

Sybase® Adaptive Server® Enterprise

New features in ASE 15.7

TABLE OF CONTENTS

- 1 Introduction
- 1 Introducing ASE 15.7
- 2 Optimized for Storage
- 3 Application Developer Productivity
- 4 Enhanced Diagnostics & Monitoring
- 4 Security Administration
- 5 Performance Enhancements
- 5 Summary

INTRODUCTION

The last few years have seen an explosive growth in the data volumes being managed by organizations. This is not only true with respect to historical data that needs to be captured and analyzed in large data warehouses, but also for operational data stores running the transactional as well as operational reporting systems. The transition from an average of a few hundred gigabytes managed by the operational store a decade ago, to the many terabytes being managed today, has happened rapidly. Various industry surveys point to a 25-30% or more annual data growth for an operational store. The average operational database now measures in terabytes.

Sybase Adaptive Server Enterprise 15 was designed to meet the challenges of terabyte-scale operational environments. Such environments need the performance and scale to meet the most demanding Wall Street trading systems, but at the same time the ability to run intra-day and end-of-day reporting and operational decision support systems. This ability to handle “mixed workloads”— the happy coexistence of demanding transactional systems and operational reporting systems — is where ASE 15 excels. ASE 15 introduced significant technology updates throughout the ASE stack, from a revamped query processing engine to support for semantic partitions, all technologies precisely designed to seamlessly handle terabytes of data.

Sybase ASE 15 was introduced in September 2005. Since then ASE 15.0.1, ASE 15.0.2, ASE 15.0.3 and ASE 15.5 have added an abundance of new capabilities to the product line. Most notably, ASE 15.0.2 introduced many new features sought after by application developers and database administrators such as user-defined SQL functions, instead of triggers, row locked system catalogs, xmltable, several security and encryption enhancements, new statistical aggregate functions, and numerous other performance, monitoring and diagnostics features.

ASE 15.0.3 introduced state-of-the art Java support by providing the ability to plug in an off-the-shelf Java Runtime Environment (JRE) and JVM components, such as J2SE, into ASE. To enhance replication performance, SQL statement replication was added, which along with Sybase Replication Server, greatly speeds up replication of large batch jobs. In addition, the QPTune utility as well as the compatibility mode (for ASE 12.5x) for the query processor, allows customers to better manage their migrations to ASE 15.

A fully integrated in-memory database capability was added in ASE 15.5. This seamless integration of an in-memory database within a traditional disk based database was a first in the industry. In-memory databases are being used for caching data (e.g. reference and compliance data in trading systems), processing of large volumes of market data, processing of large batch jobs where only derived results will eventually need to persist, as well as a fast cache for ASE temporary objects.

Sybase also introduced a shared-disk-cluster edition to meet the continuous availability demands of mission critical enterprise class ASE installations. The ASE 15.0.1 CE (Cluster Edition) release was the culmination of several years of R&D. Sybase’s continued aggressive investment in its database technologies has been rewarded by impressive growth in its acquisition of new customers and growth in market share.

INTRODUCING ASE 15.7

The market appeal and reach for Sybase ASE has dramatically increased with the 2010 acquisition of Sybase by SAP, the leader in business applications. SAP Business suite applications, as well as most all SAP applications running on SAP Netweaver technology, now run on Sybase ASE. These are indeed exciting times for Sybase ASE.

Sybase ASE 15.7 is the latest release in the ASE 15x product line; and it’s a feature-packed release. This document provides an overview of some of these new capabilities, and the benefits customers can expect from moving to ASE 15.7.

ASE 15.7 offers important enhancements for:

- Optimizing for Storage
- Application Developer Productivity
- Performance
- Enhanced Diagnostics & Monitoring
- Security Administration

OPTIMIZED FOR STORAGE

Data growth is a serious challenge for enterprises. It can be attributed to a number of reasons: regulations requiring longer data retention, increased number of transactions, growth in unstructured data (images, text, audio and video), or the need to maintain data for longer periods of time for decision support systems. While disk itself has become cheaper, the growth in data continues to cause serious challenges for data centers.

The cost associated with data is not singular and limited to a single copy of the data. On the contrary, one unit of data on average translates into a minimum of 3-5 copies of it within the data center itself, as well as additional copies beyond the data center. Copies of the production data exist in development and testing systems, back-up and archived systems, secondary decision support systems as well as offsite backup and disaster recovery systems. Storage costs quickly get compounded when hardware costs, software license costs, data center floor space and power consumption costs, data transfer and bandwidth costs, and labor costs are thrown into the overall cost of maintaining this data. Clearly the oft-heard refrain “disk is cheap” can only be stretched so far before costs start piling up.

And this does not even consider performance costs due to I/O demands and cache memory requirements when processing this data!

ASE 15.7 introduces the following to help reduce storage requirements:

- Compression: Regular data and LOBs
- In-row LOBs
- Create nonmaterialized non-NULL columns without data copy
- Ability to shrink the transaction log

Compression

In prior releases ASE customers had the ability to compress backups thus helping reduce offline storage costs. ASE 15.7 now allows in-database compression for active data sets. Both regular relational data as well as unstructured data (large objects or LOBs) can be compressed. Early internal benchmarks indicate that compression ratios can readily hit the 50-80%.

ASE 15.7 uses a number of compression strategies to achieve high compression ratios. This includes compression within a single row to compress away empty spaces/zeros in fixed length columns. At a page/block level this includes both page dictionary and page index compression strategies. Repeated data items as well as repeated sets of data items are replaced by a single reference resulting in dramatic savings for duplicated data. Data can be compressed at database, table or partition level. This allows for customers to separately make compression choices for active or inactive data.

ASE 15.7 also supports in-database LOB compression. Given that LOBs can be very large in size (up to 2GB), compression can result in very significant space savings. FastLZ and ZLib compression techniques are supported, which are lossless compression algorithms. While the former provides lower CPU usage and execution times, the latter provides higher compression ratios.

Customers have the flexibility to choose LOB compression at database, table or column level.

In-Row LOBs

Modern applications increasingly need to manage and manipulate unstructured data. This data is stored as LOB columns (e.g. text, images, unitext, XML). LOBs can vary in size from a few bytes to gigabytes. ASE optimized for LOB storage by allocating separate storage blocks that could be pointed to from the parent data row. This was very efficient for handling larger LOBs. ASE 15.7 enhances this capability by now also supporting in-row LOBs for the situations when LOBs are fairly small and can readily fit within the parent row.

This has two advantages: storage savings by optimizing for the space allocated for small LOBs, as well as performance benefits for retrieval of small LOBs by reducing the I/O overhead.

ASE seamlessly handles the migration of LOBs from in-row to off-row storage (or vice versa) to always optimize for the storage required as LOBs change in size due to updates.

ASE 15.7 offers important enhancements for:

- Optimizing for Storage
- Application Developer Productivity
- Performance
- Enhanced Diagnostics & Monitoring
- Security Administration

Non-materialized Non-NULL Columns Without Data Copy

In the past adding non-NULL columns to an existing table required a complete data copy. This can be particularly cumbersome for very large tables. ASE 15.7 now allows creation of non-materialized non-NULL columns. This means columns get created without actually making a data copy (which can require a large amount of temporary space for the data copy, equal to the original table size). This saves both space and time.

When such rows are actually updated they get instantiated (physically stored in the row) behind the scenes.

Shrinking Transaction Log

ASE 15.7 allows for the ability to shrink a transaction log. Transaction logs can often grow very large due to a number of reasons: enlarged to handle log-full situations, to support one-time operations that may require lots of space or even due to generous estimates during capacity planning.

APPLICATION DEVELOPER PRODUCTIVITY

ASE 15.7 offers numerous enhancements designed to enhance the productivity of application developers. These include:

- Simpler and more efficient LOB retrieval and manipulation
- Greater flexibility around concurrency and isolation levels
- Enhanced SQL expressiveness
- Support for scripting language drivers

ASE 15.7 greatly improves on the handling of LOBs. In addition to the support for in-row LOBs discussed above, ASE 15.7 introduces the concept of LOB locators. LOB locators are essentially handles (pointers) to LOBs cached in ASE server memory. Client applications can send and receive locators as host variables and parameter markers. When a client application needs to fetch, search, update LOBs, it communicates with the server through the LOB locators, rather than pass actual LOBs to and fro. This greatly reduces network traffic between the client and the server. It also reduces the need for client side memory.

A number of T-SQL commands are provided to work with LOBs. These commands allow easy manipulation of LOBs (truncation and concatenation of LOBs, overwriting portions of a LOB, extracting substrings in a LOB, and so on). ODBC clients can use these T-SQL commands to manipulate LOBs. The jConnect JDBC driver supports access to the locators using Blob, Clob and NClob classes.

In addition to the above, LOBs can now also be passed as parameters to stored procedures. This adds greater flexibility to manipulate LOBs within stored procedures and reduces code complexity for developers.

ASE 15.7 also provides a number of T-SQL updates that allow for greater SQL expressiveness as well as control over concurrency and isolation level semantics. Some of these include:

- Support for MERGE statement to allow for transferring rows from a source table (table, view derived table) into a target table.
- SELECT FOR UPDATE to allow selected rows to be locked for a later update
- Ability to release read cursor locks at cursor close for active transactions for all isolation levels
- Various nested select statement enhancements
- Ability to run all system procedures and select DDLs in chained transaction mode

Another area where ASE 15.7 benefits application developers is through Sybase provided and supported scripting language drivers. Initial support is provided for python, PHP and perl. Application developers can now write their database applications in these languages. This makes it easy to leverage the vast pool of developer talent fluent in these languages.

ASE 15.7 introduces the following to help reduce storage requirements:

- Compression: Regular data and LOBs
- In-row LOBs
- Create nonmaterialized non-NULL columns without data copy
- Ability to shrink the transaction log

ENHANCED DIAGNOSTICS & MONITORING

ASE 15.7 provides numerous enhancements around diagnostics and monitoring. Some of these include:

- Ability to save abstract plan information in the statement cache
- Ability to output cached plan in XML
- Ability to analyze dynamic parameters before running a query which can help to avoid inefficient query plans
- Ability to monitor lock timeouts
- Ability to display statistics and histograms for various objects, via a system procedure (T-SQL equivalent of the standalone *optdiag* utility)
- Ability to monitor cursor statements
- Ability to analyze the query plan generated to understand the various factors behind ASE's choice of the query plan (via the *sp_opt_querystats* system procedure)

These and a number of other enhancements allow application developers to probe deeply into the decisions being made by ASE optimizer and identify opportunities to further tune and optimize for performance. Monitoring tables have been enhanced to provide ever more information for the sophisticated application developers and DBAs.

ASE 15.7 also introduced the *sybdiag* utility. This is a powerful new utility that can help expedite customer case management and problem diagnosis processes. It is a Java-based tool that collects comprehensive ASE configuration and environment data. This tool generates a .zip output file comprising HTML and data files that can be unzipped and viewed in a web browser. The information collected includes operating system and environment data, ASE configuration and monitoring data, and various files and scripts. Note that it does not collect information such as logins, passwords, user lists or information from data tables.

The *sybdiag* tool can be used by customers to provide Sybase customer services & support important diagnostics information through a single command, thus reducing communication cycles and time burden on both Sybase support and customers.

SECURITY ADMINISTRATION

Several security improvements are introduced in ASE 15.7. These include:

- Object ownership change
- Support for login profiles
- Split knowledge/Dual control for protecting encryption keys (a requirement for the payments card industry data standard, PCI-DSS)
- Improved encryption for FIPS 140-2 compliance — encrypting role passwords with SHA256, and encrypting external server passwords and hidden SQL text with AES256
- Single sign-on and end to end Kerberos credential forwarding

ASE 15.7 makes it easy to change object ownership from one owner to another. This often happens when the person who created an object (table, procedure etc.) may not be the one maintaining it. This can be due to employee changing responsibilities or leaving the company, mergers and acquisitions, or outsourced development and maintenance.

ASE 15.7 introduces the ability to conveniently change object ownership.

Another powerful feature that allows for easy security administration is the support for login profiles. Before ASE 15.7 each ASE login had to be managed individually. This can be cumbersome and makes it difficult to implement changes in policies that often affect large number of users.

ASE 15.7 introduces the login profiles which are SQL-defined attribute containers of login characteristics and their values (such as default database, login script, auto activated roles, authenticate with). Groups of logins can share login profiles. These are easy to administer and provide flexibility to implement security policies and comply with changing regulations.

PERFORMANCE ENHANCEMENTS

ASE 15.7 focuses on two specific areas for performance:

- Reducing query latency, especially for dynamic SQL
- ASE kernel updates to better utilize parallel hardware

ASE 15.7 introduces various internal optimizations during query execution that reduce query execution time overhead. These include optimizations around caching and reusing information from query compilation across multiple executions, reducing metadata sent across the network for result sets, usage of a global cache for dynamic SQL to cache prepared statements to be reused across different sessions, more efficient execution of queries that do not touch a table and have no result set, efficient handling of cursor *fetch* statements, as well as ability to cache cursor query plans in the statement cache.

ASE 15.7 also provides an updated kernel that can better leverage parallel hardware. The kernel updates enhance ASE scalability on systems with very large number of processors, processor cores, and hardware threads. Updates also provide better integration of I/O event processing and reduce I/O overhead within the ASE server. These enhancements will not only improve I/O performance but also help with operation of ASE in virtualized environments.

SUMMARY

ASE 15.7 is a feature rich ASE release that significantly builds on the ASE 15x product line. ASE 15.7 is the release that SAP applications run on. ASE 15.7 not only provides a powerful database platform for SAP customers, but also provides existing Sybase ASE customers new capabilities that can be readily leveraged. ASE 15.7 is a continuation of the ASE 15x product line and existing ASE 15x customers can easily move to ASE 15.7 from earlier releases.