

Sybase ASE – Lightweight Compression for Backup Server

TABLE OF CONTENTS

- 1 Executive Summary
- 1 Objective
- 2 Performance Scenario
- 2 Setting up and Using the Backup System
- 3 System Configuration
- 3 Results
- 4 Performance
- 6 Conclusion

EXECUTIVE SUMMARY

While much of the hardware and software in an IT infrastructure may be replaceable, a customer's data may be the one thing that is irreplaceable. This makes the backup and restore functionality a critical feature of any transaction-intensive database applications. There may be a variety of different ways that a failure may occur, but in order to restore the database to its original state, a common solution is the Backup Server Compression Module.

Sybase has supported backup and restore functionality in its flagship product, Adaptive Server® Enterprise (ASE), since 1991. Sybase is well known for the inherent performance advantages that make backing up and restoring infrastructure ideal for vital mission-critical, transaction-intensive database applications. The ability to maintain high performance, reliability, and scalability for database backups is the deciding factor for many customers.

Sybase ASE supports backup and restore activity using the Backup Server Compression Module. Backup Server supports compression for database backups. Data is compressed during the dump operation based on the specified compression level. It is uncompressed at load time. The compression method supported by ASE 15.0.3 and earlier provides compression levels 1 – 9; it implements the LZW compression method, which uses ZLIB library for compressing and uncompressing data. Compression level 1 is fastest in terms of execution time while compression level 9 is best in terms of archive space. This compression method is a CPU-heavy process that slows down other applications, making it difficult for customers to make effective use of resources and provide high performance applications.

Data compression always depends on important trade-off parameters, such as execution time, CPU usage and archive space. The new compression levels 100 and 101 introduced in the ASE 15.5 Backup Server Compression Module use the FastLZ library, which implements the LZ compression method with some important optimizations. This results in faster execution time and drastically reduces CPU usage, making it a lightweight process. All this comes with a marginal increase in archive space.

Internal tests with the Backup Server Compression Module of ASE 15.5 on the Linux platform produced superior performance results. This validates the introduction of new compression levels 100 and 101 as less CPU intensive and faster than existing compression levels 1 to 9. Lightweight compression introduced in ASE 15.5 is better with respect to execution time and CPU usage, at a minimal cost of archive size, than existing compression methods.

OBJECTIVE

For most customers, a database needs to be backed up frequently enough to synchronize with updated critical data from other applications. This requires that backups run in background without affecting other applications, and in shorter execution times. To fully validate the compression performance of the Backup Server Compression Module of ASE, Sybase conducted internal tests to measure various performance parameters for data compression.

The objective is to provide a better compression technique for database backups in terms of CPU usage and execution time with only a small increase in archive space.

Performance objectives:

- Improve overall performance for compressed backups
- Reduce CPU usage to make it a lightweight compression
- Reduce execution time compared to existing compressed backup time
- Ensure only a minor increase in archive space

PERFORMANCE SCENARIO

The process for testing and measuring performance parameters for backup data compression was designed to simulate a typical backup and restore scenario, in which dump and load activity is executed in the background while the server is kept busy with other transactions and applications running on same CPU. The execution time was measured in milliseconds with start and end times noted before and after execution of dump and load activity. The percentage of CPU usage was monitored manually using the 'top' command on the Linux platform; a range of CPU usage values were noted for dump and load activities. The dump and load activities were executed using three stripes, and archive size was measured for all three archive devices using the 'ls -l' command on the Linux platform.

Each test ran for roughly 10-15 minutes, with a test suite running for an hour. Over a 10 hour period, each test suite was repeated 10 times. The design was to sequentially run dump/load tests for the new compression level 100, 101, and the traditional compression level 1 with the same database 'production' under the same environmental conditions for ASE 15.5. The performance numbers were measured for every test suite run, and the results were averaged.


SETTING UP AND USING THE BACKUP SYSTEM

The test database 'production' with 82GB spanning five different devices and separated data/log segments was used for testing performance numbers. We used the "dump database...with compression=" command to specify the compression level for each backup request. We also used "load database" command to restore the database for the test database 'production'.

Here, we compare performance results for compression level 100 against compression level 1, as our primary objective is to provide lightweight and fast compression.

Data flow for enhanced compression during backup/restore activity in ASE 15.5

Dump database...with compression=100



Backup/Restore Syntax

The following commands illustrate backing up and restoring a Sybase ASE database using the Backup Server Compression Module. The compression level to be used for compressed backups is specified by the “with compression=” option.

The existing compression supports compression levels 1 to 9 where 1 represents fast compression in terms of execution time, and level 9 is used for best compression in terms of archive space.

The new compression levels 100 and 101 introduced in ASE 15.5 adheres the same command syntax.

Compression level 0 indicates no compression.

Backup command:

```
dump database <database name> to <archive name> with compression = <compression level>
```

Restore command:

```
load database <database name> from <archive name>
```

Tests

- *dump database production to “prod1.dmp” stripe on “prod2.dmp” stripe on “prod3.dmp” with compression = 1*
Dumps database ‘production’ using traditional compression level 1 in three different stripes.
- *dump database production to “prod1_fast.dmp” stripe on “prod2_fast.dmp” stripe on “prod3_fast.dmp” with compression = 100*
Dumps database ‘production’ using new compression level 100 in three different stripes.
- *load database production from “prod1.dmp” stripe on “prod2.dmp” stripe on “prod3.dmp”*
Loads database ‘production’ using compression level 1.
- *load database production from “prod1_fast.dmp” stripe on “prod2_fast.dmp” stripe on “prod3_fast.dmp”*
Loads database ‘production’ using faster compression level 100.

SYSTEM CONFIGURATION

Server Machine	HP Proliant DL580, 4 X Intel Xeon 2.0GHz
Total Memory	8GB
Server Software	Red Hat Enterprise Linux Premium Edition 4.4 Adaptive Server Enterprise 15.5 Backup Server 15.5

RESULTS

- Comparative results
Database: *production*
Size: *82GB*
No. of stripes: *three*

ASE 15.5 Results:

Dump database 'production'		ASE 15.5 Release	
Parameter	New Compression Level 100	New Compression Level 101	Traditional Compression Level 1
ASE normal workload	826 sec	851 sec	854 sec
ASE heavy workload	870 sec	887 sec	892 sec
Size of archive device:	1.83G + 1.83G + 1.83G	1.8G + 1.8G + 1.8G	1.42G + 1.42G + 1.42G
% CPU usage range:	16 – 30	20 – 35	75 – 98

Load database 'production'		ASE 15.5 Release	
Parameter	New Compression Level 100	New Compression Level 101	Traditional Compression Level 1
ASE normal workload	1015 sec	949 sec	932 sec
% CPU usage range:	10 – 25 (rarely upto 50)	15 – 30 (rarely upto 60)	20 – 40 (rarely upto 100)

* Compression level 1 – Traditional compression in ASE versions earlier than 15.5

* Compression level 100 – New compression in ASE 15.5

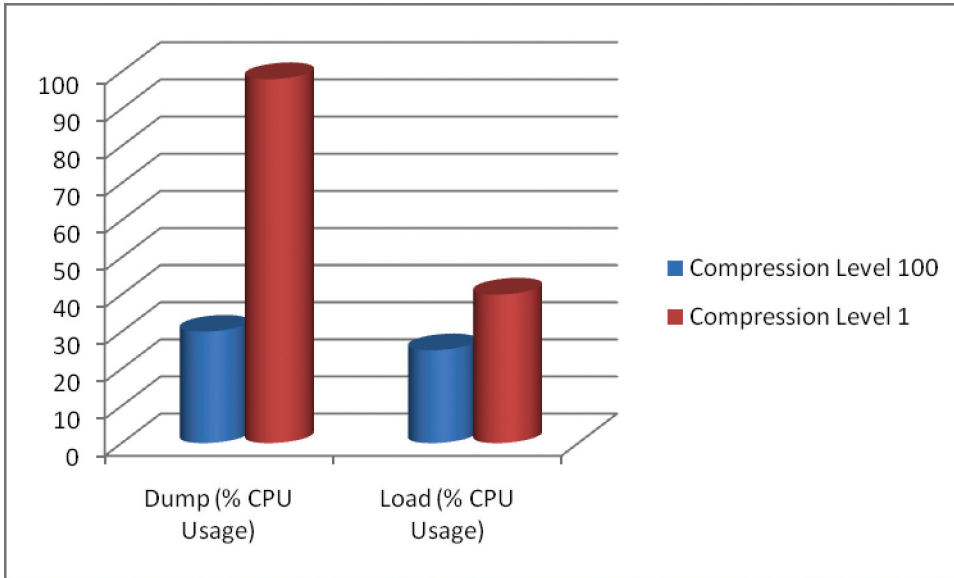
PERFORMANCE

Sybase ASE database backup and restore with the lightweight compression method improves the performance of backup and restore operations.

Backups with new compression levels introduced in ASE 15.5 drastically reduces the percentage of CPU usage for backup and restore commands, and results in minor reduction in the execution time for backups.

Internal testing with an 82GB Sybase ASE 15.5 database shows a greater improvement in backup and restore operations with the new compression levels when compared to existing traditional compression levels.

The following illustrates that the new lightweight compression in ASE 15.5 uses 50-60% less CPU, and 10-20% less execution time than traditional Sybase ASE Backup Server compression.



* Compression level 1 – Traditional compression in ASE versions earlier than 15.5

* Compression level 100 – New compression in ASE 15.5

The bar chart above maps the percentage of CPU usage for compression level 100 and level 1 for dump and load operations. The numbers show that compression level 100 is 50-60% less CPU-intensive during dump operations and 40-50% more lightweight for load operation of database 'production' when compared to level 1.

Also, the execution time for dump/load operations with compression level 100 is within a 10-15% limit when compared to compression level 1. Thus, there is nominal impact on execution time during dump/load operations with a significant reduction in CPU usage.

CONCLUSION

- **Ease of use and seamless integration**

Existing users of Sybase ASE can seamlessly perform database backup and restore operations using this new compression. The existing backup and restore commands of Sybase ASE with the compression option have been extended to accept the new compression levels 100 and 101 introduced in ASE 15.5. This provides a familiar command line interface for existing Sybase ASE users.

- **Better performance for Backup and Restore**

The new lightweight compression technique with compression level 100 performed better in terms of percentage of CPU usage and execution time at a marginal cost of archive space as compared to traditional compression level 1. The results obtained during testing showed a significant performance improvement. Lightweight compression reduced CPU usage values to 40-50% of the values obtained with pre-ASE 15.5 compression. Also, the execution time was less to 70-80% of the time needed with pre-ASE 15.5 compression. This performance improvement was at cost of a 25-30% increase in archive space.

The Backup Server Compression Module of Sybase ASE 15.5 supports a new, lightweight compress method that outperforms compression method supported in earlier versions of ASE. The new compression method provides faster and more efficient compression during backup and restore activities.