



Dell Storage Designed for Scality RING Best Practices Guide

This document contains support for the Scality solutions comprised of Dell SD630-S and the Dell SD7000-S.

Dell Engineering
March 2016

Revisions

Date	Description
March 2016	Initial release

Copyright © 2016 Dell Inc. All rights reserved. Reproduction of these materials in any manner whatsoever without the written permission of Dell Inc. is strictly forbidden. Dell and the Dell logo are trademarks of Dell Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.



Table of contents

Revisions	2
1 Introduction.....	4
1.1 Documentation resources	4
2 Tuning the Dell SD7000-S Physical Servers for performance	5
2.1 Configuring LSI RAID controller for performance	5
2.2 Using the MegaCLI RAID to configure LSI RAID controller.....	5
2.2.1 Patrol Read feature settings	5
3 Configuring Connectors.....	6
3.1 Network tuning for 10 Gb Ethernet.....	6
3.2 Adding Connectors for increased performance	6
4 Maintaining the Scalify RING	7
4.1 Detecting disk errors and replacing disks	7
4.1.1 To replace a disk.....	7
4.1.2 To replace an SSD	7
4.1.3 To change a disk as soon as failure is detected	7
4.2 Backing up and restoring configuration files	8
4.3 Adding Scalify RING capacity.....	8
4.3.1 To add a new disk to expand storage.....	8
4.3.2 To add Physical Servers to increase Scalify RING capacity.....	9
4.4 Joining nodes to the Scalify RING.....	9
4.5 Removing storage nodes.....	9
4.6 Updating firmware and patches.....	10
4.7 Upgrading Scalify software	10
4.7.1 Sites located in private networks	10
5 Getting Help	11
5.1 Contacting Dell.....	11
5.2 Locating Service Tag.....	11
5.3 Quick Resource Locator	11
5.4 Contacting Scalify	12



1 Introduction

This document is a collection of the fundamental best practices for setting up and maintaining a Scality RING solution based on Dell PowerEdge servers and Dell Storage enclosures with links to Scality and Avago Technologies web sites for obtaining more information.

The Dell Storage Designed for Scality RING solution enables deployment of an ultra-dense storage platform paired with Scality software, for object-based storage or file-based storage on the petabyte scale. Supported services include object level protocols: Scality REST; RS2/Amazon; Openstack Swift; and Cloud Data Management Interface (CDMI). Supported file level services include NFS; Server Message Block (SMB) Openstack; Openstack Cinder; and local Linux FS via FUSE. For more information see, **Dell.com/sdscalityseriesmanuals**.

1.1 Documentation resources

For a complete list of supported Dell and Scality documentation, see the *Support Matrix* for this release.

Scality provides you with login credentials to access the Scality documentation on purchase of Scality RING software.

NOTE: If you are not provided with login credentials, contact your Scality Account Representative.

These credentials are required to access Scality RING documentation, which is delivered in a version specific WIKI format.

For information about top-level documentation, go to <http://docs.scality.com>.

You can view the documents on the Scality documentation page based on the selected software version. The latest version is selected by default. For the most up to date supported version of software, see the *Support Matrix*.

To view an earlier software version, select the version number from the drop-down list, and then click **Go**.



2 Tuning the Dell SD7000-S Physical Servers for performance

Scality offers tips for tuning the Dell SD7000-S, which you can access from the Scality documentation web site, see <http://docs.scality.com>. These topics are listed in the right pane, under the section labeled **Administration**:

- **Partitioning and File system Tuning Tips**
- **Network Tuning for 10 Gb Ethernet**

2.1 Configuring LSI RAID controller for performance

The Avago Technologies portal provides drivers, tools, and documentation for LSA RAID controllers. To access the links listed in this section, you must log in to the Avago Technologies web site. Register for the Avago Technologies site at <http://www.avagotech.com/> and obtain a username and password.

2.2 Using the MegaCLI RAID to configure LSI RAID controller

The MegaRAID tool and supporting documentation is available from the Avago Technologies portal.

http://www.avagotech.com/docs-and-downloads/raid-controllers/raid-controllers-common-files/8-07-14_MegaCLI.zip

2.2.1 Patrol Read feature settings

Scality recommends that you disable the Patrol Read feature on the LSI RAID controllers. The Patrol Read feature runs periodically in the background checking for physical media errors. Patrol Read can kick off a full consistency check if any errors are found, which can negatively impacting performance. Use Patrol Read to supplement RING software error detection, which is currently available only on *object* read operations.

Scality recommends that you only enable the Patrol Read feature if you require error detection outside of read operations, such as for disk scrubbing operations.

For more information, on the Scality document web site, in the left pane select **Operations > How to > RAID How To**. In addition, documentation is available through the Avago Technologies portal.



3 Configuring Connectors

Scality documentation offers tips on configuring and tuning the Connector.

On the Scality documentation web site, in the left pane, select **Overview > RING Connectors** for detailed information on the available Connector types.

Information about tuning Connectors is found on the Scality web site, in the left pane, select **Operations > Best Practices > Tuning Tips**. Under the section **OS Tuning Guide**, click **Changing the memory allocator for Ring connectors**.

3.1 Network tuning for 10 Gb Ethernet

When using 10 Gb Ethernet adapters to link Connectors and Physical Servers, the default configuration settings of the Linux TCP stack and of the network adapter itself do not provide the best performance. To optimize performance, tune the settings as described in Scality documentation. In the left pane, select **Operations > Best Practices > Tuning Tips > Network tuning for 10 Gb Ethernet**.

3.2 Adding Connectors for increased performance

Connectors host the processes for data access (NFS, System Message Block (SMB), and Object,). Dell recommends a minimum of two Connectors and that you add more as performance requirements of the application load increase.

Connector types for file server solutions are implemented on a Dell SD630-S server. However, some Connector types for object servers, such as `sproxyd`, are instead installed directly onto the Physical Servers. See the *Deployment Guide* for more information on Connector implementation for file and object servers.

Note: When adding storage capacity, you may need to scale-up the number of Connectors to match the increased capacity.



4 Maintaining the Scality RING

To locate useful scripts for maintaining the Scality RING, go to <http://docs.scality.com>. In the left pane, go to **Operations > Scripts**. Find each script mentioned in this section listed under **Scripts** on the Scality documentation page.

4.1 Detecting disk errors and replacing disks

You can use the following disk management scripts (tools) to add new disks in empty slots (whether to increase data storage capacity or to serve as a non-volatile path (NVP) containers for object metadata) or to replace disks when the repair helper returns a REPAIR_FAILED exit code.

In the left pane, browse to **Operations > Tools**, and then in the right pane click **Disk Management Tools**. This path includes links and instructions for adding and replacing failed disk.

In the left pane, find more disk management procedures under **Operations > How to > Disk How To**. This page includes the following sections:

- **How and when to change a failed disk**
- **Check disk usage**
- **Check system logs for disk errors**
- **Restore `bizobj.bin` on a faulty SSD**

4.1.1 To replace a disk

Use the tool `scaldiskreplace` to replace a failed disk. Download the script from the Scality documentation web site.

On the Scality web site, in the left pane, select **Operations > Monitoring > Disk Management Interface**.

Note: Although you can use `scaldiskreplace` to replace a healthy disk with a disk that provides increased storage capacity or superior performance, this is not encouraged. Disk deactivation is not automatic for replacement of a healthy disk. The `scaldiskreplace` tool does not copy anything. When replacing a healthy disk, copy the disk data manually. Manual intervention by an operator is required to deactivate the disk (stop the `biziod`), delete the configured `nvp` directory (if any), and unmount (`umount`) and remove the disk being replaced.

4.1.2 To replace an SSD

Plan to replace SSD before it fails. Use this task to be proactive with SSD replacement.

On the Scality web site, in the left pane, browse to **Operations > How to > Node and `biziod` How to** and click **Planned replacement of an SSD**.

4.1.3 To change a disk as soon as failure is detected

The Scality web site contains information and tips about changing a disk as soon as a failure is detected.



On the Scality web site, in the left pane, select **Customer Support FAQ**. In the right pane, select **Logs and Errors > A disk failed! What do I do?** A procedure named **Changethediskassoonasafailureisdetected** is located here.

Useful information about disk failure is listed in the **Customer FAQ > Uncategorized > Why is our whole platform slowing down after a single disk failure**.

4.2 Backing up and restoring configuration files

The Scality documentation web site offers information on backing up and restoring a Physical Server, Connector, and Supervisor configuration files.

On the Scality web site, in the left pane, select **Install and configure > Configuration > Backup configuration files**.

The `scality-node`, `scality-rest-connector`, and `scality-supervisor` packages each install a script and a daily cron job, which performs the backup of important configuration files. The backup files are stored under the following directories:

```
/var/lib/scality-node/backup  
/var/lib/scality-rest-connector/backup  
/var/lib/scality-supervisor/backup
```

Use backup files to reinstall the packages in case the contents of the `/etc` directory are lost or damaged. If the backup files are copied to another server, use them to install the packages from scratch.

4.3 Adding Scality RING capacity

Expanding capacity to the Scality RING should occur when your Scality RING installation has reached 60 to 65 percent total capacity. This ensures that the expansion starts prior to the RING reaching 80 to 85 percent total capacity.

Monitor capacity usage and plan to grow the RING when it has reached 80 percent capacity. Scality recommends not allowing the storage pool to fill up above 90 percent before expansion.

Warning: Although the system continues to run after the storage pool fills above 90 percent, writing new data becomes an issue as the system approaches full capacity.

Consult your certified Scality sales team for detailed capacity planning.

4.3.1 To add a new disk to expand storage

Expand storage by adding disks to empty slots in a Dell SD7000-S. Disks should be added evenly to all Physical Servers, so that capacity is evenly distributed around the Scality RING. While adding drives, ensure to maintain a raw HDD to SSD capacity ratio greater than 0.6 percent.



Example: A Scality RING that consists of one 800 GB SSD and fifteen 8 TB HDDs, has a raw storage capacity ratio of $.8\text{TB}/120\text{TB} = 0.666$ percent per Physical Server. Consult your certified Scality sales team for detailed capacity planning.

Use the tool `scaldiskadd` to add disks to the available capacity. Download the script from the Scality documentation web site, in the left pane, select **Operations > Monitoring > Disk Management Interface**.

Note: Drives and SSDs need to be a RAID 0 and mounted in the Linux system prior to being added using the `scaldiskadd` command.

4.3.2 To add Physical Servers to increase Scality RING capacity

Increase capacity by adding additional Dell SD7000-S servers to the Scality RING. Each Dell SD7000-S provides two Physical Servers, where each is connected to 45 drive slots available for HDD and SSD. Only add additional Physical Servers when all Physical Servers are at full capacity. When existing servers reach full capacity, you can add one or more full capacity SD7000-S to the Scality RING.

Consult your certified Scality sales team for detailed capacity planning.

For information on adding Physical Servers see the Scality documentation web site, in the right pane, select **Overview > What is Scality Ring**. Additional configuration information is located here. In the left pane, select **Install and Configure > Configuration > Configuring a RING**, or in the left pane, select **Operations > How to > Ring How To**.

4.4 Joining nodes to the Scality RING

Automatically configure and add virtual servers (nodes) to Scality RING storage. For more information, see the Scality documentation web site, in the left pane select **Operations > Scripts > Join Nodes to Ring (Scality-auto-join.sh)**.

4.5 Removing storage nodes

To retire a Physical Server, you must shut down the server. Scality recommends that prior to physically removing the server, from the Supervisor Web UI, select the **Leave** command for each of its nodes and then the **Stop** command for the Physical Server.

Note: Nodes are automatically forced to leave the Scality RING if all the disks to which they have access are marked as full or out of service. This allows the Scality RING to isolate problem disks and redistribute objects and replicas to other disks in the Scality RING.

For more information on the **Leave** command, see Scality documentation. In the left pane, select **Overview > Ring Supervisor > Supervisor Web UI > Supervisor Operation Window**.



4.6 Updating firmware and patches

Scality recommends that you maintain up to date firmware and patches for your Scality RING solution. The most recent patch, differentiated by operation system (OS), is available from a link on the Scality documentation web site in a section titled **Downloads**. To find this section, on the Scality documentation web site, in the left pane, select **Downloads**.

Hardware must run the most currently supported firmware before you deploy it in production.

Monitor new firmware releases for all major components in the system, as firmware bugs can negatively impact the storage system.

Note: Do not apply the patches simultaneously to every storage node in the Scality RING.

4.7 Upgrading Scality software

To maintain production traffic, you must upgrade servers one at a time, waiting for each server to be completely active again before upgrading the next server.

All components can work with each other within one major Scality RING release. Scality recommends upgrading the Supervisor server first, followed by the Physical Servers, and Connectors last. If further assistance is required, contact the Scality technical services team.

4.7.1 Sites located in private networks

Scality recommends installing the upgrade packages from a repository. You can configure the Scality repository on a system, or use a custom repository. If necessary, download a Scality ISO image and loaded as a local repository.

If further assistance is required, contact the Scality technical services team.



5 Getting Help

5.1 Contacting Dell

Dell provides several online and telephone-based support and service options. If you do not have an active internet connection, you can find contact information on your purchase invoice, packing slip, bill, or Dell product catalog. Availability varies by country and product, and some services may not be available in your area. To contact Dell for sales, technical assistance, or customer-service issues:

1. Go to Dell.com/support.
 1. Select your country from the drop-down menu on the bottom right corner of the page.
 2. For customized support:
 - a. Enter your system Service Tag in the Enter your Service Tag field.
 - b. Click Submit.
The support page that lists the various support categories is displayed.
 3. For general support:
 - a. Select your product category.
 - b. Select your product segment.
 - c. Select your product.

The support page that lists the various support categories is displayed.

5.2 Locating Service Tag

Your system is identified by a unique Express Service Code and Service Tag number. This information is used by Dell to route support calls to the appropriate service provider.



Figure 1 Locating service tag

5.3 Quick Resource Locator

Use the Quick Resource Locator (QRL) to get immediate access to system information and how-to videos. This can be done by visiting Dell.com/QRL or by using your smartphone or tablet and a model-specific Quick Resource (QR) code located in your documentation. To try out the QR code, scan the following image.



Figure 2 Quick Resource Locator

5.4 Contacting Scality

U.S.A. Toll Free: +1-800-319-7090

France: +33 -184885570

Italy: +39-0662207278

Japan: +81-3-4588-1408

U.K. & International: +44 2033 181 073