Using SmartDedupe&SmartCompression of Huawei Dorado6000 V3 Storage in VMware VSI Environments

Virtualization brings data centers to a revolutionary shift. Enterprise data centers have been gradually changed from traditional data centers to virtual ones with increased quantity. In VMware virtual server infrastructure (VSI) environments, the SmartDedupe&SmartCompression feature of Huawei Dorado6000 V3 storage systems can be used to improve resource utilization and operational efficiency.

Jian Yi
Storage Solution Verification Dept, IT
2017-03-30
Issue 1.0

Huawei Technologies Co., Ltd.

Contents

1 About This Document

1.1 Overview

The SmartDedupe&SmartCompression feature of OceanStor Dorado6000 V3 all flash storage systems reduces the space occupied by redundant data, thereby improving the efficiency of the storage systems in transmission, processing, and storage. This document describes how to verify the performance of the SmartDedupe&SmartCompression feature in VMware VSI environments, providing reference for Huawei partners and users.

1.2 Intended Audience

- Huawei employees
- Partners
- Customers

1.3 Customer Benefits

The SmartDedupe&SmartCompression feature of OceanStor Dorado6000 V3 all flash storage systems efficiently reduces redundant data and its occupied space.

1.4 Key Components

Hardware:

- OceanStor Dorado6000 V3 all flash storage systems V300R100C00
- Huawei RH2288 V2 server

Application software and test tools:

- Virtualization software: VMware vSphere 6.0
- Multipath software: Huawei UltraPath 8.06.063
- Database software: Oracle Database 11.2.0.4
- Mail software: Microsoft Exchange Server 2013

- Web 2.0-based load test software
- Silly Little Oracle Benchmark (SLOB)
- Microsoft Exchange Load Generator 2013

Application operation systems:

- Red Hat Enterprise Linux Server 6.5
- Windows Server 2012 R2 x64
- SUSE Linux Enterprise Server 11SP2

1.5 Load Models

1.5.1 Database Services

SLOB, a test tool widely used in the industry, is used to simulate the service load model of Oracle online transaction processing (OLTP). The service load model at the I/O layer is a typical OLTP service model, where small data blocks are accessed at random and the read/write ratio is adjustable.

1.5.2 Mail Services

Exchange Load Generator 2013, a test tool provided by Microsoft, is used to simulate Microsoft Office Outlook 2003/2007 (Online Mode and Cached Exchange Mode) and activities of Outlook Web App clients. The test tool can generate load by using one or multiple protocols to simulate mail services.

1.5.3 Web/App Services

An Apache Web 2.0-based application and load generator is used to simulate the service model for virtual servers of web/app services. These virtual servers need service logics that can process ongoing programs and enormous requests from users. In addition, the virtual servers need to transfer dynamic resources from databases or file systems to clients. This service model is a typical web/app service model.

2 Application Scenarios

2.1 Server Virtualization

With server virtualization, applications can be deployed on virtual machines (VMs), instead of being separately deployed on a large number of physical servers. This maximizes the usage of server resources, simplifies management and maintenance, and improves scalability and reliability, bringing a significant IT revolution. According to the statistics of a research organization, more and more applications are deployed on virtual platforms. With explosively increasing services, enterprises' service data requires more storage. Due to the rapid growth of data, enterprises may be faced with the following challenges:

- Service data has a large amount of redundant data, which leads to cost increase in purchasing storage devices.
- More storage devices are deployed, complicating management and maintenance. Furthermore, a large sum of money is required for constructing equipment rooms and deploying power and cooling systems to ensure that the storage devices can work properly.

The value of redundant data is far from compensating the cost of maintaining storage devices. Therefore, enterprises are concerned about how to reduce redundant data as well as its occupied storage space without affecting the integrity and accuracy of service data.

To overcome this concern, OceanStor Dorado V3 all flash storage systems provide the SmartDedupe&SmartCompression feature. This feature reduces the storage space occupied by redundant data and maximizes the usage of storage resources, thereby improving the efficiency of the storage systems in transmission, processing, and storage.

2.2 Service Scenarios

The SmartDedupe&SmartCompression feature has more benefits in virtualization application scenarios. In virtualization application scenarios, users store massive VM files in storage devices. Generally, these files have the same operating system, generating a large amount of duplicate data. Moreover, duplicate user data exists.

Huawei OceanStor Dorado6000 V3 all flash storage systems can be interconnected with virtualization platforms of multiple vendors. Users can create various virtualization applications on the virtualization platforms as needed. Table 2-1 lists common virtualization platforms and some virtualization applications, and their generic versions that meet the compatibility requirements are supported.

Table 2-1 Common service scenarios

Virtualization Platform	VSI Application	Storage System
VMware ESXi Huawei FusionCompute Microsoft Hyper-V Citrix XenServer	• Web/App server • Oracle Database • Microsoft SQL Server • Microsoft Exchange Server • File sharing application • ERP application	Huawei OceanStor Dorado V3 all flash storage systems

2.3 Precautions

To fully use this feature without affecting performance, note the following when configuring this feature:

- You are not advised to enable compression at the application layer for virtual servers running database services. You need to place the system disks and data disks of VMs in different LUNs. For LUNs where system disks reside, only deduplication is enabled. For LUNs where data disks reside, only compression is enabled.
- Huawei OceanStor Dorado V3 all flash storage systems adopt global deduplication. That is to say, LUN quantity does not directly affect feature functionality. You can configure multiple LUNs to improve feature performance.
- For virtual servers running mail and web/app services, determine the feature configuration based on actual application and performance requirements.

3 Feature Configuration in VMware VSI Environments

3.1 Introduction to VMware VSI

Based on resource requirements, VMs can be divided into compute-intensive VMs, network-intensive VMs, and storage-intensive VMs. The three types of VMs have different storage demands on capacity, IOPS/GB, and response time. Therefore, the storage solution and feature configuration need to be customized to maximize resource usage. For enterprises, database systems are the foundation of their resource management systems and office collaboration systems. The key services of virtualization include the following three types: database services, mail services, and web/app services.

3.1.1 Database Services

VMs for database services are compute-intensive and storage-intensive. Database services carry enterprises' key information assets, so high performance and availability requirements are raised for the services. With the development of virtualization technologies, traditional database services are deployed on virtualization platforms, implementing server virtualization. This greatly reduces production and maintenance costs. Common databases include Oracle and Microsoft SQL Server databases. When deploying database VMs, consider assigning a storage pool with a high SLA level (high IOPS/GB and short I/O response time) for mission-critical applications to ensure their performance.

3.1.2 Mail Services

VMs for mail services are compute-intensive and network-intensive. For Exchange 2013, the mailbox server directly communicates with the Active Directory, client access server, and Microsoft Outlook client in the following manner:

The mailbox server obtains information about the recipient, server, and organization from the Active Directory using Lightweight Directory Access Protocol (LDAP). The client access server provides identity verification, proxy, and limited redirect services and supports the following common client access protocols: HTTP, POP, IMAP, and SMTP. The client access server is a stateless thin server, so it does not render data at all. Moreover, I/Os are never stored or queue up on the client access server.

3.1.3 Web/App Services

VMs for web/app services are network-intensive and storage-intensive. Web/App servers are running on these VMs to provide users with client access and service logic processing functions. After a web/app server receives an HTTP or TCP/IP request from a user, it starts running programs such as the web, script, and Java to process service logics. Meanwhile, the server reads local static files, or obtains dynamic data/files from or updates dynamic data/files in the back-end database server. Then, the server transfers the data to the client. On the web/app server, the size of a program or a static file ranges from dozens of megabytes to several gigabytes. During the frequent requests from users, the program and static file are generally cached to the server memory, so they are not read from disks. I/Os generated by this type of VMs come from program logs and OS memory exchange. Their requirements on I/O throughput and response time are not so stringent.

3.2 Solution

3.2.1 Architecture

In this solution, four servers and two 16 Gbit/s Fibre Channel switches are used for testing in VMware VSI environments.

The purpose of this test is to verify that the SmartDedupe&SmartCompression feature of Huawei OceanStor Dorado6000 V3 all flash storage systems can reduce the storage space occupied by redundant data.

Figure 3-1 Solution architecture

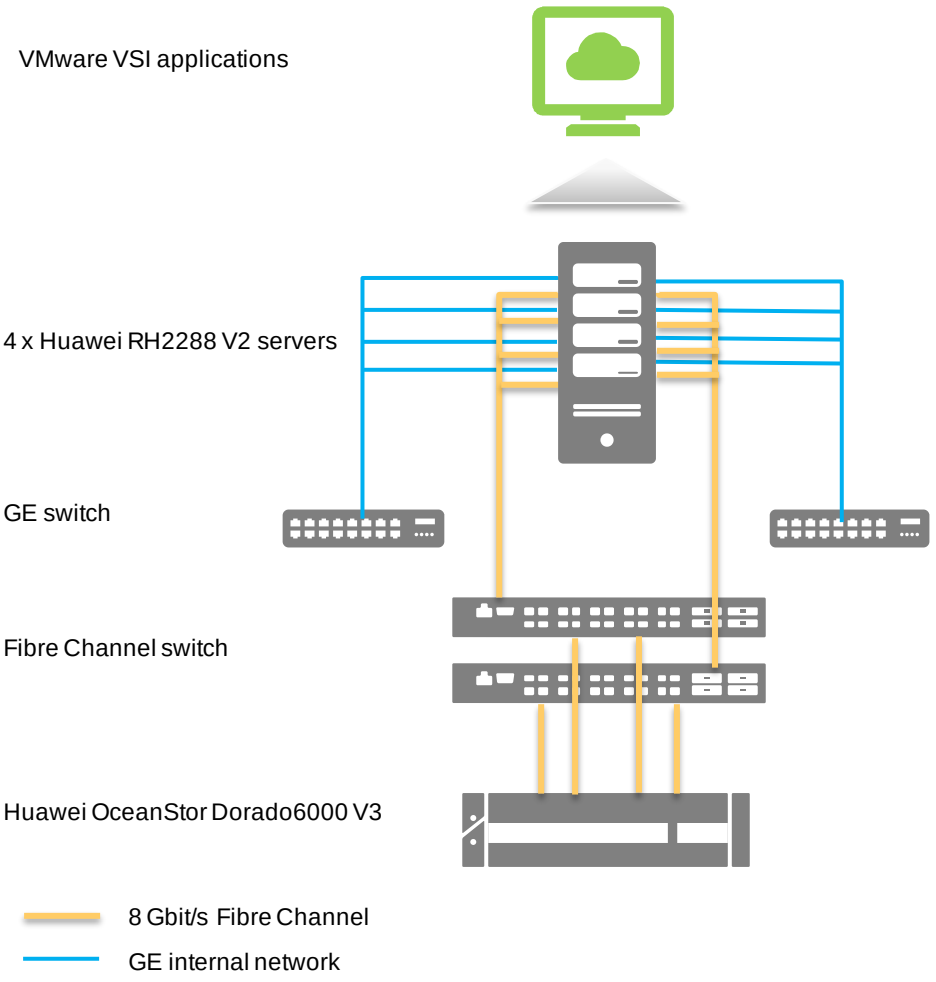

3.2.2 Hardware Configuration

Table 3-1 provides the hardware configuration in this test.

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 8 Gbit/s Fibre Channel dual-port HBAs 2 x Intel 1 Gbit/s Ethernet HBAs	4

Storage device	OceanStor Dorado6000 V3: 2 x controllers 1 x 2 U 25 slots disk enclosures 25 x 900 GB SSD disks 2 x four-port SmartIO I/O modules	1
Fibre Channel switch	SNS2224	2
GE switch	S5700	2

3.2.3 Software Configuration

Figure 3-2 provides the software configuration in the test. The generic versions of software and operating systems are used.

Table 3-2 Software configuration

Environment	Item	Software
Database	OS	Red Hat Enterprise Linux Server 6.5
	Database software	Oracle Database 11.2.0.4
Mail system	OS	Windows Server 2012
	Mail tool	Microsoft Exchange Server 2013
Web/App software	OS	SUSE Linux Enterprise Server 11SP2
Storage	OceanStor Dorado6000 V3	V003R100C00
Multipath software	UltraPath	Huawei UltraPath 8.06.063

3.2.4 Storage Configuration

A disk domain is created with a high hot spare policy. In the disk domain, a storage pool is created with RAID 6 configured. To facilitate management, system disks and data disks are separately placed for both database and mail VMs. Figure 3-2 shows the capacity and quantity of LUNs for reference.

Figure 3-2 Storage configuration

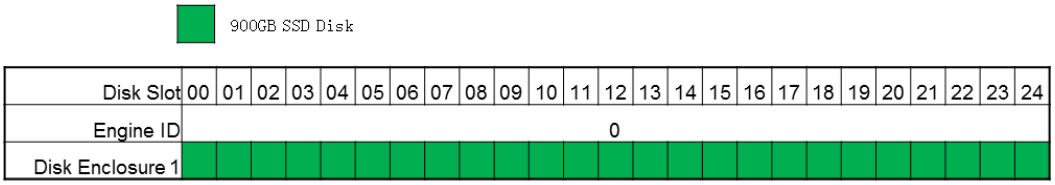

LUN	ID	SmartDedupe&SmartCompression	Capacity	Description
VSI_DB_OS_01	0	SmartDedupe	500 GB	Stores the system disks of database VMs.
VSI_DB_OS_02	1	SmartDedupe	500 GB	Stores the system disks of database VMs.
VSI_DB_DATA_01	2	SmartCompression	1 TB	Stores the data disks of database VMs.
VSI_DB_DATA_02	3	SmartCompression	1 TB	Stores the data disks of database VMs.
VSI_DB_DATA_03	4	SmartCompression	1 TB	Stores the data disks of database VMs.
VSI_DB_DATA_04	5	SmartCompression	1 TB	Stores the data disks of database VMs.
VSI_Mail_OS	6	SmartDedupe&SmartCompression	500 GB	Stores the system disks of mail VMs.
VSI_Mail_DATA_01	7	SmartDedupe&SmartCompression	7 TB	Stores the data disks of mail VMs.
VSI_Mail_DATA_02	8	SmartDedupe&SmartCompression	7 TB	Stores the data disks of mail VMs.
VSI_WEB_DB_01	9	SmartDedupe&SmartCompression	500 GB	Stores both system and data disks of

				database VMs of web/app services.
VSI_WEBServer_01	10	SmartDedupe&SmartCompression	500 GB	Stores both system and data disks of server VMs of web/app services.

3.2.5 Application Configuration

Table 3-3 Application configuration

Application	Requirement	Capacity/Type/Quantity
Database	Read/Write ratio	7:3
	Database quantity	8
Mail	Mailbox quantity	2000
	Mailbox size	2 GB
	Preset data for users	1 GB
	Number of mails sent/received per day	150
Web/App	Web server quantity	20

3.3 Verification Procedure

3.3.1 Environment Deployment

Table 3-4 Environment deployment procedure

Task	Step	Description
1		Prepare the physical environment (including hardware and the networking mode).
2		Install and configure operating systems.
	2.1	Install ESXi 6.0 on four RH2288 servers.
	2.2	Configure IP addresses.

	2.3	Install OceanStor UltraPath and restart the system.
3		Configure storage subsystems.
	3.1	Import licenses required for storage configuration.
	3.2	Configure storage resources (including disk domains, storage pools, LUNs, and LUN groups)
	3.3	Create hosts, host groups, and mapping views.
	3.4	Scan for LUNs on hosts.
	3.5	Create VMFS datastores.
4		Install and configure an Oracle database.
	4.1	Upload the Oracle 11.2.0.4 database software package.
	4.2	Install database software.
	4.3	Create a database.
	4.4	Generate data for testing.
5		Install and configure Microsoft Exchange Server.
	5.1	Configure VM storage (for example, creating directories to be mounted, storage pools, and virtual disks.)
	5.2	Install components.
	5.3	Configure the **HOSTS** file.
	5.4	Install Exchange 2013.
	5.5	Configure DAG groups and servers.
	5.6	Install Exchange Load Generator 2013.
	5.7	Initialize the mailbox to generate data for testing.
6		Install and configure a VM for web/app services.
	6.1	Deploy the web server.
	6.2	Deploy the web database and mount the file system.
	6.3	Configure the network.
	6.4	Initialize the database and file system.

Deploy the test environment based on the procedure provided in the preceding table. For details, see the related user guides or installation and configuration guides.

3.3.2 Result

Figure 3-3 shows the test result. The data reduction ratio is 5.8:1.

Figure 3-3 Test result

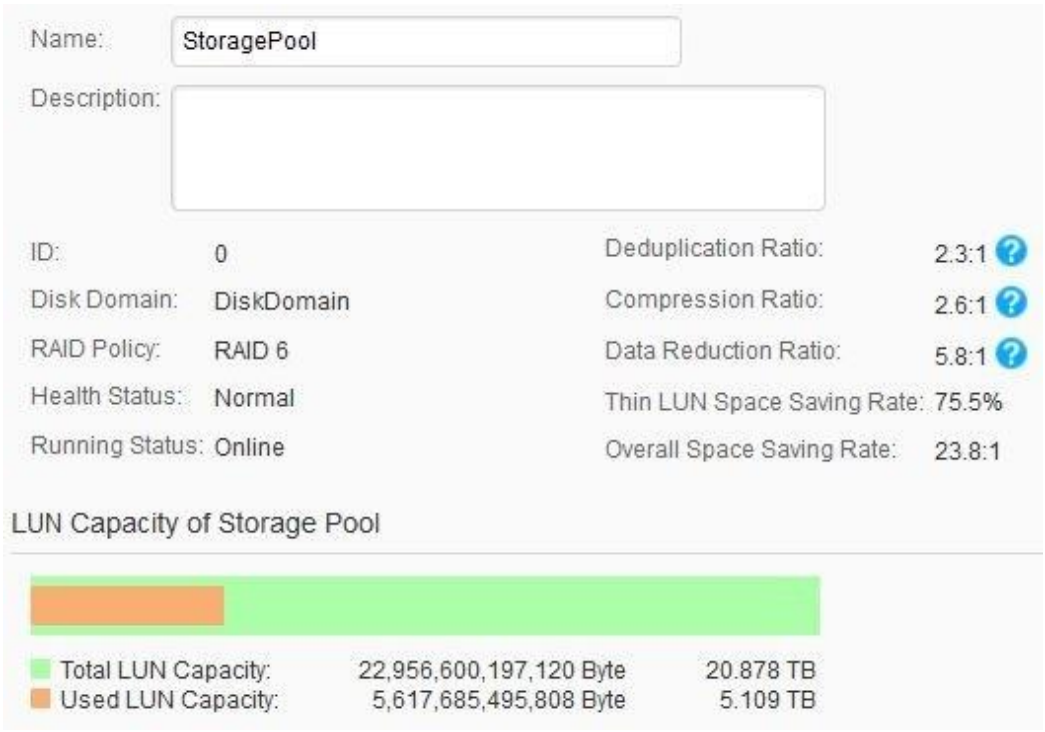

4 Appendix

4.1 Reference Documents

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

4.2 Terminology

Table 4-1 Terminology

Terminology	Description
VSI	virtual server infrastructure
SmartDedupe	Intelligent data deduplication
SmartCompression	Intelligent data compression

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