
Enhancement to the `sybdiag` utility.	The `-S` option in the `sybdiag` utility is enhanced by adding the `ssl:` prefix to support collecting diagnostic data for SSL enabled SAP ASE servers. See sybdiag for more details.

Table 4: HADR New Features

Feature	Description
Modify the connection timeout value.	You can use the `HADR connect timeout` configuration parameter to modify the connection timeout value. See HADR connect timeout for more details.
Modify the remote query timeout value.	You can use the `HADR remote query timeout` configuration parameter to modify the remote query timeout value. See HADR remote query timeout for more details.

Feature	Description
Set the memory limit for SAP Replication Server	You can use the `sap_set memory_size` command to set the memory limit for SAP Replication Server instances on all HADR nodes. You can set this property at any time, either before or after the HADR system is setup. If unset, RMA tunes the memory limit for SAP Replication Server automatically. See sap_set memory_size for more details.
New profile parameter for Fault Manager.	You can use the `ha/syb/run_odbc_threaded` profile parameter to run the shallow probes in a special thread. See Parameters that Affect the Fault Manager's Actions for more details.

3 Version Update - 16.0 SP03 PL06

SAP Adaptive Server Enterprise (SAP ASE) version 16.0 SP03 PL06 introduces several enhancements.

Table 5: 16.0 SP03 PL06 Enhancements

Feature	Description
`create table <table_name> with row_caching {on [ allow [<number> | default] ] | [for [ all | default ] ] rows ] | off}` `alter table <table_name> set row_caching {on [ allow [<number> | default] ] | [for [ all | default ] ] rows ] | off}`	You can configure the number of rows a single query can move to the IMRS (called cached data rows). `default` indicates that you are reverting to the default behavior (the default value is 100). `for all` means that all the accessed rows are moved to the IMRS. See Determining the Number of Rows a Single Query Can Move to the IMRS.
Commands that take an EX_TAB lock do not take a statement registration in IMRS-enabled databases.	When a command (for example, `reorg rebuild` in offline mode) takes an EX_TAB lock, it specifies the following: • There are no older scanners active on the object. • The utility is working on the latest version. • The IMRS pack operation cannot create a new version because the command has an EX_TAB lock. See Statement Registration and the Garbage Collector.
`sp_imrs 'pack_rows', ['aus', '<number_of_allocation_units>' | 'pct', '<percentage_of_allocation_unit_to_pack>']`	Allows you to specify the rows you want to manually pack out of the IMRS. See Using Pack Operations to Move IMRS Rows to the Page Store and sp_imrs.
`kerberos null password change` configuration parameter	Configures SAP ASE to allow Kerberos administrators to pass caller passwords as null to reset login passwords with `alter login`. See allow kerberos null password.
IMRS-enabled databases support `load transaction ... until_time` parameter	You can use load transaction with the `until_time` parameter to load an IMRS-enabled database. However, the `until_time` parameter breaks the load sequence, and it cannot be followed by `load transaction` commands. See Dumping and Loading into an IMRS-Enabled Database.

Feature	Description
`show_plan(<DBID>, <procedure_ID>, <plan_ID> [, <statement_number>])`	The `show_plan` function adds syntax for determining information about stored procedure plans using database IDs, procedure IDs, plan IDs and the statement numbers. See show_plan.
Single user mode not required for `alter database [row | version] storage off`	Previous versions of SAP ASE required the database to be in single-user mode when you issued these commands: • `alter database <database_name> row storage off` • `alter database <database_name> version storage off` SAP ASE version 16.0 SPS03 PL06 and later do not require the database to be in single-user mode when you issue these commands. See Dropping IMRS Devices.
`syslogininfo`	Adds the `syslogininfo` view, which returns information for all logins. See syslogininfo.
`sp_dump_history @operation='@dmp_end_time'`	`sp_dump_history` allows you to specify operations that apply up to a specific date. See sp_dump_history.
`sp_procxmode sproc_name [enable_dc | disable_dc]`	`sp_procxmode` • `enable_dc` – enables deferred compilation for the stored procedure. • `disable_dc` – disables deferred compilation for the stored procedure. See sp_procxmode.
`update` and `delete` commands transparently redirect to page store in IMRS-enabled databases.	`update` and `delete` commands in versions earlier than 16.0 SP03 PL06 could internally fall back and issue error 728 due to insufficient memory. In this and later versions, these commands transparently redirect to page store. See Rectifying Lack of Memory in the IMRS Cache.

Feature	Description
`statistics io` enhanced to display IMRS-related statistics	`statistics io` reports: • The number of rows read, written, or migrated in IMRS cache and page-store for data row cached (DRC) and MVCC-enabled tables. • The number of logical and physical I/Os for the `sysversions` table for on-disk MVCC-enabled tables. See Statistics for In-Memory Row Storage.
`show_prepared_statements`	Allows you to view the SQL text of prepared statements. See show_prepared_statements.
`sp_shrink`	Allows you to shrink a device or database to reuse or delete unused space. See sp_shrink.
`limit` and `offset` options in the `union` statement	The `union` statement supports the `limit` and `offset` options in either an individual query or the final result set. See union operator.
Web Services on SAP ASE PL06 uses JRE version 8.	Web Services for SAP ASE is updated to prevent security vulnerabilities, requiring that you set the SAP_JRE8 environment variable to JRE8 or later.
`sp_ssladmin refreshcert`	Allows you to activate a certificate on an SAP ASE server without having to restart the server. See sp_ssladmin and Updating an Expired Certificate.
SAP ASE supports the encryption of the master key with an external Hardware Security Module (HSM) key.	HSM is a physical device that provides secure storage for encryption keys. SAP ASE can use a single HSM key per SAP ASE instance to encrypt or decrypt master keys present in any databases. See Key Protection Using External Keystore.

Feature	Description
SAP ASE now ships JAVA RUNTIME (JRE) 8.1.046 and 7.1.060.	The following JRE versions are shipped with SAP ASE. No action is necessary, as both versions ship with the product, install automatically, and run simultaneously on the same machine. Table 6: JRE Support **Product/Feature** / **Version** / **Location** PCI/JVM / SAP JRE 8.1.046 / `$SYBASE/shared/ase/` Cockpit/Utilities / SAP JRE 8.1.046 and SAP JRE 7.1.060 / `$SYBASE/shared/` Installer/uninstaller / SAP JRE 8.1.046 / `$SYBASE/jre64/` **i Note** SAP Support is available for product installations running SAP JRE; third-party JREs are not supported. See SAP Java Runtime Environment.

Table 7: HADR New Features

Feature	Description
SQL statement replication	Both HADR systems and HADR with DR node systems support SQL statement replication. See SQL Statement Replication in the HADR Users Guide and SQL Statement Replication in the HADR System with DR Node Users Guide.
`sp_configure 'stream rep msg channel timeout'`	`sp_configure 'stream rep msg channel timeout'` allows you to set a timeout value for the message channel from RepAgent to SAP Replication Server. See sp_configure 'stream rep msg channel timeout'.
Table level replication and table level exclusion replication	Table level replication is supported when replicating data from an HADR system to an external system and vice versa. Table level exclusion replication is supported in HADR systems, HADR with external systems, and HADR with DR node systems.

Feature	Description
Resume replication to external system after standby SAP Replication Server is down.	You can use the automatic method to configure the primary RepAgent to connect to the external SAP Replication Server to enable data replication from the primary SAP ASE to the external system when the standby host is down. See Configuring Primary SAP ASE to Replicate Data to External System When the Standby Host is Down.

4 Version Update - 16.0 SP03 PL05

SAP Adaptive Server Enterprise (SAP ASE) version 16.0 SP03 PL05 introduces several enhancements.

Table 8: 16.0 SP03 PL05 Enhancements

Feature	Description
Workload profiler support for on-disk multiversion concurrency control (MVCC)	The workload profiler is enhanced to provide an estimation for in-memory row storage (IMRS) cache and the version storage size required to enable on-disk MVCC for a database.
During a checkpoint, SAP ASE flushes the changed `imrslog` buffers to avoid hangs.	During a checkpoint, SAP ASE ignores group commit operations and flushes the changed `imrslog` buffers, even if more changes are expected on the same buffer.
Parameters added to `sp_dump_history`	The parameters are: • `@operation = [listpurgefiles | purgefiles | delete]` • `version` • `stripe_name` • `dump_date` • `format`
Group commits	SAP ASE performs a group commit for `sysimrslogs` by default.
`ExportPrivateKey` utility	Adds the `-keypass` parameter, which specifies a password to access the private key. If you do not specify `-keypass`, the password used for the `keystore` is used to access the private key.
Resolving spinlock contention resolution	SAP ASE uses lockless data cache latches to resolve buffer-level spinlock contention.

Table 9: HADR New Features

Feature	Description
The HADR with DR node system allows you to recover replication when the standby host is down.	When the standby host is down in an HADR with DR node system, you can still recover the replication path from the primary site to the DR site. See Recovering Replication to DR Node When the Standby SAP Replication Server is Down

Feature	Description
SSL on HADR supports for Business Suite.	SSL on HADR supports Business Suite from the 16.0 SP03 PL05 release. **Note** SSL is supported on Linux x64, Linux x86-64, and Windows x64. See HADR Security.

5 Version Update - 16.0 SP03 PL04

SAP Adaptive Server Enterprise (SAP ASE) version 16.0 SP03 PL04 introduces several enhancements.

For additional information, see the readme files in the Release Notes Information https://launchpad.support.sap.com/#/notes/2643576.

Table 10: 16.0 SP03 PL04 Enhancements

Feature	Description
SAP JRE	The SAP Java Runtime Environment (JRE) versions have been upgraded to SAP JRE versions 7.1.055 and 8.1.037. No action is necessary, as both versions ship with the product, install automatically, and run simultaneously on the same machine. **i Note** SAP Support is available for product installations running SAP JRE; third-party JREs are not supported.
`sp_logiosize` adds the `sysimrslogs` parameter	`sp_logiosize` allows you to update the size of the I/O for the `syslogs` and `sysimrslogs` tables. See sp_logiosize.
`sp_imrs`	The command `sp_imrs clear, 'blocking_spid' 'db_name'` can now remove all types of blocking SPIDs from the system, and not just infected or rogue SPIDs. Also, the output of the command `sp_imrs show, 'blocking_spid' 'db_name'` has been enhanced with corrective measures related to unblocking the system, among others. See sp_imrs.
`sp_altermessage` adds the `<severity_level>` parameter	`<severity_level>` allows you to specify that the server writes all error message of a specific severity level to the error log and displays this error to the client. See sp_altermessage.
`decompression row threshold` configuration parameter	Determines the maximum number of columns in a table that are uncompressed using column decompression. See decompression row threshold.
Changed memory requirements for in-memory databases	SAP ASE version 16.0 SP03 PL04 has a memory footprint overhead per version of 64 bytes. However, when a DML operation (for instance, `update`) triggers the migration of a row to the in-memory row storage (IMRS), the row version takes up 96 bytes of the IMRS. See On-Disk Multiversion Concurrency Control.

Feature	Description
Backup Server	Backup Server gets the following enhancements: • Parameters and program options: The `backupserver` executable now accepts a new set of parameters that can be passed in normal and extended option modes for argument binding. • Configuration parameters and configuration file: New configuration options/parameters have been added, and they are stored permanently in a new configuration file. • Administrator user and password: A new administrator user ('sa' by default) has been added, and is the only user authorized to display or make changes to any Backup Server options. The password for this user is encrypted and saved in the new configuration file. • Language Handler: The new language handler enables Backup Server options to be passed as command-line parameters, or changed through `isql`. • SSL support: Backup server supports SSL for communication with SAP ASE, and for data transfer with other backup servers, when using remote dumps and/or loads. See Enabling SSL in the Backup Server, and backupserver.
Replication in in-memory row storage	All DML operations are now performed in in-memory row storage if replication is enabled.
Securely automate SAP ASE server administration tasks using secure user store feature	Secure user store allows you to automate database administration tasks without compromising data confidentiality by connecting to an SAP ASE server without entering a username and password in the command line. With secure user store, you can establish the connection between a client application and an SAP ASE server using a user-generated key that retrieves the connection information (username, password, host, and port number) from the secure store. You can use the `aseuserstore` utility to insert, list, and delete the connection information in the form of user keys in the secure store. The utility supports ADO.NET, ODBC, jConnect, and SAP ASE native client applications, as well as the SAP ASE utilities `isql`, `bcp`, and `defncopy`. See Securely Automate Database Administration Tasks Using Secure User Store.
Temporary precomputed result set	The temporary precomputed result set stores query redundant views that are used multiple times by the same query or by other frequently executed queries. The temporary precomputed result set is refreshed when query execution starts and stays valid until the data changes or the user session ends. See Temporary Precomputed Result Sets.
Apply resource limits for all users using any application	In addition to adding, modifying, and dropping resource limits for a specific user or application, you can also apply these to all users using any application by specifying `NULL` for `<name>` and `<apppname>` in the following system stored procedures: • sp_add_resource_limit • sp_modify_resouce_limit • sp_drop_resource_limit

Feature	Description
`monRepLogActivity` table adds the `CachedBeforeImageUsed` column.	RepAgent thread now retrieves before-image row values from the RepAgent thread memory, and puts the number of times it retrieves these values from the memory into the `CachedBeforeImageUsed` column. This improves performance by saving RepAgent from having to perform an auxiliary log scan. See monRepLogActivity.
`drop index` adds the `with nowait` parameter.	The `with nowait` option of the `drop index` command allows you to drop indexes instantaneously. See drop index.

HADR New Features

- HADR supports the replication of databases enabled for in-memory row storage (IMRS). The following changes are introduced to support IMRS.

Feature	Description
RepAgent support for the `sysimrslogs` scanner	The `sysimrslogs` scanner reads logs of IMRS-enabled databases. See Monitoring the Scanners.
System table support for the `sysimrslogs` scanner	The following system tables provide diagnostic data for the `sysimrslogs` scanner: ○ monRepSchemaCache ○ monRepScannersTotalTime ○ monRepLogActivity ○ monRepStreamStatistics ○ monRepScanners

- With SAP Replication Server automatic memory management, you can replicate multiple databases (certificated with 20) at the same time with low total cost of ownership (TCO) and simple configuration steps, which ensures you to deploy and manage multiple pairs of databases in the HADR environment. See Automatic Memory Management for more details.
- RepAgent adds two new configuration parameters to control whether to slow down user tasks at RepAgent startup:
  - `max user task slowdown`
  - `initial log scan percent`

  See RepAgent Configuration Parameters for details.
- The HADR system allows you to configure the primary SAP ASE, so that it can connect to the external replication system directly and resume data replication when the standby host is down, without suffering data loss. See Configuring Primary SAP ASE to Replicate Data to External System When the Standby Host is Down.
- When replicating out of an HADR system, you can change the default HADR administrator user name (DR_admin) to your own login name with which you set up the HADR through the

`cap_filter_dr_admin_name` configuration parameter. This ensures that the external SAP Replication Server can filter out all commands issued under your login name. See Replication Server Configuration Parameters for details.

- HADR system adds the new `removehadr` utility to clean up an HADR environment. See Cleaning Up an HADR Environment Using the removehadr Utility for details.

6 Version Update - 16.0 SP03 PL03

SAP ASE version 16.0 SP03 PL03 introduces a number of enhancements.

For additional information, see the readme files in the Release Notes Information https://launchpad.support.sap.com/#/notes/2550339.

Table 11: SAP ASE 16.0 SP03 PL03 Enhancements

Feature	Description

The `password_random` function includes a new parameter for creating limitations on a random password.	The `password_random` function includes a new parameter that you can use to specify a sub-string that is not to be used as part of a random password when the `disallow simple passwords` password complexity check is set to true.
	For example, to generate a random password that does not contain the sub-string `am3`, execute:
	`disallow simple passwords:  1`
	`minimum password length:  3`
	`select password_random(4,"am3")`
	See password_random.

Allow the movement of the first 256 pages of a database to another device	When you shrink a database with the `alter database ... off` command, you can move the logic pages 0–255 (the first allocation unit of a database) to another device. See alter database and Shrinking Databases.
`set statistics ioplan {on | off}`	This new command, when set to on, displays the IO statistics per plan for each view. See set.
Inline table UDFs can be converted to parameterized views.	Either of these methods internally changes inline table user-defined functions to parameterized views: • Specify the `inline` parameter when using the `create function` command. • Enabling the `inline table functions` configuration parameter. Inline table UDFs eliminates run-time overhead and improves execution performance. See Parameterized Views.
The `sysrvroles` table adds status bits 8 and 16.	These status bits determine if the user defined role being activated is mutually exclusive. See syssrvroles.
Sessions in MVCC preallocated memory.	When using multiversion concurrency control (MVCC), the server preallocates memory from the global fragment pool when a session starts, so new statements need not allocate or free the memory. This memory is freed when the session is closed. See Multiversion Concurrency Control

Feature	Description
SAP ASE binaries are signed with SHA-256 encryption	Use of SHA-256 hashing algorithm to digitally sign SAP ASE executables: To provide enhanced security and integrity of the executable files, now SAP ASE binaries are signed with a stronger hashing algorithm – SHA-256.
Memory footprint overhead optimized for on-disk MVCC	The memory footprint overhead per version for on-disk multiversion concurrency control (on-disk MVCC) has been reduced to 96 bytes. See On-Disk Multiversion Concurrency Control.
Enhanced preemption techniques implemented for the IMRS garbage collector	Long-running, and inactive transactions are now preempted by the in-memory row storage (IMRS) garbage collector. This enables the garbage collector to clear unused versions from the IMRS memory. See Using the Garbage Collector to Reclaim Memory.
Adds three system procedures for SAP ASE monitoring: • `sp_w` • `sp_sp` • `sp_p`	These new system procedures are the simplified versions of the existing `sp_who` and `sp_showplan` system procedures. `sp_w` returns information about all current SAP ASE users and processes. `sp_sp` and `sp_p` display the concise and extended output of the query plan respectively. See the commands sp_w, sp_sp, and sp_p.
HADR adds a standalone DR node.	SAP ASE HADR with DR node system adds a standalone third node as the DR server to an existing HADR cluster. See HADR System with DR Node Users Guide

<table>
<tr><th>Feature</th><th>Description</th></tr>
<tr><td>Both the HADR pair and DR node support up to five customized directories for database, transfer log, log, configuration, and backup files.</td><td>
i Note
This functionality applies to the Linux platforms only.

You can configure the following parameters in the response file to customize the file directories:

<table>
<tr><th>Parameter</th><th>Description</th><th>Sample Value</th></tr>
<tr><td>dm_database_file_directory</td><td>The database file directory</td><td>$SYBASE/<SID>/database</td></tr>
<tr><td>dm_translog_file_directory</td><td>The translog file directory</td><td>$SYBASE/<SID>/translog</td></tr>
<tr><td>dm_log_file_directory</td><td>The log file directory</td><td>$SYBASE/<SID>/log</td></tr>
<tr><td>dm_config_file_directory</td><td>The configuration file directory</td><td>$SYBASE/<SID>/cfg</td></tr>
<tr><td>dm_backup_file_directory_for_database</td><td>The backup file directory for database</td><td>$SYBASE/<SID>/backup</td></tr>
</table>

See Customizing File Directories.
</td></tr>
<tr><td>SAP ASE prints out the bind variables in error logs.</td><td>SAP ASE now prints out the values of bind variables in error logs when deadlock happens, instead of just printing out bind variables for SQL statements and participating specified process IDs that cause the deadlock. This enhancement makes it easier to identify objects involved in a deadlock.</td></tr>
<tr><td>The sp_hadr_admin system procedure adds these parameters:
• sp_hadr_admin … add_application_interface
• sp_hadr_admin … drop_application_interface
• sp_hadr_admin … list_application_interface</td><td>Allows you to add, drop, and list applications interfaces to an HADR cluster. Applications use this interface to connect to SAP ASE.

See sp_hadr_admin Syntax.</td></tr>
<tr><td>setuphadr command uses openssl utility provided with SQL Anywhere.</td><td>The setuphadr command now uses the openssl utility provided with SQL Anywhere for certificate management.</td></tr>
</table>

Feature	Description
The `sp_addtype` system stored procedure is enhanced to reuse the IDs of the deleted datatypes when creating new user-defined datatypes.	When creating user-defined datatypes, each newly created datatype is associated with a unique datatype ID. The maximum number of the datatype ID is 32768. Before this enhancement, when ID 32768 was reached, no new user-defined datatype could be created. However, with the enhancement, when creating new user-defined datatypes, the IDs of deleted datatypes can be reused for newly created user-defined datatypes. This means, even if ID 32768 is already created, you can still create new datatypes as long as you do not exceed the maximum number of datatype IDs.

7 Version Update - 16.0 SP03 PL02

SAP ASE version 16.0 SP03 PL02 introduces a number of enhancements.

For additional information, see the readme files in the Release Notes Information https://launchpad.support.sap.com/#/notes/2550339.

Table 12: 16.0 SP03 PL02 Enhancements

Feature	Description
SAP ASE SQLScript adds the `LOCATE` function.	The `LOCATE` function returns the position of a substring within a string. See LOCATE Function (String)
Resolve task hanging on `kctCheckAsync()` function.	When running in threaded kernel mode, a task may hang on the function `kctCheckAsync()`. To resolve this issue, the `async poll timeout` configuration parameter has been introduced. The task can be killed using the configuration parameter, or using the kill command when the trace flag 7866 is enabled. See async poll timeout.
Allow `select` queries to access `sysimrslogs`	You can run `select` queries against the `sysimrslogs` system table. Previous versions did not allow you to view this table. See sysimrslogs.
Allow `quiesce database`, `mount`, and `unmount` commands on mirrored databases	You can issue a `quiesce database`, `mount`, and `unmount` commands with external dumps on servers that have mirroring enabled. See quiesce database, mount, and unmount.
Adds parameters to `sp_helpconfig` for the `HCB index memory pool size` configuration parameter	Adds additional arguments for the `estimate` parameter to run against the `HCB index memory pool size`, allowing it to override default values. See sp_helpconfig and HCB index memory pool size.
The `monSpinlockActivity` monitoring table including information about the buffer-level spinlock contention.	The output is reported as `bufspin <cache_name>`. See Buffer-Level Spinlock Contention (bufspin cache_name).
`dbcc` adds the `hotbufs` parameter.	Issuing `dbcc hotbufs` lists pages in the cache with high contention. See Buffer-Level Spinlock Contention (bufspin cache_name) and dbcc.

8 Version Update - 16.0 SP03 PL01

SAP ASE version 16.0 SP03 PL01 introduces enhancements to existing procedures.

For additional information, see the readme files in the Release Notes Information https://launchpad.support.sap.com/#/notes/2512467.

Table 13: 16.0 SP03 PL01 Enhancements

Feature	Description
`sp_imrslog_thresholdaction`	The `sp_imrslog_thresholdaction` procedure creates a threshold for the minimum number of free pages in a database. See sp_imrslog_thresholdaction

9 Version 16.0 SP03

SAP ASE version 16.0 SP03 builds on the speed and scalability themes of ASE 16.0 SP02 by extending its OLTP performance capabilities.

The key features for this release are:

- In-Memory XOLTP:
  - In-memory row storage (IMRS) stores transactionally active data, providing a high-performance storage repository for hot data (hot tables or hot partitions).
  - Multi-Version Concurrency Control (MVCC) eliminates read and write conflicts.
  - Hash Cached BTree (HCB) allows caching of frequently accessed unique index keys for faster access.
- SAP HANA Integration:
  - The initial feature release of SAP HANA SQL schema and SQLscript support allows developers to develop in a common SAP SQL dialect and structure across SAP DBMS platforms.
  - SAP HANA-style session context provides support for SAP ASE and SAP HANA compatibility.
- Security Enhancements:
  - Common Crypto Library (CCL) provides a standard and secure SSL implementation.
  - On Demand Encryption allows encryption of sensitive commands on demand, as well as the automatic encryption of passwords.
  - Granular Auditing provides support for greater flexibility in limiting which logins or sessions are audited – particularly for previously global level audited events.

These are the new features and enhancements that are available:

Feature	Description
In-Memory Row Storage [page 28]	The in-memory row storage (IMRS) cache stores transactionally active data, providing a storage repository for hot data and extends the table's standard, page-based storage.
SAP ASE Support for SAP HANA SQLScript [page 39]	Use the SQLScript structured query language and SAP HANA SQL extensions used by SAP HANA to build application logic using SAP HANA SQLScript syntax directly in SAP ASE.
Granular Auditing [page 43]	Audit databases at the login level or audit logins granted a particular role.
SAP Common Crypto Library Support [page 43]	SAP ASE now uses the proprietary SAP Common Crypto Library (FIPS 140-2 compliant) as the SSL provider and cryptography library.
Simplified Native Access Plans Improvements [page 37]	Simplified native access plan has been enhanced to improve the overall workload performance of the identified workload.
Enable Buffered IO for Backup Server [page 37]	Allow Backup Server to use buffered I/O on database devices during load database and load transaction operations.

Feature	Description
Query Optimization [page 36]	Track CPU, execution, compile, and elapsed time for queries.
Optimize the Return Value of a Query [page 35]	Optimize the return value of a query by removing unnecessary fields from the `select` and `select into` statements.
JAVA RUNTIME (JRE) Support [page 38]	SAP ASE now ships JAVA RUNTIME (JRE) 8.1.026.
SAP HANA Session Context [page 40]	The SAP HANA style session context is provided to support SAP ASE and SAP HANA compatibility.
intersect and except Parameters [page 45]	Use the `intersect` parameter to return the common data of two or more queries, and the `except` parameter to subtract one query result set from another.
Table-Valued Parameter [page 47]	Use the table type to define stored procedure parameters, representing tabular results, and pass multiple rows of data to a stored procedure without creating temporary tables.
Global Temporary Tables [page 47]	Global temporary tables can be used to create persistent metadata across sessions, and private data per session, with data being automatically deleted at the end of the session.
HADR New Features [page 48]	New commands and utilities are available for HADR.

For additional information, see the readme files in the Release Notes Information https://launchpad.support.sap.com/#/notes/2489727.

9.1 Performance Improvements and System Administration

New and changed features for performance and system administration.

9.1.1 In-Memory Row Storage

The in-memory row storage cache stores transactionally active data, providing a storage repository for hot data (hot tables or hot partitions) and extending the table's standard, page-based storage.

The in-memory row storage (IMRS) cache supports hot tables (where the entire table is memory-resident, and is constantly scanned or updated) and hot data (a portion of the table, usually small, is transactionally active and memory-resident). The IMRS cache is an extension of the database device space allocated to the database, providing additional in-memory storage to store frequently accessed rows from the IMRS-enabled database.

Not all tables, or even all rows in a table, reside in the IMRS cache. Instead, only a small number of hot tables, or hot rows, in an OLTP schema use the IMRS cache. The server inserts all new data into the IMRS cache, which it accesses and updates directly in-memory.

A disk-resident, fully durable database enabled for IMRS is a traditional database, supporting all the ACID properties of a regular disk-resident database, and providing full support for `load database` and `load transaction`. Although an in-memory database is entirely in-memory, only a small portion of the IMRS-enabled database is in-memory.

In-memory row storage is described in the In-Memory Database Users Guide.

The workload profiler is described in Managing the Workload with sp_wlprofiler.

System Changes

IMRS introduces these system changes.

System Change	Description
Changed commands: • `alter database ... [[,] row_caching {on | off} [ for all tables ]] ...` • `alter database ... set [[,] {row_caching | snapshot_isolation} }on | off} [for all tables]]...` • `alter index ... | [, index_hash_caching = {on [, bucket_count = <number> ] | off | default}] ...` • `alter table ... set [[,] row_caching {on | off | default}] [[,] snapshot_isolation {on | off}] row storage {off} ...` • `create database ... row storage on cache_name with row_caching {on | off} ...` • `create table... with row_caching {on | off}]...` • `disk init... type = 'inmemory' | 'imrslog' | 'nvcache'...` • `select ... into with [[,] row_caching {on | off}] [[,] snapshot_isolation {on | off}]`	Commands add syntax and support for: • Creating tables and databases for IMRS • Altering table and databases to use the IMRS • Adding support for multiversion concurrency control (MVCC) • Managing, monitoring, and supporting the IMRS environment See Reference Manual: Commands

System Change	Description
Changed system procedures: • `sp_cacheconfig [<cachename> ... row_storage ...` • `sp_dboption` • `sp_help` • `sp_helpcache` • `sp_helpdb` • `sp_spaceusage` • `sp_spaceused`	System procedures add syntax and support for: • Configuring SAP ASE for row storage • Monitoring and managing the IMRS See Reference Manual: Procedures.
New system procedures: • `sp_imrs` • `sp_wlprofiler`	Provides support for: • Managing and monitoring an IMRS cache • Monitoring and configuring resources for the IMRS See Reference Manual: Procedures.
Altered system tables: • `sysattributes` • `sysconfigures` • `sysdatabases` • `sysdevices` • `sysobjects` • `syssegments`	The changes are: • Adds the `RC object_type`. • Adds the encrkeyid column. • Adds the value 16 to the `status2` column. • Adds the `status5` column • Reports information about the `imrslogsegment`. See System Tables.
New monitoring tables: • `monIMRSCache` • `monIMRSCacheActivity` • `monIMRSGCTasks` • `monIMRSPartitionActivity`	• Provides monitoring and usage metrics for IMRS caches, • Stores statistics relating to in-memory data caches. • Provides information about garbage collection tasks that are running and the statistics they collect, such as total memory freed large object garbage collection tasks, total disk space freed by deleting the obsolete LOB versions, and so on. • Provides information about the activity for in-memory row storage caches across individual partitions or objects. See Monitoring Tables.

System Change	Description
Altered functions: • `curunreservedpgs` • `lct_admin` • `row_count` • `object_attr` • `dol_downgrade_check` • `has_sysvers_tran`	• Returns the number of free pages for imrslog on-disk row storage devices. • Extends the output from `list_dump_fs` parameter to report the use of IMRS in future database or transaction dumps. • Reports information about IMRS. See Reference Manual: Building Blocks.
New configuration parameters: • `enable mem scale` • `number of imrs gc tasks per db` • `number of pack tasks per db` • `number of lob gc tasks per db` • `imrs cache utilization`	• Enables or disables the Mem Scale licensable option (introduced in version SP02, but required for IMRS). • Determines the default number of IMRS garbage collector tasks spawned during a server start or when an IMRS-enabled databases are brought online. • Determines the number of imrs_pack threads per database • Determines the number of LOB garbage collector tasks spawned during a server start or when an IMRS-enabled databases that contain LOB data are brought online. • Specifies a threshold, as a percentage, for IMRS cache utilization, above which the server does not allow `insert`, `update`, and `select` commands to store rows in the row storage. See Reference Manual: Configuration Parameters.

9.1.1.1 Multiversion Concurrency Control

Multiversion concurrency control (MVCC) allows the server to lock rows for writing in one session while granting access to unaltered versions of these rows in another session. That is, reads never wait for writes, and writes never wait for reads.

SAP ASE offers multiversion concurrency control that uses in-memory row versioning, where all inserts, updates, and deletes are performed in the in-memory row storage or on-disk multiversion concurrency control (on-disk MVCC), which stores row versions in a temporary database specified when the database is created or altered. Multiversion concurrency control works the same regardless of the method you use.

See Multiversion Concurrency Control.

System Changes

MVCC introduces these system changes.

System Change	Description
Changed commands: • `alter database... [version storage {on <tempdb_name> | off]... [[,] snapshot_isolation { on [ using version storage ] | off } [ for all tables ]]` • `alter table [, snapshot_isolation { on [ using version storage ] | off }| off }]` • `create database... [ version storage on <temporary_database>] ... [ [,] snapshot_isolation { on [{ row | version} storage ] | off } ]` • `create table... row_caching {on | off}]... snapshot_isolation {on [using version storage] | off}` • `disk init ... [, type = 'inmemory' | 'imrslog' | 'nvcache']...`	• Specifies that you are altering an existing database for MVCC or to use snapshot isolation level. • Specifies that you are altering existing tables for MVCC or to use snapshot isolation level • Specifies that you are creating a database for MVCC or with snapshot isolation level • Specifies that you are creating a table that uses MVCC or uses snapshot isolation level • Specifies that you are creating a device for the `imrslog` • Configures row caching for the current session. See Reference Manual: Commands
Changed system procedures: • `sp_cacheconfig [<cachename> ... [row_storage]`	• Creates a cache of type `row_storage` that can be used for versioning to support MVCC See Reference Manual: Procedures.

System Change	Description
New configuration parameters: • `number of imrs gc tasks per db` • `number of lob gc tasks per db` • `number of reexecutions`	• Determines the default number of IMRS garbage collector tasks spawned during a server start or when an IMRS-enabled databases are brought online. • Determines the number of LOB garbage collector tasks spawned during a server start or when an IMRS-enabled databases that contain LOB data are brought online. • Specifies the maximum number of internal reexecutions of the DMLs the query processor can attempt when it receives write conflicts. See Reference Manual: Configuration Parameters.
New option for the `pssinfo` function.	The option `has_sysversions_tran` returns 1 if the process includes any on-disk MVCC transactions for a given database. See pssinfo

9.1.1.2 Hash Cached BTree Index Support

Hash cached B-tree indexes consist of a a standard disk-based, B-tree index and an additional in-memory lock-free hash table.

Initially, the server creates only the B-tree index for all data rows, and the hash table is built at runtime according to the current workload. The server indexes data rows accessed by point queries using the hash table. After it is cached, the hash table provides the data row ID directly, reducing the path to the RID for the qualified data row from the traditional B-tree index and improving performance. Applications with numerous repeated point queries benefit from using hash-cached B-trees.

The in-memory hash table is used only for hot rows (that is, portions of a table, usually small, that are transactionally active and memory resident) from in-memory row storage. Changes to the in-memory hash table are not logged, and locks are not taken when it is changed, which help scale the performance, especially in highly- concurrent DML environments.

Hash-Cache BTree Indexes are described in Performance and Tuning Series: Locking and Concurrency Control.

Hash cached BTree indexes introduce these system changes.

System Changes

Hash cached BTree indexes introduce these system changes.

System Change	Description

Changed commands:
`create index ...index_hash_caching = {on [, bucket_count = <number> ] | off | default} max_rows_per_page = <num_rows>,`
`alter index ...[, index_hash_caching = {on [, bucket_count = <number> ] | off | default}]`
	• Creates hash cache BTree indexes
	• Alters existing indexes to use hash cached BTree.
	See Reference Manual: Commands

New configuration parameters: • `enable HCB index` • `HCB index auto tuning` • `HCB index memory pool size` • `HCB auto tuning interval` • `point query rate threshold` • `hash table hit rate threshold` • `number of hcb gc tasks per db`	Enables or disables SAP ASE for hash-cached BTree indexes. See Reference Manual: Configuration Parameters.

9.1.2 Optimize the Return Value of a Query

Use the `select...limit` and `offset` parameters to limit the number or rows in your query results, or omit a specified number of rows before the beginning of the result set.

You can use `limit` and `offset` on these statement types:

- `select` statements
- Deviced tables
- Union
- Non-correlated subqueries

System Changes

This feature introduces these system changes.

System Change	Description
`select limit/offset`	This feature adds a `limit` option which allows you to limit the number of rows returned from a query, and an `offset` option which allows you to omit a specified number of rows before the beginning of the result set. Using both `limit` and `offset` skips both rows as well as limit the rows returned. For more information, see select and Pagination Queries Using limit and offset .

9.1.3 Query Optimization

Track CPU, execution, compile, and elapse time for queries.

New global variables and the `QueryOptimizationTime` monitoring table column, help diagnostics, or tracking query processing.

System Changes

This feature introduces these system changes.

System Change	Description
• `@@compiletime` • `@@elapsedtime` • `@@cputime` • `@@runtime`	Global variables that show the: • Compilation time of the previously compiled query. • Elapsed time taken by the current run of the query. • CPU time taken by the current run of the query. • Time taken to execute the query (in `exec_lava`). See Global Variables.
`QueryOptimizationTime`	These monitoring tables have been extended to include the `QueryOptimizationTime` column: • `monProcessActivity` • `monProcessStatement` • `monSysStatement` See Monitoring Tables

Related Information

Query Optimization Diagnostics and Tracking

9.1.4 Simplified Native Access Plans Improvements

Simplified native access plans introduced in 16.0 SP02 is query plan execution feature that allows for faster execution of extreme online transaction processing (XOLTP) queries. The simplified native access plan feature has been enhanced for SAP ASE 16.0 SP03 to improve the overall workload performance of the identified workload.

The simplified native access plans features requires the ASE_MEMSCALE license, and is enabled using:

```
sp_configure "streamlined dynamic SQL", 1
```

These are the enhancements for 16.0 SP03:

- Procedure engine – simplified native access plans support usage is extended for the SAP ASE procedure engine. This extension allows more efficient security and rvm check calls.
- SNAP expression engine improvement – Native expression evaluation generation is improved and extended.
- Tighter integration with access layer – Generated callbacks for access layer reduce round-trip cost between query processing and the access layer.
- Asynchronous Code generation and JIT compilation – SNAP code generation and JIT compilation have been improved to be more efficient and scalable.

9.1.5 Enable Buffered IO for Backup Server

The new configuration option `enable buffered io for load` allows Backup Server to use buffered I/O on database devices during `load database` and `load transaction`, regardless of the database devices configuration settings in SAP ASE.

By default, Backup Server uses the same mechanisms as SAP ASE when opening a device. If a database device uses synchronous I/O or direct I/O, Backup Server uses these same settings. When you set the configuration option `enable buffered io for load`, SAP ASE requests Backup Server to open the database devices using buffered I/O during the load process, independently of the devices being configured to use DSYNC or DIRECTIO. Once the `load` command completes, all the in-memory buffers are flushed to disk. This option results in a heavy usage of the operating system file cache, and, in some cases, can speed up load operations.

System Changes

This feature introduces these system changes.

System Change	Description
`enable buffered io for load`	Enables Backup Server to use buffered I/O on database devices during `load database` and `load transaction`. See Reference Manual: Configuration Parameters.

9.2 Installation

New, changed, and removed features related to installation and update.

9.2.1 JAVA RUNTIME (JRE) Support

SAP ASE now ships JAVA RUNTIME (JRE) 8.1.026.

The following JRE versions are shipped with SAP ASE.

Table 14: JRE Support

Product/Feature	Version	Location
PCI/JVM	SAP JRE 8.1.026	`$SYBASE/shared/ase/`
Cockpit/Utilities	SAP JRE 8.1.026 and SAP JRE 7.1.048	`$SYBASE/shared/`
Installer/uninstaller	SAP JRE 8.1.026	`$SYBASE/jre64/`

9.2.2 Configurable Data Cache Size Increase

The configurable cache size is increased from 2TB to 4TB (4TB is the total memory configurable in SAP ASE).

SAP recommends that you create a cache greater than 2TB during boot time to avoid fragmented memory and reduce the cache creation time.

To support a maximum cache size of 4TB, the unit value of `sysconfigures.value2` is increased from 1KB to 2KB. The `value` column retains its `int` datatype with this change.

However, with this change, SAP ASE represents cache size with half of the number used earlier, allowing SAP ASE to create maximum data cache size up to 4 TB instead of earlier maximum of 2 TB.

For example, earlier cache size of 1000 MB was represented as 102400 of 1KB size, now it is represented as 51200 of 2KB size.

For example, creating a cache of 100MB in earlier versions required 102400 pages. For versions 16.0 SP02 PL06 and later, a cache of 100MB requires 51200 pages of 2K size.

This change also affects downgrade because earlier versions expect a cache that is based on 1KB pages instead of the new 2K pages, and because earlier versions cannot manage a cache size of greater than 2TB.

Downgrade considerations:

- If you want to downgrade to a version earlier than 16.0 SP02 PL06, you need to configure the cache size to less than 2TB in the current version, and then downgrade.
- If the cache size is already lower than 2TB in current version:
  - After you downgrade, run `sp_cacheconfig` with the size value of the current version. This allows the run value and configuration value to be displayed as the same value. If you do not run `sp_cacheconfig`, the functionality is not affected, but the displayed run and configuration values for the cache size will be different.
  - If there are large pools configured, after you downgrade, run `sp_poolconfig` with the size value of the current version for all the large pools. This allows run value and configuration values to be the same.

> **i Note**
>
> This feature was introduced in 16.0 SP02 PL06 on Linux only. In 16.0 SP03, this features is available on all platforms.

9.3 SAP HANA Integration

New features related to SAP HANA and SAP ASE integration.

9.3.1 SAP ASE Support for SAP HANA SQLScript

SAP supports its own scripting language; SQLScript in stored procedure and SQL extensions based on the ANSI SQL standards.

SAP ASE supports the SQLScript structured query language and SAP HANA SQL extensions used by SAP HANA by adding a SQLScript parser, allowing you to build application logic using SAP HANA SQLScript syntax directly in SAP ASE.

To use SAP HANA SQLScript and SQL extensions, create a unique database in SAP ASE that supports a parser specifically designed for SQLScript. You can then develop SQL scripts, or modify existing Transact-SQL scripts using SQLScript extensions. You can then execute these scripts in the SAP ASE SQLScript database, or on any SAP HANA database.

For more information and a list of supported SQLScript extensions in SAP ASE, see the SAP ASE SQLScript Reference.

9.3.2 SAP HANA Session Context

SAP HANA-style session context is provided to support SAP ASE and SAP HANA compatibility.

Applications on a database server can limit access to the data. For example, you can code an application so that it knows which users are allowed to update specific data. The attributes that enable this coding comprise an application context.

System Changes

This feature introduces these system changes:

System Change	Description

`set`	Use the `set` command to define a session variable. For example:
	`set 'abc' = '123'`
	Use the `set` command to rewrite a session variable. For example:
	`set 'abc' = '456'`

`unset`	Use `unset` to delete a session variable. For example: `unset 'abc'`
`session_context`	Use the `session_context` function to retrieve a variable. For example: `select session_context('abc')` See Reference Manual: Building Blocks

9.4 Security

New and changed security features.

9.4.1 Encryption

SAP ASE now supports the encryption of specific sensitive commands on demand, as well as the automatic encryption of passwords and commands when using the `bcp` utility.

9.4.1.1 On-Demand Encryption

On-demand encryption allows you to issue encrypted commands and system procedures as needed, both for a full `isql` session, as well as on an ad hoc basis when you want to run individual commands that contain sensitive data.

When you use on-demand encryption, the client encrypts the command before sending it to the server. The server decrypts the command before processing it. On-demand encryption is only available as long as both the client and server is running SAP ASE versions 16.0 SP03 or later; no special permissions or additional licenses are necessary.

System Changes

This feature introduces these system changes.

System Change	Description
`go encrypt`	When you add `encrypt` to the `go` command terminator, SAP ASE encrypts just the preceding command. Statements are encrypted individually, so subsequent commands, if sent without `encrypt`, are sent unencrypted.
`isql --command_encryption`	Encrypts every command sent during an `isql` session. See isql
`isql --conceal`	Hides and encrypts input during an `isql` session, so that sensitive information such as passwords and PIN codes are concealed. See isql
`@@prev_batch_encrypted`	Checks the encryption status of the most recently executed batch of commands in the current user session. A return of 0 means the server received the command as plain text; a return of 1 means the commands were received encrypted. See Global Variables
`set`	`set remote command encryption {on|off}` Enables or disables encryption during the session for commands and system procedures sent when using CIS. Only the servers need to be capable of on-demand encryption. This applies to the following: • `connect to <server_name>` • `execute <server_name>...<proc_name>` • `sp_autoconnect` • `sp_passthru` • `sp_remotesql` See set

Related Information

Encryption

9.4.1.2 Bulk Copy Encryption

As of SAP ASE 16.0 SP03 and later, the `bcp` utility automatically encrypts the passwords and commands specified by the `initstring`, `colpasswd`, and `keypasswd` arguments.

System Changes

This feature introduces these system changes.

System Change	Description
`bcp`	The following arguments for the `bcp` utility are automatically encrypted: • `--initstring "<TSQL_command>"` • `--colpasswd [[[<db_name>.[<owner>].]<table_name>.]<column_name>[<password>]]` • `--keypasswd [[<db_name>.[<owner>].]<key_name>[<password>]]` **Note** There is no command line option to turn encryption on or off for these commands. See bcp, and Encryption

Related Information

Encryption

9.4.2 SAP Common Crypto Library Support

SAP ASE uses the proprietary SAP Common Crypto Library (FIPS 140-2 compliant) as the SSL provider and cryptography library. OpenSSL is no longer supported for encryption and decryption.

System Changes

This feature introduces these system changes.

System Change	Description
`certauth, certpk12, certreg`	These utilities are deprecated as SAP ASE no longer supports OpenSSL for SSL certificates. As a replacement, SAP ASE now includes SAP Common Crypto Library for cryptographic services, such as encryption and decryption.
`ExportPrivateKey`	You can use the `ExportPrivateKey` utility to encrypt a server's private key in Public-Key Cryptography Standards (PKCS) 8 format. Bundled with the SAP ASE installer, you can use the `ExportPrivateKey` utility with keytool (a key and certificate management tool) to generate an SSL certificate. You can use the utility to create SSL certificates on both Windows and Linux platforms with Java Runtime Environment (JRE) version 7 or later. See ExportPrivateKey

9.4.3 Granular Auditing

Auditing has been enhanced to allow auditing at the login level, or logins granted a particular role, which reduces the number of audit records generated.

Auditing is supported at the login and role level (both system roles and user-defined roles) for:

- All global options.
- The new global audit option `all`. You can use `all` to audit all actions for a specified login name and role name.

Audit Log in Attempts and All Actions at the Login and Role Level

Audit the attempts to log on to the server for particular login or role. For example:

```
sp_audit "login", "admin_role", "all", "on"
```

Auditing at the role level for both system roles and user-defined roles enables the logins granted the particular role to be audited. The content of the audit record is not changed; just the number of audit records generated, which reduces the number of audit records produced.

In this example, the `all` audit option is used to audit all actions for logins who have been granted the role `db_owner`:

```
sp_audit "all", "db_owner", "all", "on"
```

Audit at the Login or Role Level for Global Options

All global audit options support auditing at login and role level (both system roles and user-defined roles). For example:

```
sp_audit "config_history", "doc_role", "all", "on"
```

Display the Audit Status

Use `sp_displayaudit` to display the status at the login or role level. For example:

```
sp_displayaudit "login", "lyang"

Subject Name     Audit Option     Value     Type of Subject
-------------- ---------------- --------- -----------------
lyang                login        on           login
(1 row affected) (return status = 0)
```

```
sp_displayaudit "login", "db_owner"

Subject Name     Audit Option     Value     Type of Subject
-------------- ---------------- --------- -----------------
db_owner               login          on       role
(1 row affected)
(return status = 0)
```

System Changes

This feature introduces these system changes.

System Change	Description
`sp_auditing`	All global options have been enhanced to support auditing at login and role level (both system roles and user-defined roles). The `login` option has been enhanced to generate audit records of log in attempts for a particular login or logins granted system or user-defined roles. The `all` option has been enhanced to generate audit actions for all actions for specified login name and role name. In versions before 16.0 SP03, the `all` option only supported system roles. See sp_audit.
`sp_displayaudit`	`sp_displayaudit` has been enhanced to support the status of auditing at the login and role level (both system roles and user-defined roles). See sp_displayaudit.
`sysauditoptions`	The columns `id` and `status` have been added to `sysauditoptions` table to indicate audit status. See sysauditoptions.

9.5 Transact-SQL Language

New and changed Transact-SQL statements.

9.5.1 intersect and except Parameters

Use the `intersect` parameter to return the common data of two or more queries, and the `except` parameters to subtract one query result set from another.

Use the `intersect` parameter to return a single result set of a `select` statement for common data of two or more queries. Duplicate rows are eliminated from the result set, and only the rows selected by all queries or data sets are returned. A record that exists in one query and not in the other is omitted from the `intersect` results.

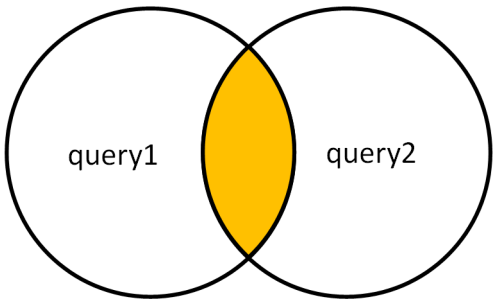

Use the `except` parameter to return a result set based on two select statements, in which the rows from the first select statement that are not available in second select statement, are returned. All rows from the second query are subtracted from the first query.

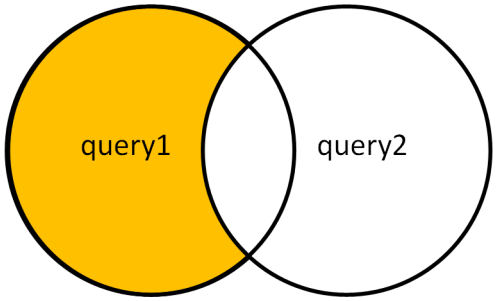

You may use `except` and `minus` interchangeably; `except` is used throughout the documentation for simplicity only.

System Changes

This feature introduces these system changes.

System Change	Description
`intersect` `except`	The parameter, `intersect` returns a single result set of a `select` statement for common data of two or more queries. The parameter `except` returns rows from the first select statement that are not returned by a second select statement. See select, and Using intersect and except.

9.5.2 Table-Valued Parameter

A table type is an extension datatype used to define stored procedure parameters representing tabular results. It allows you to pass multiple rows of data to a stored procedure without creating temporary tables.

System Changes

This feature introduces these system changes:

System Change	Description
`create type`	Creates a user-defined table type. See create type
`drop type`	Drops a user-defined table type. See drop type

Related Information

Table Type Extension Datatype

9.5.3 Global Temporary Tables

A global temporary table contains persistent metadata across sessions, and private data per session, with data being automatically deleted at the end of the session, even if the session ends abnormally.

Global temporary tables are globally available to any connection once created. This means the metadata of the table is shared across sessions. The data in a global temporary table is session-specific, meaning data inserted by a session can only be accessed by that session. The table is dropped when the last connection using it is closed.

System Changes

This feature introduces these system changes.

System Change	Description
`create global temporary table`	The `create global temporary table` statement allows you to create tables that are globally available to any connection. See create global temporary table.

9.6 HADR New Features

SAP ASE version 16.0 SP03 introduces the following new features for HADR.

SAP Replication Server

SAP Replication Server now supports the replication of stored procedures and SQL statements into or out of an HADR system. See Replicating Data In and Out of an Existing HADR System.

System Changes

This table describes SAP Replication Server system changes:

System Change	Description
`sysadmin sqm_purge_queue`	Addition of the `[, check_only]` lets you purge queues without necessarily hibernating on the Replication Server. See `sqm_purge_queue` in the *SAP Replication Server Reference Manual*.

This table describes RMA system changes:

System Change	Description
`sap_collect_log`	Command to collect the ASE/RS/RMA logs of different hosts in an HADR system to one local directory. See sap_collect_log.
`sap_suspend_component` and `sap_resume_component`	Commands to suspend or resume components on a replication path. See sap_suspend_component and sap_resume_component.

9.7 System Administrative Changes

Version 16.0 SP03 adds a number of system administrative changes.

System Changes

System Change	Description
`sysconfigures.value`	To support a maximum data cache size of 4TB, the unit value of `sysconfigures.value2` is increased from 1KB to 2KB. The `value` column retains its `int` datatype with this change. With this change, SAP ASE requires half the number of pages of size 2K to create a cache of the same size, allowing SAP ASE to create a maximum cache size of 4TB instead of the earlier maximum of 2TB. For example, creating a cache of 100MB in earlier versions required 102400 pages. For versions 16.0 SP02 PL06 and later, a cache of 100MB requires 51200 pages of 2K size. Procedures that you created on earlier versions of SAP ASE that directly query `sysconfigures.value` must multiply this value by 2 to get the actual cache size.
`sp_add_resource_limit . . . idle_time`	Limits the amount of time a process remains idle, without impacting other long-running processes that are actively performing work. See sp_add_resource_limit.
`CAP_IPC_LOCK`	Enabling lock shared memory on Red Hat Linux Enterprise configures the server to lock shared memory to physical memory in 2k pages, and enables the server to use granular privileges to lock shared memory with the `CAP_IPC_LOCK` capability setting. Locking Shared Memory.
`syslogsTruncpointsProcessed`	The `monRepTruncationPoint` monitoring table adds this column. If in-memory row storage is enabled, displays the number of times a truncation point was moved in `syslogs`. If IMRS is not enabled, has the same value as `TruncpointsProcessed`. monRepTruncationPoint.
`sysimrslogsTruncpointsProcessed`	The `monRepTruncationPoint` monitoring table adds this column. If in-memory row storage is enabled, displays the number of times a truncation point was moved in `sysimrslogs`. If IMRS is not enabled, displays 0. monRepTruncationPoint.

System Change	Description
`sp_chgattribute <view_name>`	Enables or disables forced view materialization. sp_chgattribute.

9.8 Deprecated Features

Features of SAP ASE that are deprecated.

Table 15: SAP ASE Deprecated Features

Deprecated Feature	End of Life	Description
OpenSSL	16.0 SP03	Common Crypto Library replaces OpenSSL (for ASE) and CSI (for RS).
Developer Edition	16.0 SP03	The SAP ASE Developer Edition is no longer supported. As an alternative, you can download and evaluate a non-production environment edition of SAP ASE Enterprise Edition. For more information, see Evaluate SAP Adaptive Server Enterprise.
Backup Server 32-bit	January 2018	If you are planning to develop new applications with SAP ASE, avoid using Backup Server 32-bit. Backup Server 64-bit is available as a migration path.
TLI Interface	January 2018	There are two interfaces file entry formats, TLI and TCP. If you are planning to develop new applications with SAP ASE, avoid using the TLI interface.
Kernel Process Mode	January 2019	Customers planning to develop new applications with SAP ASE should avoid using this feature. Before the final support date, this feature will continue to work and will be supported, but there will be no enhancements. After the final support date, this feature may not be explicitly removed, but may not work and there will be no support.

10 Version Update - 16.0 SP02 PL06

SAP ASE version 16.0 SP02 PL06 introduces enhancements to existing features and options.

For additional information, see the readme files in the Release Notes Information https://launchpad.support.sap.com/#/notes/2479721 and https://launchpad.support.sap.com/#/notes/2504483.

Table 16: 16.0 SP02 PL06 Enhancements

Feature	Description

HADR Enhancements	New for HADR:
	• SSL security is enabled for the following HADR scenarios:
	  ○ Inside an HADR system. See Enabling SSL for the HADR System.
	  ○ For an external replication system. See Configuring SSL for External Replication.
	  ○ For the Fault Manager. See Configuring the Fault Manager in an SSL-Enabled HADR Environment.
	i Note
	SSL is supported on platforms Linux x64 on x86-64, and Windows x64.
	• Support for the Fault Manager for Business Suite. See Using the Fault Manager with Business Suite
	• Two profile parameters for Fault Manager setup:
	  ○ `ha/syb/bootstrap_retries`.
	  ○ `ha/syb/trace_file_size`.
	See Parameters for Fault Manager Setup
	• Command to display detailed information about the current version of the Component Interface (CI) in use by the SAP Replication Server.
	See admin version, ci.

Enable or Disable RVM (Reference Validation Mechanism) Checks	The new configuration parameter, `enable lightweight rvm` enables and disables reference validation mechanism checks to determine if there are no permission changes between executions of the same compiled plans, streamlining the query execution path. For more information, see enable lightweight rvm.

Feature	Description

Support for Cache Memory up to 4TB	**i Note**
	This feature is available in 16.0 SP02 PL06, but has been removed for 16.0 SP02 PL06 HF1.
	
	(Linux Only) The configurable cache size is increased from 2TB to 4TB (4TB is the total memory configurable in SAP ASE). SAP recommends that you create a cache greater than 2TB during boot time to avoid fragmented memory and reduce the cache creation time.
	
	To support a maximum cache size of 4TB, the unit value of `sysconfigures.value` is increased from 1KB to 2KB. The `value` column retains its `int` datatype with this change.
	
	With this change, SAP ASE requires half the number of pages of size 2BK to create a cache of the same size, allowing SAP ASE to create a maximum cache size of 4TB instead of the earlier maximum of 2TB. For example, creating a cache of 100MB in earlier versions required 102400 pages. For versions 16.0 SP02 PL06 and later, a cache of 100MB requires 51200 pages of 2KB size.
	
	Procedures that you created on earlier versions of SAP ASE that directly query `sysconfigures.value` must multiply this value by 2 to get the actual cache size.
	
	The behavior of existing stored procedures such as `sp_cacheconfig` and `sp_helptext` remain unchanged, as they implement the change of unit size from 1KB to 2 KB.
	
	This change also affects downgrade because earlier versions expect a cache that is based on 1KB pages instead of the new 2K pages, and because earlier versions cannot manage a cache size of greater than 2TB.
	
	Downgrade considerations:
	• If you want to downgrade to a version earlier than 16.0 SP02 PL06, you need to configure the cache size to less than 2TB in the current version, and then downgrade.
	• If the cache size is already lower than 2TB in current version:
	  ○ After you downgrade, run `sp_cacheconfig` with the size value of the current version. This allows the run value and configuration value to be displayed as the same value. If you do not run `sp_cacheconfig`, the functionality is not affected, but the displayed run and configuration values for the cache size will be different.
	  ○ If there are large pools configured, after you downgrade, run `sp_poolconfig` with the size value of the current version for all the large pools. This allows run value and configuration values to be the same.

Enable Implicit Conversion Between Character Types	The `extend implicit conversion` configuration parameter allows implicit conversion between character types and arithmetic types. The supported character types include `char` and `varchar` and arithmetic types include `smallint`, `int`, `bigint`, `numeric` and `decimal`. See extend implicit conversion

11 Version Update - 16.0 SP02 PL05

SAP ASE version 16.0 SP02 PL05 introduces enhancements to existing features and options.

For additional information, see the readme files in the Release Notes Information https://launchpad.support.sap.com/#/notes/2405541.

Table 17: 16.0 SP02 PL05 Enhancements

Feature	Description

HADR Support for External Replication	External replication allows you to replicate data from an HADR system to an external Replication Server, or from an external Replication Server to the HADR system.
	
	External replication includes changes to these commands and system tables:
	
	• New RMA Commands
	`sap_disable_external_replication`,
	`sap_disable_external_replication`,
	`sap_enable_external_replication` , `sap_status spq_agent`,
	`sap_tune_rat`, `sap_failover_drain_to_er`
	• Changed RMA Commands
	`sap_disable_replication`, `sap_enable_replication`,
	`sap_materialize`, `sap_status resource` , `sap_tune_rs`
	• Changed RMA Tables
	`hadrStatusResource`, `monRepSyncTaskStatistics`
	• New Replication Server Commands
	`configure spq_agent enable` , `configure spq_agent disable`,
	`resume spq_agent` , `suspend spq_agent`, `activate all spq_agent`,
	`activate all spq_agent`, `deactivate all spq_agent`, `create connection`, `drop connection`, `alter connection`
	• Changed Replication Server System Tables
	`rs_databases.attributes`, `rs_databases.src_status`,
	`rs_spqratgroup`, `rs_spqratstate`, `rs_spqratconfig`
	
	For more information, see Replicating Data In and Out of an Existing HADR System .

Feature	Description
SySAM License Changes	If you are using a served license model, you must install the SySAM license server version 2.4. SAP ASE 16.0 SP02 PL05 and SySAM 2.4 use FlexNet Publisher version 11.14. You are allowed a grace period of 90 days when the requisite number of licenses is not met. SAP ASE issues warnings when the requisite number of licenses are not available and provides the expiration duration time, up to the maximum of 90 days. Previously, the maximum was 30 days. SySAM 2.4 is not supported on PowerLinux. A separate SySAM license server and utilities installer is provided which allows you to install SySAM 2.4 on machine or operating system other than that used by the product. If you are installing a SySAM license server on UNIX into a directory that contains an existing SySAM license server, run: `SYSAM-2_0/bin/sysam configure none` For more information, see Installing the SySAM License Server and Platform Support.
Optimize the Return Value of a Query	You can optimize the return value of a query by removing unnecessary fields from the `select` and `select into` statements. This feature adds a `limit` option which allows you to limit the number of rows returned from a query, and an `offset` option which allows you to omit a specified number of rows before the beginning of the result set. Using both `limit` and `offset` skips both rows as well as limit the rows returned. You can use `limit` and `offset` on these statement types: • `select` statements • Deviced tables • Union • Non-correlated subqueries For more information, see select and Pagination Queries Using limit and offset .
Restore Source Database Options	The `restore database options` configuration parameter restores database options that are set by `create database`, `alter database`, and `sp_dboption` when you load a database or transaction dump. For more information, see restore database options .
Restore Database Devices	The new `disk reinit name` parameter allows you to recreate a database device that has been removed (planned or unplanned) from the operating system. For more information, see disk reinit and Restoring Database Devices with disk reinit.

<table>
<tr><th>Feature</th><th>Description</th></tr>
<tr><td>Capacity Counters Increased</td><td>The capacity of some counters have been increased by using 64 bit integers.
The columns for the following tables have been changed from <code>int</code> to <code>bigint</code>:<table><tr><th><code>monProcessActivity</code></th><th><code>monProcessStatement</code></th><th><code>monSysStatement</code></th></tr><tr><td><code>LogicalReads</code></td><td><code>LogicalReads</code></td><td><code>LogicalReads</code></td></tr><tr><td><code>PhysicalReads</code></td><td><code>PhysicalReads</code></td><td><code>PhysicalReads</code></td></tr><tr><td><code>PagesRead</code></td><td><code>PagesModified</code></td><td><code>PagesModified</code></td></tr><tr><td><code>PhysicalWrites</code></td><td></td><td></td></tr><tr><td><code>PagesWritten</code></td><td></td><td></td></tr></table></td></tr>
<tr><td>Workload Analyzer Enhancement</td><td>The SAP ASE cockpit Workload Analyzer Replay Comparison Report dashboard has been enhanced. The enhancements to the Replay Comparison Dashboard include:
You can modify a repository connection.
Display the capture Status.
Delete a replay from both the repository database, as well as its associated PCAP file.
Analyze workload replay information such as requests from IP addresses, a histogram displaying the timing and number of requests, and a summary of additional basic replay information.
View the log level preferences.
For more information, see SAP ASE cockpit Capturing, Analyzing, and Replaying Workloads documentation.</td></tr>
<tr><td>New Function for Cursor Information</td><td>The new <code>cursorinfo</code> function returns information about a cursor (including its status, whether a cursor is open or read only, and so on), and allows you to specify the nesting level.
For syntax , see cursorinfo.</td></tr>
<tr><td>New Parameter Allows Transaction Rollback</td><td>A new configuration parameter allows a transaction rollback when a non-recoverable error is raised. The new parameter <code>allow db suspect on rollback error</code>, allows a forceful shutdown of the database and marks the database suspect if a non-recoverable error is hit during a transaction rollback. If enabled, this provides the ability to execute a last dump transaction, which allows an archive of all the transactions up to the point the rollback failed.
For syntax, see sp_dboption.</td></tr>
</table>

Feature	Description
CSI libraries and OpenSSL Versions	Ensure that you are using the CSI libraries and OpenSSL versions required for your security needs. Multiple versions of CSI are provided to prevent issues where multiple products are installed in the same `$SYBASE` location. Learn how to list the dependent CSI libraries using the `ldd` command and determine the OpenSSL version. Upgrade older version of installed SAP products which have an indirect dependency through OCS to older CSI libraries and recompiled and relinked with the latest OCS libraries provided with the installed version. For more information, see OpenSSL Versions.
Accelerate Database Checking	Accelerate checking on a database by skipping tables with text and image columns. If most of the tables in your database do not contain text or image columns, use the following command to skip checking on those tables and only check tables that have text and image columns: `dbcc checkdb(dbid | dbname [, "skip_tables_with_no_text" ])` For additional information about handling text and image error, see Correcting Text and Image Errors.
New Commands to Control Privileges	Two new commands allow you to control privileges for object owners. In previous releases, object owners have all privileges on the objects they own. The commands `allow` and `deny` allow you to control object privileges: • `allow` removes privilege restrictions on all tables for all owners. • `deny` removes privileges from object owners on all tables. Use `sp_displayaudit` to displays the status of audit process for `allow` and `deny`. Use `sp_helprotect` to display `deny` permissions. For more information, see allow, deny, sp_displayaudit, and sp_helprotect.

Feature	Description
Enhancements for Exception Handling	Enhancements for exception handling (secondary failure in the signal handler) and stack overflow checks in signal handler. The following enhancements have been introduced for all UNIX platforms: • Check stack overflow in the signal handler: Two new configuration parameters (`sigstack min size` and `sigstack csmd min size`) are introduced to configure the minimum amount of stack that is required to handle an address violation signal without overflowing the stack. By default, the check to prevent stack overflow in signal handler is enabled. Use trace flag 2913 to disable the stack overflow in the signal handler. The `sigstack csmd min size` parameter is used to capture a CSMD during signal handling; however, in this case more stack is required. If the amount of available stack upon entering signal handler is less than the configured value, the task becomes inoperative. If required, you can increase or decrease the stack size. • Handle a secondary failure in the signal handler: Enhancements have been introduced in the signal handler to handle a secondary exception. Use trace flag 2912 to disable the feature. **i Note** It's recommended to contact SAP technical support to make changes to these values. For more information, see sigstack csmd min size and sigstack min size .
Enhancements for `wait on uncommitted insert`	The `wait on uncommitted insert` configuration parameter has been enhanced to allow control of the wait behavior for `insert`. When set to the value of 2, the `insert` command adheres to `wait on uncommitted insert`. For more information and examples, see wait on uncommitted insert .
Enhancements for `dump database` and `dump transaction`	A label is added to `dump database` and `dump transaction` to mark the dump. When you rematerialize databases using the `dump` and `load` commands in a replication system, SAP Replication Server can distinguish a labeled dump from scheduled system dumps, which may interfere with the rematerialization process. The syntax is: `label = '<label_name>'` For more information, see dump database, dump transaction, and Adding Labels to Dumps.

12 Version Update - 16.0 SP02 PL04

SAP ASE version 16.0 SP02 PL04 introduces enhancements to existing features and options.

For additional information, see the readme files in the Release Notes Information https://launchpad.support.sap.com/#/notes/2336153 and https://launchpad.support.sap.com/#/notes/2365138.

Table 18: 16.0 SP02 PL04 Enhancements

Feature	Description
Backup Server Archive API made public	Backup Server archive API routines used by Backup Server to issue I/O requests to an archive byte stream have been made public. For more information, see Reference Manual: Backup Server Archive API.
`transfer table ... to` will now write output files only to a defined output directory tree rather than any directory writable by the running server.	As a security precaution, the `transfer table ... to` command is no longer permitted to write an output file to any directory that is writable by the running server. Instead, output files must be written to a specified output directory or to one of its subdirectories. For more information, see Transfer Tables to a Directory and sp_file_path.
New command options to resolve mismatch of IDs for allocation pages	The database ID is stored in the allocation pages of the database and remains unchanged if you change the dbid of the database when you mount it onto a server. To resolve a mismatch of IDs, you can scan and fix the database IDs for all the allocation pages using `dbcc checkalloc with fixdbid`. The `fixdbid` option runs a minimal check in the database and will fix any possible database ID mismatch for all allocation pages. The `allow_dbid_mismatch` option has also been added. Use this option to avoid raising a device validation error due to multiple database IDs found in the allocation pages. For more information, see dbcc and mount.
Set the TLS protocol to only use TLS 1.2	You can now set the Transport Layer Security (TLS) / Secure Sockets Layer (SSL) network protocol to only use TLS 1.2. To set the TLS version to TLS1.2, use: `sp_ssladmin setprotocol,'TLS1.2'` Also, you can use the new global variable `@@ssl_protocol`, or the procedure `sp_ssladmin lsprotocol` to identify the protocol negotiated during a session. For more information, see sp_ssladmin and Display or Set the TLS Protocol.

Feature	Description
New Web Services properties to control ciphers and TLS protocols	Two new Web Services properties have been added that allow you to disable specific cipher suites and TLS protocol to better manage security of SSL. These two properties are optional. • `com.sybase.ase.ws.producer.ssl.exciphersuites` – specifies the weak cipher suites to disable to prevent SSL attacks and browser errors. • `com.sybase.ase.ws.producer.ssl.exprotocols` – specifies the TLS protocol versions to disable to prevent SSL attacks and browser errors. For more information, see ws_properties.
New parameter allows you to delete the underlying operating system file when dropping a device	The new optional parameter `dropfile` allows you delete the underlying operating system file when dropping a device. For more information, see sp_dropdevice.
New configuration parameter for parameterized SQL text	Using `abstract plan dynamic replace` allows you to ignore the plan hint in the plan clause, and force the query to use the plan hint in `sysqueryplans`. For more information, see abstract plan dynamic replace.
New options for the `pssinfo` function	The `pssinfo` function supports these new options for `<pss_fields>`: • `retstat` – return status of the last executed stored procedure. • `client_cap_largeident` – returns 1 if the client supports large identifiers. • `trace_spid` – returns the spid the `set tracefile` command is tracing. • `tracefile_is_set` – returns a non-zero value if the spid is being traced. • `trace_fname` – returns the tracefile opened for capture. • `suid` – SUID of the spid. • `progname` – name of the client driver program, supplied at log in. • `progvers` – version of the client driver program, supplied at log in. • `tdsvers` – version of the client TDS, as an SAP ASE client. • `client_progname` – version of the currently active client program. • `client_progvers` – version of the currently active client program. • `client_tdsvers` – version of the currently active client TDS. See pssinfo.
New command option that allows you to skip the beginning of chained transaction for select queries	Use the new `set` option `no_chained_tran_for_l0_select` when `set chained` is enabled to skip the beginning of chained transaction for select queries being executed at isolation level 0. For more information, see set.

Feature	Description
Security Enhancements to the ASE PCA/JVM Module	The SAP ASE PCA/JVM module that integrates the target JavaVM/JRE no longer allows the JavaVM to load shared objects that are not listed in the `${SYBASE}/shared/ase/jvm.wl` white-list. The shared objects listed within the white-list provide support for JRE specific native methods. This list is the minimum required set of shared objects allowing the standard ASE_JAVA functional support. If additional shared objects are required to support specific application needs, the white-list `jvm.wl` file can be either removed or the required shared objects can be added as needed. The white-list is currently supported on UNIX/Linux releases only. The `${SYBASE}/${SYBASE_ASE}/lib/pca/sybasert.jar` that adds extensions to the target JRE, including the provided ASE PCA/JVM security manager, was further restricted to prevent the change of `system.properties`.

13 Version Update - 16.0 SP02 PL03

SAP ASE version 16.0 SP02 PL03 introduces enhancements to existing features.

For additional information, see the readme files in the Release Notes Information https://launchpad.support.sap.com/#/notes/2296685 and https://launchpad.support.sap.com/#/notes/2305941.

Table 19: 16.0 SP02 PL03 Enhancements

Feature	Description
Date-time Adjustment for Workload Replay	SAP ASE 16.0 SP02 PL03 introduces date and time adjustment for workload replay (workload analyzer option). You can now configure the SAP ASE cockpit so that when the replay begins, the date and time of the replay SAP ASE server are set to the time at which the capture originally started. By default, the date and time on the replay SAP ASE server is the actual date and time at which the replay occurs. You can use SAP ASE 15.7 SP137 to capture a workload executed on your SAP ASE server and then analyze and replay the workload using the 16.0 SP02 SAP ASE cockpit. This can help you discover any changes to your database design or server configuration that can improve performance in 16.0 SP02. SAP Adaptive Server Cockpit. For more information, see Reference Manual: Backup Server Archive API.

14 Version Update - 16.0 SP02 PL02

SAP ASE version 16.0 SP02 PL02 introduces new features.

For additional information, see the readme files in the Release Notes Information https://launchpad.support.sap.com/#/notes/2260570.

Table 20: 16.0 SP02 PL02 Enhancements

Feature	Description
Non-volatile caching	You can use named caches on solid-state drives to benefit from the performance benefits associated with non-volatile caching. Configuring an NV cache for a database and associating named caches with the database allows those objects to automatically map to the non-volatile cache. Once the NV cache is created and database is bound to it, the transition is smooth and no special action is required from the database administrator to use the NV cache. This NV Cache feature requires the MemScale licensed option. See *System Administration Guide: Volume 2 > Configuring Data Caches > Managing NV Cache* for details.

15 Version Update - 16.0 SP02 PL01

SAP ASE version 16.0 SP02 PL01 focuses on stability and quality. There are no new features.

For additional information, see the readme files in the Release Notes Information https://launchpad.support.sap.com/#/notes/2222850.

16 Version 16.0 SP02

SAP ASE version 16.0 SP02, built on the foundation of the SAP ASE version 16.0, expands the capabilities of extreme online transaction processing (XOLTP) by focusing on performance, availability and ease of use, as well as further integration with SAP HANA and SAP Business Suite.

SAP ASE version 16.0 SP02 offers reduced latency and improved throughput through features such as in-memory computing and intelligent data placement. Transactions and queries run faster with support for hardware-based locking, compiled queries, lockless data cache, latch free indexes, and temperature-based data management. A majority of these features are available through the SAP ASE database MemScale option.

Additionally, a new unified high-availability and disaster recovery (HA/DR) architecture enhances system availability by delivering continuity of service and zero data loss within one powerful and straightforward solution. These capabilities are delivered as part of SAP ASE database always-on option.

The SAP ASE workload analyzer option allows you to capture, analyze, and replay the production workload in a non-disruptive manner. You can capture a workload on your production system, replay it on a testing environment, and quickly analyze the impact configuration changes may have on application performance. By using real-life scenarios to determine the optimal configuration for your database, you ensure that your new configuration choices will have a positive impact on your production environment

Feature	Description
Database MemScale Option [page 66]	The Database MemScale option contains a set of features such as Latch-Free Indexes, Simplified Native Access Plan, and Transactional Memory. These features offer increased throughput and reduced latency for OLTP workloads and reporting queries.
Latch-Free Indexes [page 67]	Allows you to create indexes that do not require latches during traversals or scans, which can reduce contention.
Simplified Native Access Plans on Linux [page 68]	(Linux only) Allows the query plan to compile lava execution plans into native code, which is then invoked directly, allowing for faster execution of extreme online transaction processing (XOLTP) queries.
Lockless Data Cache [page 69]	Allows the buffer manager to perform search, keep, and unkeep operations without holding a spinlock on a cache, which can reduce contention.
Transactional Memory [page 70]	(Intel Haswell EX processor on Linux only) Provides transactional memory support to reduce spinlock contention in highly concurrent environments, and increases throughput.

Feature	Description
Database Always-On Option [page 71]	Using SAP ASE replication technology, the always-on option is the high availability solution for SAP ASE. The always-on option consists of two or more servers: one designated as the primary, on which all transaction processing takes place; the others as companions, which act as warm or hot standbys for the primary server and which use Replication Server to maintain copies of designated databases on the primary server.
Workload Analyzer Option [page 71]	Allows SAP ASE to capture client transactions and replay them in the same order they were received. You can perform what-if scenarios, analyze queries, compare response times before and after configuration changes, and so on.
Data Store Access Management [page 72]	Allows you to display and manage the physical storage of their data, offering a comprehensive way of managing the data's physical characteristics, and enabling database administrators to move data between various classes of physical storage, and to change the compression level of the stored data, in order to achieve an optimal balance between storage cost and access speed.
SAP Adaptive Server Enterprise Cockpit [page 73]	SAP ASE cockpit replaces SAP Control Center as a graphical administration tool for on-board management and monitoring of SAP ASE. SAP ASE cockpit supports the newly added features such as Workload Analyzer and DSAM.
Full-Text Auditing [page 73]	Provides support for full-text DDL auditing. The full text of a command is displayed in the extrainfo column of audit tables. Sensitive parameter values are masked.
Bucket Pool [page 73]	The bucket fragment manager allows you to tune the memory manager according to the varying needs of different server components.
Table User-Defined Functions [page 75]	You can create both inline and multi-statement table user-defined functions.
Table Variables [page 76]	A table variable, declared in a SQL batch, contains not just a single value, but an entire table (that is, zero or more rows consisting of one or more 10 columns).
Union-Derived Tables in Star Joins [page 76]	Include union-derived tables as dimension tables combined with a fact table hint.
64-Bit Backup Server [page 76]	Provides support for a 64-bit Backup Server.
Binding Backup Server to a Socket [page 76]	The new configuration parameter `bind backupserver address` allows you to bind Backup Server to a socket address.

Feature	Description
Job Scheduler Restart [page 77]	The new configuration parameter `js restart delay` sets the delay period between two Job Scheduler auto restart attempts after abnormal shutdown of Job Scheduler.
Enable Resolve as Owner [page 77]	Enables the resolve as owner functionality with the `enable resolve as owner` configuration parameter.
Limiting User Locks Per Session [page 77]	Introduces the `lock_count` limittype operator for `sp_add_resource_limit` allowing you to limit the number of logical locks used simultaneously by a user process.
Model Database Backlink Flag Set After dbcc shrinkdb_setup [page 77]	0x01 is not set in `sysdatabases.status` for newly created databases, or after running `dbcc shrinkdb_setup` on the specified database.
set_appcontext Changes [page 78]	Increases the character length and allows you to overwrite existing application contexts.

16.1 Database MemScale Option

The MemScale option offers enhanced performance for transactions and queries.

The Database MemScale option includes these features:

- Latch-Free Indexes – Latch-Free Indexes [page 67]
- Simplified Native Access Plans – Simplified Native Access Plans on Linux [page 68]
- Transactional Memory – Transactional Memory [page 70]
- In Memory Database option features – In-Memory Users Guide

When you enable an individual feature that is part of the MemScale option, the license manager verifies that you have the ASE_MEMSCALE license.

System Changes

This option introduces system changes.

System Change	Description
`enable mem scale`	In order to enable this option, obtain the license, then set the `enable mem scale` configuration parameter. See *System Administration Guide: Volume 1 > enable mem scale*.

16.2 Latch-Free Indexes

Latch-free indexes allow you to create indexes that do not require latches during traversals or scans, which can reduce contention.

Scans or traversals on latched indexes use the buffer manager to retrieve an index page, which can result in contention when opposing threads try to acquire the same latch. Latch-free indexes use in-memory mapping tables instead of buffer managers to retrieve an index page. You can use latch-free indexes only on data-only locked user tables.

Latch-Free Indexes are part of the MemScale option. See *SAP ASE Options* in the Installation Guide for your platform.

Before using this feature, set the `enable mem scale` configuration parameter. Next, set the `enable LFB index` configuration parameter.

See *Performance & Tuning Series: Locking and Concurrency Control > Indexes > Using Latch-Free Indexes*.

System Changes

This feature introduces system changes.

System Change	Description

`alter index`
`create index`
`alter table`
`create table`
	This feature introduces the command parameter `latch_free_index`, which allows you to indicates the latch policy of the index.
	See the *Reference Manual: Commands*.

`enable LFB index`	The new configuration parameter, `enable LFB index` enables or disables latch-free indexes. See *Reference Manual: Configuration Parameters > enable LFB index*.
`monLatchFreeIndex`	The new monitoring table, `monLatchFreeIndex` collects information related to latch free indexes. See *Reference Manual: Tables > Monitoring Tables > monLatchFreeIndex*.

16.3 Simplified Native Access Plans on Linux

(Linux x86-64/amd64) Simplified native access plans, a new query plan execution feature that compiles lava execution plans into native code, which is then invoked directly, allowing for faster execution of extreme online transaction processing (XOLTP) queries. This compiled queries feature is part of the MemScale option.

Simplified native access plans are transparent to the user. When enabled, SAP ASE uses simplified native access plans whenever qualified plans are available.

This feature supports:

- Statement cache plans, including those used for dynamic SQL (prepared statements).
- Plans with user stored procedures.
- Datarows locking table/index serial scan, with optional dynamic partition elimination
- `Scan` and `PtnScan` (partition acan) operators
- `Restrict` operator
- `Emit` operator, including `EmitSndOp` and `EmitNoSndOp`
- `OrScan` operator
- Nested loop join operator without worktable usage (for example, no reformatting)
- `NaryNLJ` operator
- Scalar aggregate operator
- `insert`, `update`, `delete` operator
- `union all` operator
- `ScalarOp` operator

See *Performance and Tuning Series: Query Processing and Abstract Plans > Understanding Query Processing > Lava Query Execution Engine > Simplified Native Access Plans on Linux*.

System Changes

This feature introduces these system changes.

System Change	Description
`simplified native access plan`	A new configuration parameter, `simplified native access plan` sets the parameter to greater than 0 to turn the feature on. The default is 0 (zero), which disables the feature. The configuration parameter is dynamic; you do not need to restart SAP ASE. **i Note** You must also enable the `streamlined dynamic SQL` configuration parameter to use the simplified native access plan feature: `sp_configure "streamlined dynamic SQL", 1` You can also turn on "`streamlined dynamic SQL`" by setting "`enable functionality group`" to 1. See *Reference Manual: Configuration Parameters > simplified native access plan*.
`monState`	Two new columns, `SnapsGenerated` and `SnapsExecuted`, show the number of native plans generated and executed, respectively. See *Reference Manual: Tables > Monitoring Tables > monState.*
`monCachedProcedures`	Four new columns, `SnapCodegenTime`, `SnapJITTime`, `SnapExecutionTime`, and `SnapExecutionCount` provide information about individual generation and execution of simplified native access plans. See *Reference Manual: Tables > Monitoring Tables > monCachedProcedures.*

16.4 Lockless Data Cache

A lockless data cache allows the buffer manager to perform search, keep, and unkeep operations without holding a spinlock on a cache. After a page is recorded in the memory, it is wired to the respective cache, so that any subsequent access to this buffer does not require a spinlock.

Binding an object to a lockless data cache reduces contention and improves performance in high core-count systems. The performance improvement is most significant for small tables with few pages that are accessed frequently. For optimal performance, the size of the cache should be larger than the size of the object bound to it.

Configure a lockless data cache like any other cache by using `sp_cacheconfig` with the new cache type parameter "`lockless data cache`", or by editing a server `config` file. You can dynamically update an existing cache to "`lockless data cache`" as long as its replacement policy is `relaxed`. If replacement policy is `strict` or `default`, restart the server for the change to take effect.

For more information about lockless data caches, see the *System Administration Guide: Volume 2 > Configuring Data Caches > Lockless Data Cache* and *Performance & Tuning: Basics > Memory Use and Performance > Configuring the Data Cache to Improve Performance > Reduce Spinlock Contention on a Named Data Cache*.

System Changes

This feature introduces system changes.

System Change	Description
`sp_cacheconfig`	For information about the "`lockless data cache`" parameter for `sp_cacheconfig`, see the *Reference Manual: Procedures > System Stored Procedures > sp_cacheconfig*.
`sp_sysmon`	The feature introduces new `Cache Hits` counters in `sp_sysmon` output for `Data Cache Management`: `ELC`, `With Spinlock Held`, and `W/o Spinlock Held`. For more information about these counters, see the *Performance & Tuning: Monitoring with sp_sysmon > Monitoring Performance with sp_sysmon > Data Cache Management > Cache Management by Cache > Cache Search, Hit, and Miss Information*. For updated output of `sp_sysmon` for the `Data Cache Management` categories, see the *Performance & Tuning: Monitoring with sp_sysmon > Monitoring Performance with sp_sysmon > Data Cache Management > Sample Output of Data Cache Management Categories*.

16.5 Transactional Memory

Spinlock contention within lock manager and buffer manager can impact SAP ASE throughput and response time. In high performance, high core situations, use the transactional memory feature to reduce this contention.

Transactional memory is part of the MemScale option and is currently supported on Haswell and Broadwell processors on Linux. See *SAP ASE Options* in the Installation Guide for your platform.

Enable transactional memory by first setting the `enable mem scale` configuration parameter. Then, set the `enable transactional memory` configuration parameter.

See *Performance and Tuning Series: Locking and Concurrency Control > Locking Configuration and Tuning > Locks and Performance> Reduce Lock Contention*.

System Changes

This feature introduces this system change.

System Change	Description
`enable transactional memory`	A new configuration parameter, `enable transactional memory` sets the parameter 1 to turn the feature on. The default is 0 (zero), which disables the feature. The configuration parameter is dynamic; you do not need to restart SAP ASE. See *System Administration Guide: Volume 1 > enable transactional memory*.

16.6 Database Always-On Option

The always-on option consists of two or more servers: one designated as the primary, on which all transaction processing takes place; the others as companions, which act as warm or hot standbys for the primary server and contain copies of designated databases from the primary server.

In an always-on system, servers are separate entities, and data is replicated from the primary server to the standby server. If the primary server fails, a standby server is promoted to the role of primary server either manually or automatically. Once the promotion is complete, clients can reconnect to the new primary server, and see data as on the previous primary server.

Servers can be separated geographically, which makes an always-on system capable of withstanding the loss of an entire computing facility. As a high availability system, the always-on option offers a zero data loss solution with very low down time, and includes automatic failover with synchronous replication. As a disaster recovery system, the always-on option offers near zero data loss solution with minimal downtime. Disaster recovery requires manual intervention and uses asynchronous replication. SAP recommends that a high availability system uses two sites that are within a few tens of kilometers of each other. Disaster recovery sites can be hundreds of kilometers apart.

Additionally, an always-on option system includes Replication Server, which synchronizes the databases between the primary and standby servers. SAP ASE uses the Replication Agent to communicate with Replication Server, and Replication Server uses Open Client connectivity to communicate with the standby SAP ASE.

See the *HADR Users Guide*, and *SAP Adaptive Server Enterprise Cockpit > Manage SAP ASE > Always-On (HADR) Option*.

16.7 Workload Analyzer Option

The SAP Adaptive Server Enterprise, enterprise edition, workload analyzer option is available in SAP ASE cockpit, which enables the capture, analysis, and replay of a production workload non-disruptively in a Web-based tool environment.

SAP ASE workload analyzer enables the capture and replay of a production workload and utilizes it to diagnose problems and understand and manage configuration changes proactively.

See the *Workload Analyzer Users Guide*, and *SAP Adaptive Server Enterprise Cockpit > Manage SAP ASE > Capturing, Analyzing, and Replaying Workloads*.

For a video overview of the workload analyzer option, see https://www.youtube.com/watch?v=w02qo9WHLsw .

16.8 Data Store Access Management

Data Store Access Management (DSAM) provides a mean for customers to display and manage the physical storage of their data, offering a comprehensive way of managing the data's physical characteristics.

DSAM permits database administrators to move data between various classes of physical storage, and to change the compression level of the stored data, in order to achieve an optimal balance between storage cost and access speed. You must purchase and install the Partitions option to use DSDAM.

DSAM benefits include:

- Data usage monitoring.
- Moving data to faster devices or to archive devices based on data access activity.
- Enhance performance by keeping active data on high performance storage.
- Reduce costs (enabling more effective use of storage assets).

See the *Data Store Access Management Users Guide*, and *SAP Adaptive Server Enterprise Cockpit > Manage SAP ASE > Data Store Access Management*.

System Changes

This feature introduces new syntax changes.

System Change	Description
`monDeviceSegmentIO`	A monitoring table that displays the DSAM collection information by device and segment. See *Data Store Access Management Users Guide > Default Configuration > DSAM > monDeviceSegmentIO Monitoring Table* .
`monDeviceSegmentUsage`	A monitoring table that displays pages used in a database by device and segment. See *Data Store Access Management Users Guide > Default Configuration > DSAM > monDeviceSegmentUsage Monitoring Table*.
DSAM API	DSAM API is a shared API that can be utilized by third party for partitions access information. The DSAM API is composed of views and stored procedures. See *Data Store Access Management Users Guide > DSAM API*.

16.9 SAP Adaptive Server Enterprise Cockpit

SAP Adaptive Server Enterprise cockpit (SAP ASE cockpit) is a new graphical administration tool for on-board management and monitoring of SAP ASE. SAP ASE cockpit supports SAP ASE version 16.0 SP 02.

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

See the *SAP Adaptive Server Enterprise Cockpit* documentation.

For information about installing the SAP ASE cockpit, see the *SAP Adaptive Server Enterprise Installation Guide* for your platform.

16.10 Full-Text Auditing

The full text of commands is displayed in the `extrainfo` column of audit tables. Sensitive values such as passwords are masked when parameter values are displayed.

In this example, you enable auditing for `role` option and alter the role "role1" by adding password:

```
alter role role1 add passwd "sap78475984ulkmf"
```

The audit table then contains a record including the full text of the command:

```
sa_role sso_role oper_role sybase_ts_role mon_role; alter role role1
   add passwd "******"; ; ; ; ; sa/ase;
```

The password is masked using a fixed-length string of asterisks.

For general information about full-text auditing, see *Security Administration Guide* > *Auditing* > *Manage the Audit System* > *Configure the Audit System* > *Full Text Auditing*.

For a full list of commands supported by full-text auditing, see *Security Administration Guide* > *Auditing* > *Audit Tables* > *Values in Event and Extrainfo Columns*.

16.11 Bucket Pool

The bucket fragment manager allows you to tune the memory manager according to the varying needs of different server components.

Each bucket fragment pool in the system has an entry in the configuration file, similar to:

```
[Bucket Pool:Network Buffers]
        enabled = 1
```

```
        bucket size = 32
        seed = 0
        maxfill = 0
        bucket size = 64
        seed = 0
        maxfill = 0
        bucket size = 96
        seed = 0
        maxfill = 0
        bucket size = 128
        seed = 0
        maxfill = 0
        bucket size = 160
        seed = 0
        maxfill = 0
        bucket size = 192
        seed = 0
        maxfill = 0
        bucket size = 224
        seed = 0
        maxfill = 0
        bucket size = 256
        seed = 0
        maxfill = 0
        bucket size = 288
        seed = 0
        maxfill = 0
        bucket size = 320
        seed = 0
        maxfill = 0
        bucket size = 352
        seed = 0
        maxfill = 0
        instances = 1
        instance seed = 0
        defragment = 0
        autotune = 0
        overhead = 0
```

Exact entries vary according to the bucket pool options.

Since the tuning process requires the knowledge of internal functionality, you can modify bucket pool configuration with the guidance of SAP ASE support.

See *System Administration Guide: Volume 1 > Setting Configuration Parameters > Bucket Pool Configuration Parameter Group*.

System Changes

This feature introduces system changes.

System Change	Description
`monBucketPool`	For more information about this new monitoring table, see *Reference Manual: Tables > Monitoring Tables > monBucketPool*.

System Change	Description
`sp_sysmon`	For information about the `Bucketpool Activity` report in `sp_sysmon` output, see *Performance & Tuning Series: Monitoring with sp_sysmon > Monitoring with sp_sysmon >Kernel Utilization > Bucketpool Activity*.

16.12 Table User-Defined Functions

You can create both inline and multi-statement table user-defined functions.

To create inline table-valued functions:

```
create function [ <owner>. ] <function_name>
    ([{ <@parameter_name> [as] <parameter_datatype> [ = default]}
        [ ,...n ]])
    returns table
          <table_type_definition>
    [ with recompile] ]
    [ as ]
    return <select_statement>
```

To create multi-statement table-valued functions:

```
create function [ <owner>. ] <function_name>
    ([{ <@parameter_name> [as] <parameter_datatype> [ = default]}
        [ ,...n ]])
    returns [<@return_variable>] table
          ( { <column_definition> | <table_constraint> } [ ,...<n> ] )
    [ with recompile]
    [ as ]
    begin
    <function_body>
    return
    end
```

> **Note**
>
> Table variables cannot be parameters to table user-defined functions.

See *Transact-SQL Users Guide > Transact_SQL Functions > Table User-Defined Functions*.

16.13 Table Variables

A table variable, declared in a SQL batch, contains not just a single value, but an entire table (that is, zero or more rows consisting of one or more columns).

Declare table variables at the beginning of a batch. In place of the usual datatype for a variable, the keyword `table`, followed by a list of column declarations must be given, similar to that given during a `create table` statement. For example:

```
declare @tv table (c1 int, c2 char(3))
```

> **i** Note
>
> Table variables cannot be parameters to table user-defined functions.

See *Transact-SQL Users Guide > Batches and Control-of-Flow Language> Local Variables > Declaring a Table as a Variable* and *Reference Manual: Building Blocks > Table Variables*.

16.14 Union-Derived Tables in Star Joins

You can include union-derived tables as dimension tables combined with a fact table hint.

See the *Performance and Tuning: Query Processing and Abstract Plans > Controlling Optimization > Query Plan Optimization with Star Joins > Star Join Hint*.

16.15 64-Bit Backup Server

SAP ASE versions 16.0 SP02 and later support a 64-bit version of Backup Server. Versions of SAP ASE earlier than 16.0 SP02 supported only a 32-bit version of Backup Server.

16.16 Binding Backup Server to a Socket

Use the `bind backupserver address` configuration parameter to bind Backup Server to a socket address.

See *System Administration Guide, Volume 2 > Developing a Backup and Recovery Plan > Using Backup Server for Backup and Recovery >Binding Backup Server to a Socket*.

See *System Administration Guide: Volume 1 > bind backupserver address*.

16.17 Job Scheduler Restart

The `js restart delay` configuration parameter sets the delay period between two Job Scheduler auto restart attempts after abnormal shutdown of Job Scheduler.

See *System Administration Guide: Volume 1 > js restart delay*.

16.18 Enable Resolve as Owner

Earlier versions of SAP ASE enabled the resolve as owner functionality with the `enable functionality group` configuration parameter. SAP ASE 16.0 SP02 and later enable the resolve as owner functionality with the `enable resolve as owner` configuration parameter.

See *System Administration Guide: Volume 1> enable resolve as owner*.

16.19 Limiting User Locks Per Session

The `sp_add_resource_limit...lock_count` limittype parameter allows you to limit the number of logical locks used simultaneously by a user process.

See the *Reference Manual: Procedures > sp_add_resource_limit* and *System Administration Guide, Volume 2 > Limiting Access to Server Resources > Understanding Limit Types > Limiting User Locks Per Session*.

16.20 Model Database Backlink Flag Set After dbcc shrinkdb_setup

Versions of SAP ASE 16.0 SP02 and later do not use the database level text back link flag. That is, 0x01 is not set in `sysdatabases.status` for newly-created databases, or after running `dbcc shrinkdb_setup` on the specified database.

See *Transact-SQL Users Guide > Shrinking Databases > Shrink Operations on Databases That Contain Text or Image Data*.

See *System Administrators Guide: Volume 2 > Transact-SQL Functions > Built-In Functions > Text and Image Functions*.

16.21 set_appcontext Changes

The character limitation length value is increased to 512 for `set_appcontext`.

See *Reference Manual: Building Blocks > Transact-SQL Functions > set_appcontext*.

17 Version 16.0 SP01

SAP Adaptive Server Enterprise version 16.0 SP01 introduces new features and enhancements

Feature	Description

Dynamic Histograms	Dynamic histograms allow you to specify columns for which a new range cell is dynamically maintained at the end of an existing histogram. For DML operations on this column, the new range cell includes the maximum value and dynamically maintains a row count. The boundaries of the original cells of the histogram remain unchanged, but the weights of the original cells are proportionally decreased as the weight of the new cell is increased. The syntax to enable dynamic histograms is:
	`update statistics <table_name> (<column_name>, ...)`
	`...`
	`[using dynamic_histogram [ on | off ]`
	`...`
	See *Performance and Tuning Series: Query Processing and Abstract Plans > Using Statistics to Improve Performance > Using Dynamic Histograms*.

Mount and Unmount Encrypted Databases	Mount and unmount encrypted databases the same as mounting unencrypted databases (the master key and database encryption keys stored in the `master` database must be the same). Earlier versions of SAP ASE did not allow you to mount or unmount encrypted databases. See *Database Encryption > Database Encryption > Mount an Encrypted Database*.
`select into` Support for Multistatement Transactions	Use the `select into` command against permanent and temporary tables in multistatement transactions (for example, in SQL batches, stored procedures, triggers, and so on). See the *Transact-SQL Users Guide > Transactions: Maintain Data Consistency and Recovery > Using select into in Multistatement Transactions*.
Compressed Index Support	Use the `create database ... index_compression` parameter to create databases that include compressed indexes. Use the `alter database ... set index_compression` to enable or disable index compression on existing databases. See the *Compression Users Guide > Index Compression > Create or Altering Databases to Use Index Compression*
`sp_help terse`	Determine the amount and content of information displayed by `sp_help`.
`sp_helpartition terse`	Determine the amount and content of information displayed by `sp_helpartition`.
`sp_helpindex terse`	Determine the amount and content of information displayed by `sp_helpindex`.

18 Version 16.0

SAP® Adaptive Server® Enterprise version 16.0 introduces new features and enhancements.

18.1 Increased Data Availability with Partition Locking

The new partition-level locking feature increases data availability by providing locking to a finer granularity, which allows access to other partitions for concurrent DDL and DML statements.

- Partition locking as an attribute of a partitioned table
- Schema locks, which allow enhanced partition-level operations to update table schema or metadata by achieving isolation from concurrent operations
- Lock compatibility and lock sufficiency, which support issues of locking and concurrency
- Partition-level locking, which allows access to data in tables while operations on a subset of partitions in a table (partition-level operations) are active

18.2 CIS Support for HANA Server

SAP ASE version 16.0 adds the native ODBC interface to the Component Integration Services (CIS) so you can connect directly to a Hana server from SAP ASE.

18.3 Relaxed Query Limits

SAP ASE 16.0 increases some column and table restrictions.

18.4 Query Plan Optimization with Star Joins

SAP ASE 16.0 introduces the `use fact_table` abstract plan hint, which specifies the central fact table in a star join query and triggers special query plan optimization strategies for the star join query.

The `use fact_table` abstract plan hint allows the query processor to choose a parallel hash join plan for the star join query. Parallel plans enable the query processor to push bloom filter probes (which allow for faster

joins between dimension and fact tables) below the `Xchg` operator, further reducing the number of qualifying rows from fact tables.

18.5 Query Performance Improvements

SAP ASE version 16.0 adds query performance improvements.

- Using dynamic instead of static thread assignments. Dynamic thread assignment allows SAP ASE to execute parallel query plans faster and with fewer resources. SAP ASE applies dynamic thread assignment to parallel lava query plans that are generated for `select` queries.
- Improving the performance for queries run on tables with compressed data, including removing columns required only by the search argument from the substitution list.
- Improves the performance on certain parallel queries that include `SORT` operators.
- Implements performance improvements and a reduction in resource usage for certain parallel query plans involving `HASH JOIN` operators.

18.6 Full-Text Auditing

SAP ASE versions 16.0 and later support full-text DML auditing, printing parameter names and values with sensitive parameters masked for DML when you enable the DML auditing options (including `<table_access>` and `<view_access>`).

SAP ASE records the full-text audit information for these commands:

- `select`
- `insert`
- `delete`
- `update`
- `select into`

18.7 Auditing for Authorization Checks Inside Stored Procedures

The audit option `sproc_auth` enables auditing for authorization checks that are performed inside system stored procedures.

18.8 Replacing Object Definitions

Replace existing object definitions with new definitions while preserving the original name, object ID, auditing options, and permissions.

The `create or replace` functionality creates a new object, or replaces an existing object with the same name. The `or replace` clause allows SAP ASE to implicitly drop and re-create an existing object of the same name and type within the database. The `or replace` clause changes the definition of the object, but preserves existing security attributes such as permissions and auditing options, allowing you to avoid having to save granted permissions and re-grant the permissions after the object is dropped and created.

18.9 Query Plan and Execution Statistics in HTML

A new Dynamic Thread Assignment model has been introduced to enhance the use of worker threads.

18.10 Index Compression

Compression in a relational database allows more efficiently storage of data, reduced memory consumption, and improve performance due to lower I/O demands.

Index compression supports:

- Index leaf page compression
- Both DOL and APL index leaf page formats
- Compression at the table, index, and local index partition levels

18.11 SAP JVM Support

SAP ASE versions 16.0 and later use SAP JRE to support Java applications.

18.12 Full Database Encryption

SAP ASE version 16.0 introduces the ability to encrypt entire databases, providing protection for an entire database without affecting existing applications.

Once you encrypt a database, all of its data, indexes, and transaction logs become encrypted. This encryption is transparent, so that users can perform operations on tables, indexes, and so on, as usual, without noticing any differences.

Earlier versions of SAP ASE allow only column encryption.

18.13 Scalability Enhancements and Features

SAP ASE version 16.0 improves scalability by improving run-time logging and lock, metadata, and latch management.

18.14 Monitoring Threshold-Based Events

SAP ASE version 16.0 and later include the ability to configure, record, and list threshold events.

18.15 Multiple Triggers

Create multiple triggers, as well as specify the order in which the triggers are fired after statement execution.

18.16 Residual Data Removal

Remove residual data, which strengthens the security of database data.

Some database operations that delete space do not always physically erase the data. This can pose a security threat, as this residual data may be visible to a user using the `dbcc` utility. To prevent this, SAP ASE version 16.0 introduces a feature that zeros out residual data without additional work when the user performs these database operations.

You can mark data as sensitive, and configure SAP ASE to erase its residual data after performing delete or update operations.

18.17 Configuration History Tracking

SAP ASE version 16.0 add the `sp_confighistory` system procedure, which allows you to track changes to the server configuration.

`sp_confighistory` records changes to SAP ASE configurations, including which configuration option was changed, the old and new values, which user made the change, and when the change was made. Changes are stored in the `ch_events` view of `sybsecurity`, which contains one row for each configuration change event.

18.18 Cyclic Redundancy Checks for dump database

SAP ASE adds a cyclic redundancy check for accidental changes to raw data for database or transaction dumps created with compression to check and for verification that the compression blocks can be correctly read and decompressed.

18.19 Calculating the Transaction Log Growth Rate

SAP ASE version 16.0 adds the sp_logging_rate system procedure, which calculates the transaction log growth rate for a specified time period.

18.20 System Changes for SAP ASE Version 16.0

SAP ASE 16.0 contains system changes.

Global Variables

New global variables:

Variable	Description
`@@trigger_name`	Returns the name of the trigger that is currently executing
`@@tranrollback`	Returns the type of rollback encountered, if any.

Configuration Parameters

New parameters:

Parameter	Description
enable utility lvl 0 scan wait	Allows you to run `alter table ... add | drop partition` commands while Adaptive Server runs isolation level 0 scans.
large allocation auto tune	Configures SAP ASE preallocate large amounts of memory for query execution, which reduces procedure cache contention
engine local cache percent	Allows you to modify the engine local cache as a percentage of procedure cache.
enable large chunk elc	Enables large allocation in the engine local cache.
aggressive task stealing	Sets the SAP ASE scheduler task stealing policy to aggressive.
max network peek depth	Specifies how many levels deep SAP ASE peeks into a connections operating system receive buffer for a pending cancel.
user log cache queue size	Determines whether a queuing strategy is used for logging.
threshold event monitoring	Enable or disable threshold-event recording.
threshold event max messages	Determines the number of events stores in the `monThresholdEvent` table.

Commands

Changed commands:

Command	Description
alter database	Fully encrypt an existing database.
alter index	Change the compression state of future index inserts or updates using the `index_compression` clause.
alter table	For: • Index compression – change the compression state of future index inserts or updates using the `index_compression` clause. • Multiple triggers – the table owner can disable any or all of the multiple triggers defined on that table. • Residual data removal – automatically remove residual data following delete operations.
create archive database	Create a fully encrypted archive database.

Command	Description
create database	Create a fully encrypted database.
create default	Replace an object's definition using the `or replace` clause.
create encryption key	Create a database encryption key.
create function	Replace an object's definition using the `or replace` clause.
create function (SQLJ)	Replace a user-defined SQLJ function's definition using the `or replace` clause.
create index	Compress an index or index partition with the `index_compression` clause.
create procedure	Replace an object's definition using the `or replace` clause.
create procedure (SQLJ)	Replace a SQLJ procedure definition using the `or replace` clause.
create rule	Replace an object's definition using the `or replace` clause.
create table	For: • Index compression – compress indexes on a specified table using the `index_compression` clause. • Residual data removal – Create a new table that automatically removes residual data when deletions occur.
create trigger on order	For: • Multiple triggers – Create multiple triggers, as well as specify the order in which the triggers are fired after you execute a command. • `or replace` – replace an object's definition using create trigger using the `or replace` clause.
create view	Replace an object's definition using create view using the `or replace` clause.
drop encryption key	Delete the database encryption key from the `sysencryptkeys` table in the `master` database
drop trigger	Use to remove or replace an existing trigger, or multiple triggers.
dump database	Adds a cyclic redundancy check for accidental changes to raw data for database or transaction dumps created with compression.
kill	Adds the `with force` parameter if you cannot terminate the process with the regular `kill <spid>` parameter.
load database	Adds a cyclic redundancy check for accidental changes to raw data for compressed database or transaction dumps.
select	Issuing select statements that reference only @variables, @@global variables, and constants on SAP ASE version 16.0 and later in chained mode do not start new transactions.
select into	Creates an index compressed table by selecting from an existing table.
set	For:

Command	Description
	• Multiple Triggers – `set show_trigger_execution` prints to the error log the name of the trigger as it is fired, to help you identify where any processing problems are occurring. • Residual data removal – `set erase_residual_data {on | off}` enables the ability to erase residual data.

Functions

New built-in function:

Function	Description
`dbencryption_status`	Reports on the encryption/decryption status and progress of a database.

System Stored Procedures

New system stored procedures:

Procedure	Description
`sp_confighistory`	Creates the `ch_events` view and displays changes made to SAP ASE configuration.
`sp_ dropglockpromote_ptn`	Removes partition lock promotion values.
`sp_droprowlockpromote_ptn`	Removes row partition lock promotion threshold values from a database or table.
`sp_helptrigger`	Use to: • View a list of all the triggers for each of the DML actions against a table • See the orders of those triggers that specified the `order` clause when created
`sp_jsconfigure`	Configures the Job Scheduler Agent.
`sp_loggging_rate`	Calculates the transaction log growth rate for the specified time period.
`sp_setpglockpromote_ptn`	Sets page-lock promotion thresholds for partitions at the server, database, and the table.
`sp_setrowlockpromote_ptn`	Sets row-lock promotion thresholds for partitions at the server, database, and the table.

Changed system stored procedures:

Procedure	Description
`sp_audit`	Adds these auditing options:

Procedure	Description
	• `config_history` – enables or disables auditing for configuration history. • `sproc_auth` – enables or disables auditing for authorization checks that are performed inside system stored procedures.
sp_chgattribute	Adds the `pnt_locking` table attribute, which enables and disables partition-level locking.
sp_ clusterlockusage	Output enhanced to print cluster lock usage specific to partition lock information.
sp_dboption	Removes residual data from delete operations in SAP ASE using the `erase residual data` parameter.
sp_depends	Lists the multiple triggers associated with the table.
sp_encryption	Supports encrypted databases
sp_familylock	Output enhanced to include the `partitionid` column.
sp_helpdb	Output enhanced to include information about encrypted databases
sp_lock	Output enhanced to include the `partitionid` column.

System Tables

New system table:

Table	Description
ch_events	Contains one row for each configuration change event. `ch_events` is located in the `sysmgmtdb` database.

Changed system tables:

Table	Description
sysattributes	There are no new columns, however the encrypted database feature introduces 43, a new class that signifies database encryption.
sysconstraints	There are no new columns, however now includes information about multiple triggers.
sysobjects	Changes what `sysobjects` saves for multiple triggers.

Monitoring Tables

New monitoring table:

Table	Description
monThresholdEvent	Includes one row for each event recorded by SAP ASE.

Changed monitoring tables:

Table	Description
monCachedStatement	• Updates metrics for some columns in 5 second increments. • Increments the `UseCount` column when statement begin execution • Increments the value for columns that describe maximums for currently executing statements.
monDeadLock	The `partitionid` column displays the unique identifier for the partition.
monLocks	The `partitionid` column displays the unique identifier for the partition.
monOpenObjectActivity	Adds a number of columns
monOpenPartitionActivity	Adds a number of columns
monProcess	ClientDriverVersion column displays the connectivity driver used by the client program.
monRepLogActivity, **monRepScannersTotalTime**, and **monRepSenders**	Require that you enable the `activate monitoring` configuration parameter to start collecting monitoring data.
monRepScanners	Changes the Status column.
monRepScannersTotalTime	changes the name of the `MRPBootstrapTime` column to `BootstrapTime`.
monSysExecutionTime	(Cluster Edition only) The `ExecutionTime` column displays values for NetworkIO, DeviceIO, and CIPCIO.
monTables	The `Description` column in SAP ASE 16.0 and later supports 512 characters. Previous releases supported 255 characters.
monTableColumns	• The `Description` column in SAP ASE version 16.0 and later supports 512 characters. Previous versions supported 255 characters. • The `Label` column in SAP ASE version 16.0 and later supports 150 characters. Previous versions support 50 characters.

Utilities

Changed utilities:

Utility	Description
ddlgen	Supports encrypted databases
sybmigrate	Supports encrypted databases

Utility	Description
sybrestore	Supports restoring an SAP ASE after a master database corruption

19 Version 15.7

For a list of new and enhanced features in SAP Adaptive Server Enterprise 15.7, see SAP Adaptive Server Enterprise 15.7

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.

www.sap.com/contactsap

THE BEST RUN SAP