Using SmartDedupe&SmartCompression of Huawei Dorado6000 V3 Storage in Database Environments

Databases are widely adopted in the enterprise, as they have become the foundation of all business systems. Data stored in databases is the enterprise's core asset. Enterprises are grappling with the challenges presented by explosive data growth, which significantly raises storage costs and complexity. Business data is often filled with significant amounts of redundant information, taking up valuable disk capacity.

Wang Zhilei

Storage Solution Development Dept, IT, Huawei Enterprise BG

2017-03-30

Issue 1.0

Huawei Technologies Co., Ltd.

Contents

1 About This Document

1.1 Overview

Huawei OceanStor Dorado6000 V3 all flash storage systems (Dorado6000 V3) use SmartDedupe&SmartCompression, smart data deduplication and data compression technology, to reduce storage space occupied by redundant data and improve data transfer, processing, and storage efficiencies of storage systems. This document introduces SmartDedupe&SmartCompression and explains how it can be leveraged in database environments to realize data reduction and save storage space.

1.2 Intended Audience

- Huawei employees
- Partners
- Customers

1.3 Customer Benefits

SmartDedupe&SmartCompression of Dorado6000 V3 maximizes storage efficiency by decreasing the amount of physical storage required to house database service data.

1.4 Key Components

- Hardware
  - Huawei Dorado6000 V3 all flash storage systems V300R100C00
  - Huawei RH2288V2 servers
- Software
  - Multipathing software: UltraPath for Linux 8.01.024
  - Database software: Oracle Database 11.2.0.4
  - Cluster software: Oracle RAC 11.2.0.4
  - Test tool: SLOB2.2

- Operating system
  Red Hat Enterprise Linux 6.5

1.5 Load Model

This document describes how to use the Silly Little Oracle Benchmark (SLOB) testing tool to simulate the Oracle online transaction processing (OLTP) service load model. SLOB is a framework to test platform performance and stability. It provides a methodology to determine whether the underlying platform of high-performance Oracle databases is correctly deployed. SLOB uses SQL in an Oracle database to generate maximum physical I/Os while utilizing the least possible amount of host CPU. Underlying hardware is driven to the maximum performance to analyze restrictions of platform deployment on Oracle database performance. On the I/O layer, the load model provides the random access of small data blocks at varying read:write ratios, representing a typical OLTP load model.

2 Application Scenario

2.1 Application Scenario

This document applies to the following database scenario (actually measured). Test results vary with different scenarios and Huawei reserves the right of final interpretation.

- Oracle 11.2.0.3 and later

2.2 Precautions

To ensure deduplication and compression effects without decreasing performance, meet the following requirements in actual configurations:

- For databases, you are advised to only enable the compression function because the unique timestamp marks generated during storing Oracle database tables do not allow deduplication.
- Huawei OceanStor Dorado V3 all flash storage systems adopt global deduplication in which LUN quantity will not affect data reduction. You can configure multiple LUNs to increase concurrency and improve performance.
- The compression function is not enabled for Oracle applications.

3 Feature Configuration in OLTP Environments

3.1 Solution

3.1.1 Architecture

The solution uses four servers and two 16 Gbit/s Fiber Channel switches to build a dual-switch network and two dual-node Oracle Real Application Cluster (RAC) clusters for a test.

The test is executed to evaluate how SmartDedupe&SmartCompression of OceanStor Dorado6000 V3 storage reduces redundant service data and improves storage efficiency.

Figure 3-1 Architecture

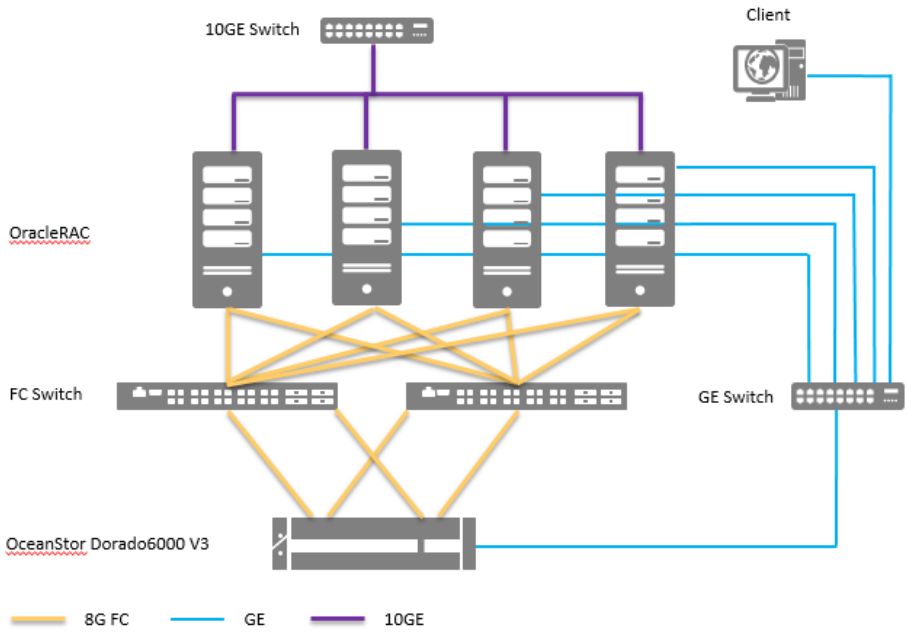

3.1.2 Hardware Configuration

Table 3-1 lists hardware configuration.

Table 3-1 Hardware configuration

Device	Component	Quantity
Server	RH2288 V2: 256 GB memory 2 x Intel(R) Xeon(R)E5-2660 CPUs 2 x Qlogic 16 Gbit/s FC dual-port HBAs 1 x Intel 10 Gbit/s Ethernet HBA	4
Storage	OceanStor Dorado6000 V3: 2 x controllers 1 x 2 U 25-slot disk enclosure 25 x 3.6 TB SSD disks 2 x four-port SmartIO I/O modules	1
Fibre Channel switch	SNS2224	2
GE switch	S5700	1

3.1.3 Software Configuration

Table 3-2 lists software configuration.

Table 3-2 Software configuration

Item	Software
Operating system	Red Hat Enterprise Linux 6.5
Multipathing software	Huawei UltraPath 8.06.063
Database cluster software	Oracle Grid Infrastructure 11.2.0.4
Database software	Oracle Database 11.2.0.4
Testing tool	SLOB2.2
OceanStor Dorado6000 V3	V003R100C00

3.1.4 Storage Configuration

Dorado6000 V3 is used as the storage and is configured as follows: creating one disk domain using the low hot spare policy, one storage pool in RAID 6 from the disk domain, and twenty-one 20 TB thin LUNs (all configured with the compression function). Table 3-3 is an

example of configurations. You are not required to strictly follow the configurations but must enable the deduplication and compression functions.

Table 3-3 Storage configuration example

Item	LUN for Data	LUN for Online Logs	LUN for Arbitration	LUN for Archive Logs
LUN	8 x 200 GB thin LUN	8 x 200 GB thin LUN	3 x 5 GB thin LUN	8 x 200 GB thin LUN
Compression	Enabled	Enabled	Enabled	Enabled
Deduplication	Disabled	Disabled	Disabled	Disabled

3.2 Verification Procedure

3.2.1 Environment Deployment

Table 3-4 Environment deployment procedure

Task	Step	Description
1		**Prepare the physical environment (hardware and networking).**
2		**Install and configure the operating system.**
	2.1	Install Red Hat Enterprise Linux 6.5 Basic Server on four RH2288 servers.
	2.2	Configure the IP address (**/etc/sysconfig/network-scripts/ifcfg-eth0 & ifcfg-eth1**).
	2.3	Configure the host name (**/etc/sysconfig/network**).
	2.4	Configure the host name resolution (**/etc/hosts**).
	2.5	Configure I/O scheduling algorithms and TPH (**/boot/grub/menu.lst**).
	2.6	Create the user and user groups (user **oracle**, group **dba** and **oinstall**).
	2.7	Create a directory for installing target databases (**$ORACLE_BASE, $GRID_HOME**).
	2.8	Configure SSH mutual trust (**/home/oracle/.ssh/authenticated_keys**).
	2.9	Configure environment variables (**/home/oracle/.bash_profile**).

	2.10	Configure kernel parameters (**/etc/sysctl.conf**).
	2.11	Configure user limits (**/etc/security/limits**).
	2.12	Configure firewalls (**/etc/selinux/config** and **iptables**).
	2.13	Configure the NTP time service (**/etc/sysconfig/ntpd**).
	2.14	Configure the software installation source (**/etc/yum.repos.d/rhel-source.repo**).
	2.15	Install necessary software (**yum install …**) and the Oracle JRE.
	2.16	Install OceanStor UltraPath and reboot the system.
3		**Configure a storage subsystem.**
	3.1	Import licenses for storage configuration.
	3.2	Configure storage resources (disk domains, storage pools, LUNs, and LUN groups).
	3.3	Create hosts, host groups, and mapping views.
	3.4	Scan for LUNs on the hosts.
	3.5	Configure a UDEV policy (**/etc/udev/rules.d/99-huawei-devices.rules**).
	3.6	Set the I/O policy of the block device to **noop** (**/sys/block/sd*/queue/scheduler**).
4		**Install and configure an Oracle cluster and database.**
	4.1	Upload and decompress Oracle Grid 11.2.0.4 and Oracle Database software packages.
	4.2	Install and configure cluster software (**grid/runInstaller**).
	4.3	Install database software (**database/runInstaller**).
5		**Create a test database.**
	5.1	Create an ASM disk group (**asmca**).
	5.2	Create a database (**dbca**).
	5.3	Generate test data.
	5.4	Disable the compression function for Oracle applications.
6		**Use SLOB to test.**
	6.1	Check data configurations.
	6.2	Create a test tablespace.
	6.3	Use SLOB to generate test data.
	6.4	Conduct SLOB tests at varying read:write ratios.

For details about deployment operations, see relevant Huawei user guides or Oracle installation and configuration guides.

3.2.2 Data Compression Testing

Testing methodology:

1. Verify the compression ratio of newly generated data.
2. Obtain the compression ratio (50% update ratio) in a long-duration SLOB test to check whether the compression ratio changes.

Step 1 Check data configurations.

Check and configure database instance parameters by following instructions in Oracle documents to ensure that database configurations meet the following requirements:

- The instance has 10 groups of online logs. Each group contains two 128 MB log files that are stored in **+DATA**.
- Temporary tablespace exists with sufficient free space.
- UNDO tablespace exists with sufficient free space.

Table 3-5 Instance parameter settings

Parameter	Value
db_create_file_dest	+DATA
db_block_size	8192
db_file_multiblock_read_count	128
fast_start_mttr_target	30
db_writer_processes	2
db_cache_size	256M
use_large_pages	TRUE

Step 2 Create a test tablespace.

On SLOB, run the **ts.sql** script under the **misc** file to create a test tablespace named **IOPS**.

```
[root@node1 ~]# su - oracle
[oracle@node1 ~]$ sqlplus / as sysdba

SQL*Plus: Release 11.2.0.4.0 Production on Tue Mar 7 15:19:19 2017

Copyright (c) 1982, 2013, Oracle.  All rights reserved.


Connected to:
Oracle Database 11g Enterprise Edition Release 11.2.0.4.0 - 64bit Production
With the Partitioning, Real Application Clusters and Automatic Storage Management options

SQL> @/home/SLOB/misc/ts.sql
SQL> set timing on
SQL> drop tablespace IOPS including contents and datafiles;
drop tablespace IOPS including contents and datafiles
*
ERROR at line 1:
ORA-00959: tablespace 'IOPS' does not exist


Elapsed: 00:00:00.01
SQL>
SQL> create BIGFILE tablespace IOPS datafile size 1G
  2  NOLOGGING ONLINE PERMANENT EXTENT MANAGEMENT LOCAL AUTOALLOCATE SEGMENT SPACE MANAGEMENT AUTO ;

Tablespace created.

Elapsed: 00:00:01.62
SQL>
SQL> alter tablespace IOPS autoextend on next 200m maxsize unlimited;

Tablespace altered.

Elapsed: 00:00:00.00
SQL> exit;
Disconnected from Oracle Database 11g Enterprise Edition Release 11.2.0.4.0 - 64bit Production
With the Partitioning, Real Application Clusters and Automatic Storage Management options
[oracle@node1 ~]$
```

Step 3 Load 1 TB of test data.

Set data loading parameters in the **slob.conf** file and run the **setup.sh** script to load the SLOB test data into the tablespace. After the tablespace is loaded with the test data, record the data compression ratio on the storage.

```
[oracle@node1 SLOB]$ vi slob.conf
[oracle@node1 SLOB]$ sh ./setup.sh IOPS 128
NOTIFY  : 2017.03.07-15:26:50 :
NOTIFY  : 2017.03.07-15:26:51 : Begin SLOB setup.
NOTIFY  : 2017.03.07-15:26:51 : Load parameters from slob.conf:

SCALE: 1000000 (1000000 blocks)
LOAD_PARALLEL_DEGREE: 16
ADMIN_SQLNET_SERVICE: ""
SQLNET_SERVICE_BASE: ""

Connect strings to be used:
ADMIN_CONNECT_STRING: "/ as sysdba"
NON_ADMIN_CONNECT_STRING: " "

NOTIFY  : 2017.03.07-15:26:51 : Testing connectivity to the instance to validate slob.conf settings
NOTIFY  : 2017.03.07-15:26:51 : Testing Admin connect using "/ as sysdba"
NOTIFY  : 2017.03.07-15:26:51 : Successful test connection: "sqlplus -L / as sysdba"
NOTIFY  : 2017.03.07-15:26:51 : Dropping prior SLOB schemas. This may take a while if there is a large number of old schemas.
NOTIFY  : 2017.03.07-15:27:04 : Previous SLOB schemas have been removed
NOTIFY  : 2017.03.07-15:27:04 : Preparing to load 128 schema(s) into tablespace: IOPS
NOTIFY  : 2017.03.07-15:27:04 : Loading user1 schema
NOTIFY  : 2017.03.07-15:28:33 : Finished loading, indexing and gathering statistics on user1 schema in 89 seconds
NOTIFY  : 2017.03.07-15:28:33 : Commencing multiple, concurrent schema creation and loading
NOTIFY  : 2017.03.07-15:28:35 : Waiting for background batch 1. Loading up to user17
```

Table 3-6 Parameter settings for data loading

Parameter	Value
SCALE	1000000
LOAD_PARALLE_DEGREE	16

Figure 3-2 Initial data compression

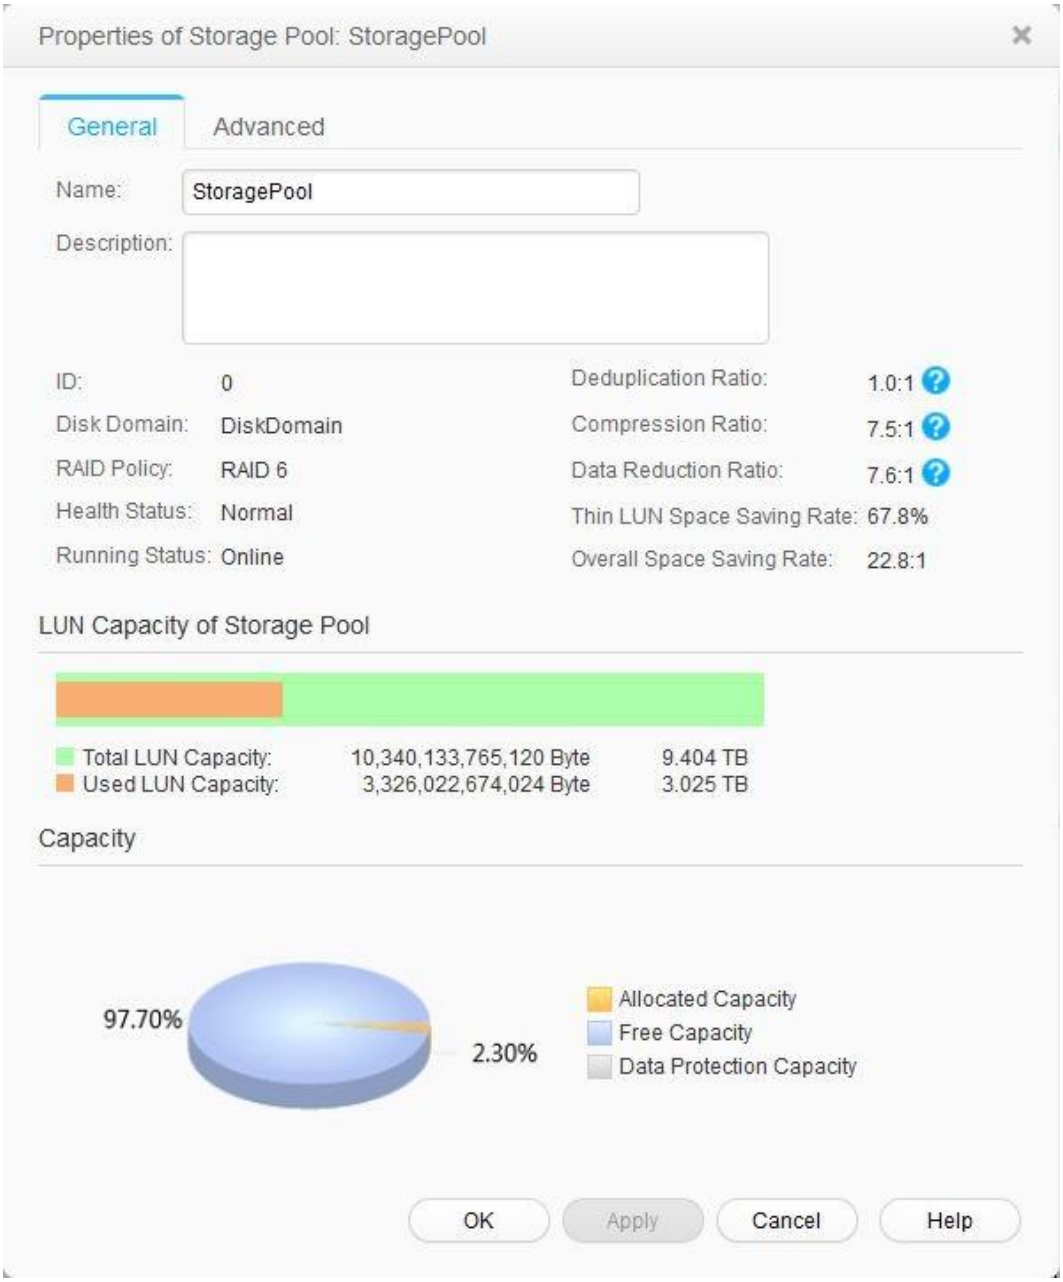

Step 4 Conduct a 48-hour load test.

After setting load testing parameters in the **slob.conf** file, run the **runit.sh** script to conduct a long-duration load test. During the test, observe the compression ratio on storage.

```
[oracle@node1 SLOB]$ vi slob.conf
[oracle@node1 SLOB]$ sh ./runit.sh 128
NOTIFY|2017-03-07 16:27:07|
NOTIFY|2017-03-07 16:27:07|Conducting SLOB pre-test checks.

UPDATE_PCT: 50
RUN_TIME: 172800
WORK_LOOP: 0
SCALE: 1000000 (1000000 blocks)
WORK_UNIT: 64
REDO_STRESS: LITE
HOT_SCHEMA_FREQUENCY: 0
DO_HOTSPOT: FALSE
HOTSPOT_MB: 8
HOTSPOT_OFFSET_MB: 16
HOTSPOT_FREQUENCY: 3
THINK_TM_FREQUENCY: 0
THINK_TM_MIN: 0.2
THINK_TM_MAX: 0.3

THREADS_PER_SCHEMA: 1

ADMIN_SQLNET_SERVICE: ""
SQLNET_SERVICE_BASE: ""
SQLNET_SERVICE_MAX: ""

Connect strings to be used:
admin_connect_string: "/ as sysdba"
non_admin_connect_string: ""
admin_conn: "sqlplus -L / as sysdba"


NOTIFY|2017-03-07 16:27:09|Testing SYSDBA connectivity to the instance to validate slob.conf settings.
NOTIFY|2017-03-07 16:27:09|Testing connectivity. Command: "sqlplus -L / as sysdba"
NOTIFY|2017-03-07 16:27:09|Testing connectivity. Command: "sqlplus -L user1/user1"
NOTIFY|2017-03-07 16:27:09|Testing connectivity. Command: "sqlplus -L user128/user128"
NOTIFY|2017-03-07 16:27:10|Performing redo log switch.
NOTIFY|2017-03-07 16:27:16|Redo log switch complete.
NOTIFY|2017-03-07 16:27:16|Setting up trigger mechanism.
NOTIFY|2017-03-07 16:27:26|Running iostat, vmstat and mpstat on current host--in background.
NOTIFY|2017-03-07 16:27:26|Connecting 1 sessions to 128 schema(s) ...
NOTIFY|2017-03-07 16:27:30|
NOTIFY|2017-03-07 16:27:30|Pausing for 5 seconds before triggering the test.
NOTIFY|2017-03-07 16:27:35|Executing AWR "before snap" procedure. Connect string is "sqlplus -S -L / as sysdba"
NOTIFY|2017-03-07 16:27:42|
NOTIFY|2017-03-07 16:27:42|Triggering the test.
NOTIFY|2017-03-07 16:27:42|List of monitored sqlplus PIDs written to /tmp/.25296_slob_pids.out
NOTIFY|2017-03-07 16:27:47|Waiting for 172792 seconds before monitoring running processes (for exit).
```

Table 3-7 Load test parameter settings

Parameter	Value	Description
SCALE	1000000	
RUN_TIME	172800	
UPDATE_PCT	50	
THREADS_PER_SCHEMA	1	
WORK_UNIT	64	
DO_HOESPOT	FALSE	
THINK_TIME_FREQUENCY	0	

Figure 3-3 Long-duration data compression

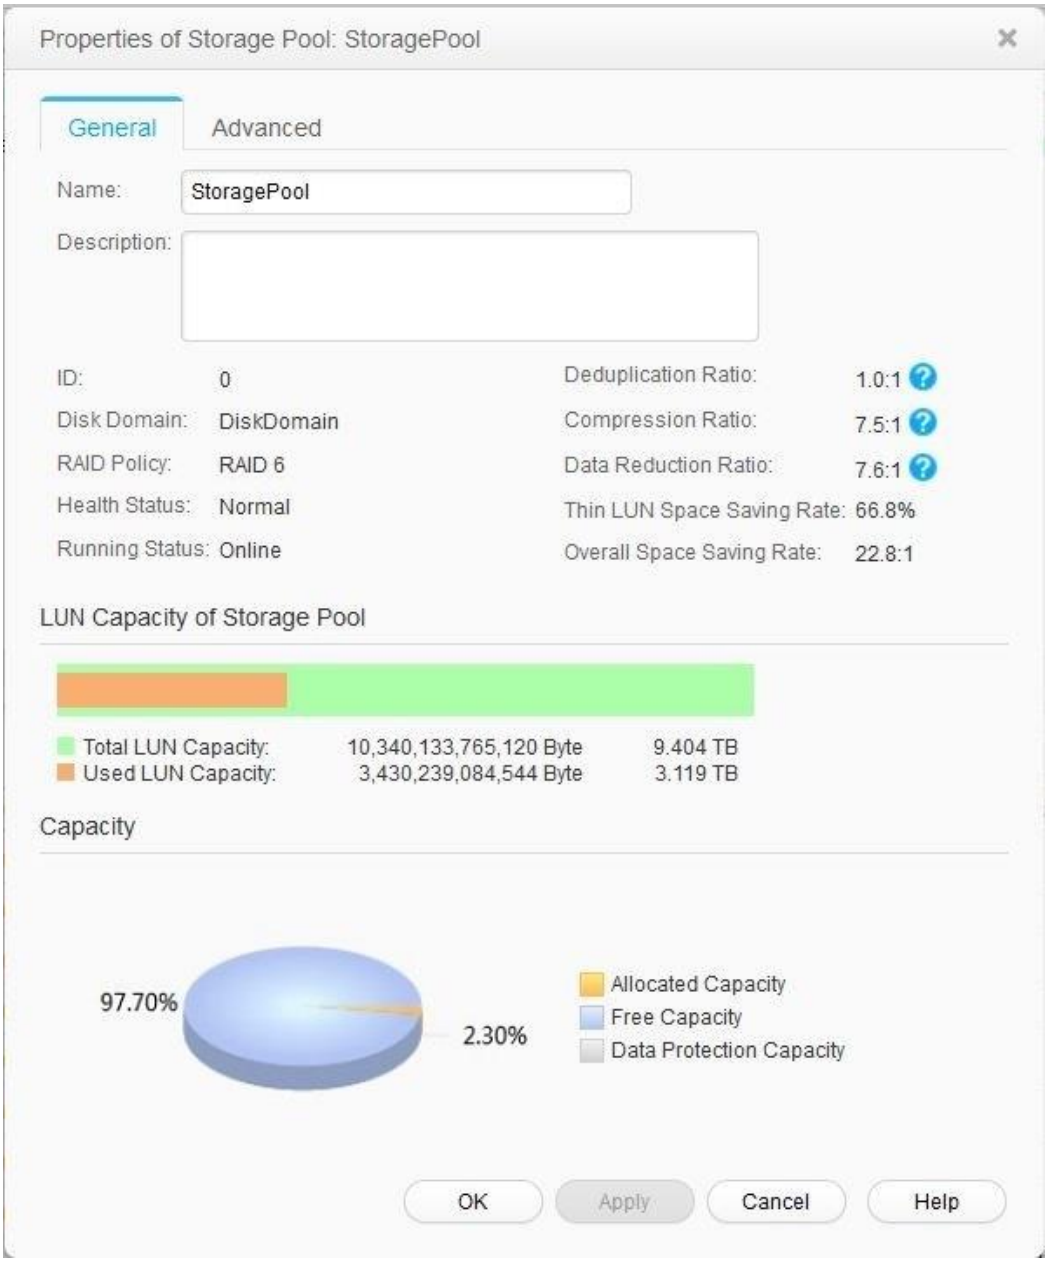

Conclusion:

When the SLOB testing tool is used in database environments, the initial and long-duration compression ratios are both 7.6, indicating that the compression ratio is almost not affected by the long-duration (48 hours) load test.

----End

4 Appendix

4.1 Reference

HUAWEI OceanStor V3 Converged Storage Systems SmartDedupe and SmartCompression Technical White Paper

4.2 Terminology

Table 4-1 Terminology

Acronym and Abbreviation	Full Name
RAC	Real Application Cluster
ASM	Automatic Storage Management
OLTP	online transaction processing
OLAP	online analytical processing

Notice

The purchased products, services and features are stipulated by the contract made between Huawei and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Huawei Technologies Co., Ltd.

Address: Huawei Industrial Base
Bantian, Longgang
Shenzhen 518129
People's Republic of China

Website: http://www.huawei.com

Email: support@huawei.com