

DR Series System Appliance

## Best Practices Guide

Setup, Replication and Networking

Quest Engineering

November 2017



© 2017 Quest Software Inc.

## ALL RIGHTS RESERVED.

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND

This guide contains proprietary information protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of Quest Software Inc.

The information in this document is provided in connection with Quest Software products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of Quest Software products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, QUEST SOFTWARE ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL QUEST SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF QUEST SOFTWARE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Quest Software makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. Quest Software does not make any commitment to update the information contained in this document.

If you have any questions regarding your potential use of this material, contact:

Quest Software Inc.

Attn: LEGAL Dept

4 Polaris Way

Aliso Viejo, CA 92656

Refer to our Web site (<https://www.quest.com>) for regional and international office information.

## Patents

Quest Software is proud of our advanced technology. Patents and pending patents may apply to this product. For the most current information about applicable patents for this product, please visit our website at <https://www.quest.com/legal>.

## Trademarks

Quest, the Quest logo, and Join the Innovation are trademarks and registered trademarks of Quest Software Inc. For a complete list of Quest marks, visit <https://www.quest.com/legal/trademark-information.aspx>. Microsoft®, Windows®, Windows Server®, Internet Explorer®, MS-DOS®, Windows Vista® and Active Directory® are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Red Hat® and Red Hat® Enterprise Linux® are registered trademarks of Red Hat, Inc. in the United States and/or other countries. Novell® and SUSE® are registered trademarks of Novell Inc. in the United States and other countries. Zmanda is a trademark of Zmanda Incorporated in the USA. All other trademarks and registered trademarks are property of their respective owners.

## Legend

 **WARNING:** A WARNING icon indicates a potential for property damage, personal injury, or death

 **CAUTION:** A CAUTION icon indicates potential damage to hardware or loss of data if instructions are not followed.

 **IMPORTANT, NOTE, TIP, MOBILE, or VIDEO:** An information icon indicates supporting information.

# Contents

<b>Introduction</b> .....	<b>6</b>
Administration guides and hardware owner's manuals .....	6
Best practice guides.....	6
Case studies .....	7
<b>System setup</b> .....	<b>8</b>
Hardware.....	8
Expansion shelves .....	8
Initial out-of-the-box setup .....	8
<b>Replication setup and planning</b> .....	<b>13</b>
Replication considerations .....	13
Replicating containers.....	14
Calculate replication interval.....	14
Bandwidth optimization.....	15
Replication encryption .....	15
Addressing packet Loss .....	16
Replication tuning.....	16
High bandwidth conditions.....	16
Reducing bandwidth concerns .....	16
Domain access.....	17
<b>Networking</b> .....	<b>18</b>
Networking capabilities .....	18
Supported network cards .....	18
Network interface card bonding .....	19
Secure network separation .....	22
Scenario 1.....	22
Scenario 2.....	28

Scenario 3.....	31
Scenario 4.....	34
Scenario 5.....	37
Troubleshooting .....	39
How to achieve peak performance with the DR6000.....	40
Guidelines.....	40
How to achieve peak performance with the DR6300.....	43
Guidelines.....	43
<b>Resources .....</b>	<b>46</b>
DR Series manuals .....	46
Quest support.....	46

# Executive summary

---

This document contains some of the best practices for deploying, configuring, and maintaining a DR Series system appliance in a production environment. Following these best practices can help ensure the product is optimally configured for a given environment.



**NOTE:** This guide is not a replacement for the system administrator's guide. In most cases, this guide provides only a high level description of the problem and solution. Please refer to the *DR Series System Administrator's Guide* for the specific options and steps required to implement the recommended actions.

---

# Introduction

In addition to this document, we recommend the documents described below for the advanced user.

**i** **NOTE:** Please consider the system administrator's guide as a pre-requisite for this guide, or at least as a reference guide on how to complete the tasks listed in this guide. The *DR Series System Administrator's Guide* is included in the list below.

You can download DR Series product documentation, setup guides, best practice guides, and user guides at [support.quest.com/DR-Series](https://support.quest.com/DR-Series) by selecting your specific DR Series system model.

## Administration guides and hardware owner's manuals

- Quest DR Series System Administrator Guide
- Quest DR4100 Systems Owner's Manual
- Quest DR6000 Systems Owner's Manual
- Quest DR4300e System Owner's Manual
- Quest DR4300 System Owner's Manual
- Quest DR6300 System Owner's Manual
- Quest DR Series System Command Line Reference Guide

## Best practice guides

- Best Practices for Setting up NetVault Backup Native Virtual Tape Library (nVTL)
- Best Practices for Setting up NetVault SmartDisk
- DR Series Disk Backup Appliance Setup Guide for CommVault Simpana 10

- [Setting up EMC Networker on the Quest DR Series Disk Backup Appliance Through CIFS or VTL](#)
- [Setting up EMC Networker on the Quest Series Disk Backup Appliance Through NFS](#)
- [Setting up Veeam on the Quest DR Series Disk Backup Appliance](#)
- [Setting up vRanger on the Quest DR Series Disk Backup Appliance](#)
- [Setup Guide for Symantec Backup Exec 2010 R3](#)

## Case studies

- [Haggar Case Study](#)
- [CGD Case Study](#)
- [TAGAL Steel Case Study](#)
- [DCIG DR4X00 and EqualLogic Better Together](#)
- [ESG DR4X00 Lab Validation](#)
- [Pacific BioSciences Case Study](#)

---

# System setup

## Hardware

When the DR Series system ships to a customer's site and is powered on for the first time, the system still needs to complete a background initialization (init) of the RAID. During this background init, the system may seem sluggish or write slower than expected. This will resolve itself within 24 hours.

## Expansion shelves

The proper boot procedure for the DR appliance is to first power on the DR expansion shelf/shelves and then connect it to the DR appliance. After connected to the appliance, power on the DR appliance and ensure the expansion license(s) are available or ready to add.

## Initial out-of-the-box setup

This section outlines the various items that must be configured when the DR appliance first powered on.

## Registration

The DR Series system has the ability to notify administrators of recent updates. You can register the DR Series appliance and ask for updates. As new releases are posted on the web, update notifications will be emailed alerting the administrator of new upgrades that are available. It is recommended that administrators enable this option during registration.

## Setting up alerts

The DR Series system has built-in monitoring through Quest Open Manage, which is installed on the appliance. All critical software and hardware alerts are sent as alerts via email. All hardware related events are also sent as a trap over SNMP to the configured monitoring server.

Quest highly recommends that alerts are configured on the DR Series appliance and are sent to a global distribution list so that they can be monitored and resolved quickly.

## Password reset

The DR Series appliance has the ability to allow the administrator to reset the administrator password. To ensure security, administrators should setup password reset immediately. This ensures that in the future if the admin password is forgotten it can be securely reset.

## Joining the domain

If the DR Series system is to be joined to an Active Directory domain it is recommended that this action is performed from the start. Doing so will allow ACLs to be applied and domain users can be used to access the data from the backup application.

### Adding the DR Series system to Active Directory

Log on to the DR Series system GUI, click Active Directory in the side panel, click Join in the upper right hand corner and enter in the name of the domain and credentials to add the DR Series system to the domain.

It will now be possible to logon to the DR appliance using the GUI\CLI via Domain\user for any users that are in the global group.

To allow multiple groups to logon to the DR appliance using Active Directory do the following:

- 1 Create a new global group in Active Directory.
- 2 Add each group to be allowed to access the DR to this global group.
- 3 Add the new global group to the DR using the following command from the CLI:

```
authenticate --add --login_group "domain\group"
```

Users that are part of the selected AD group will be able to logon to the CLI and GUI to administer the device.

### Setting ACLs and inheritance

It is best to set ACLs and inheritance when the system is first setup. By default, every user has access to all data. Attempting to change permissions and inheritance later is very time consuming.

## Set time

If the DR Series system is not joined to an Active Directory domain, it is best to configure the DR Series system to use NTP. If the DR Series system is joined to an Active Directory it will automatically sync its time.

## Storage group creation

You have a capability to organize your data by creating storage groups and creating containers within those storage groups on your DR Series system. A storage group allows you to create separate storage policies, for example, for different data groups and the different capacities utilized on a single DR Series system.

Example: You have the ability to create logical storage groups so that backup data is organized by department, workgroup, customer, and so on. Storage groups define boundaries of backup data so that data from one group cannot cross into any other data. For example, engineering data can be in one or more containers within storage group A, with customer service data in containers in Storage Group B.

Refer to the following considerations when creating Storage Groups:

- Only 32 containers are allowed on a single appliance.
- Storage\_group names can only have the characters [a-z,A-Z,0-9,'\_','-'].

### Storage groups limits

- The storage group limit is 4 in the DR4000, DR4100 DR6000 and DR4300e.
- The storage group limit is 8 in the DR4300 and DR6300.

## Container creation

The DR Series appliance uses containers to store backup data. These containers are segmented folders that have individual protocols, security permissions, marker types and connection types.

Refer to the following considerations when creating containers:

- Only 32/64/128 containers allowed on a DR4000, DR4100, DR4300e/DR6000, DR4300/DR63000 appliances respectively.
- Access protocols are set at a container level (such as, NFS, CIFS, OST, RDS, NDMP, ISCSI, FC).
- Security is set at a container level (Locking down Via IP, Unix/Windows ACLs, etc.).
- Markers are set at the container level (None, Auto, CommVault, Networker, etc.).
- OST quotas are set on a container level.
- Container Names cannot contain spaces.

## Container limits

There are several different strategies in which to approach container creation. In general, it is better to have as few containers as possible to maintain ease of management. This section is designed to assist in developing an optimal container strategy based on your organization's needs.

With most configurations where replication is not required, it is common to have a single container if the administrators have a single Data Management Application (DMA). When replication is required, container creation can become more complex, so it is best to choose what containers should and should not be replicated as well as prioritizing what data should be replicated first.

Below are four scenarios to help with picking the proper container creation strategy:

- Scenario 1: Separate data to be replicated vs. data not to be replicated
- Scenario 2: Separate data that has a higher value to replicate vs. lower value
- Scenario 3: Separate different types of data into different containers
- Scenario 4: Separate different DMA types into different containers

### Scenario 1: Separate data to be replicated vs. data not to be replicated

Robert has Exchange data that is required to have 2 copies, with 1 copy maintained offsite. Robert also has VM data, which is NOT required to have 2 copies.

**Recommendation:** Robert should have the following two containers:

- Container 1. For the Exchange data so that it can be replicated each week off site.
- Container 2. For the local VM data so that it is not replicated and does not take up valuable WAN bandwidth.

### Scenario 2: Separate data that has a higher value to replicate vs. lower value

Robert also has intellectual property (IP) that he would like to have replicated offsite each day.

**Recommendation:** Robert should create a third container and enable replication schedules on all containers (giving more time to the third container) to ensure replication of the IP container competes each day.

### Scenario 3: Separate different types of data into different containers

Assume that Robert also has SQL data that bypasses the DMA and is written directly to the DR Series system.

**Recommendation:** Robert should create a fourth container to allow the container to be locked down just to that SQL server, as well as to allow independent access which does not interfere with the DMA.

### Scenario 4: Separate different DMA types into different containers

Robert also has a VM infrastructure that he wishes to protect using Quest vRanger. This is in addition to his other DMA which is used to protect their physical servers.

**Recommendation:** Create a fifth container for vRanger data. This will allow the most flexibility in terms of locking down the NFS/CIFS share for vRanger, it also allows for the most flexibility in terms of replicating the data offsite in the future.

## Access protocols

Access protocols are set on a per container basis when a container is created. In some cases these cannot be changed. The protocol options are options are as follows:

- CIFS or NFS only
- CIFS and NFS together (although cross protocol support is not supported)
- OST Only
- RDS Only
- VTL (NDMP, ISCSI, FC)

It is possible to add or remove CIFS/NFS access. However, once the container has RDS/OST added it is not changeable.

Security is set at a container level (Locking down Via IP, Unix/Windows ACLs, etc.)

For CIFS shares it is recommended that the shares are set with the most restrictive ACLs, and further locked down by the IP or DNS name of machines that are allowed to connect to that container.

For NFS shares, it is recommended that root user be set to nobody and that NFS shares are further locked down by the IP or DNS name of the machines that are allowed to connect to that container.

## Marker Support

Many data management applications (DMAs) add metadata into the backup stream to enable them to find, validate and restore data they wrote into the file. This metadata makes the data appear unique to dedupe enabled storage. In order to properly dedupe the data, the markers need to be removed before the stream is processed.

In the DR Series system, markers are set on a container level and are set to auto by default. This allows all known detected markers to be stripped before the data is processed. In situations where the DMA does not have markers, or the DMA is known, a slight performance increase can be had by setting the markers correctly.

If the DMA in use is BackupExec or Netbackup the marker should be set to none. For other DMAs such as Commvault, Networker, Arcserve, or any of the other supported markers, set the marker to the type that corresponds to the appropriate DMA.



**NOTE:** Changing the marker type after data is ingested could make the data appear unique, causing dedupe savings to be adversely impacted until all ingested data with the previous marker is removed. To avoid this problem, the proper marker should be set on containers from the start.

---

# Replication setup and planning

The DR appliance provides robust replication capabilities to provide a complete backup solution for multi-site environments. With WAN optimized replication, only unique data is transferred to reduce network traffic and improve recovery times. Replication can also be scheduled to occur during non- peak periods, and prioritizes ingest data over replication to ensure optimal backup windows. The following sections cover the various considerations and planning that should be taken into account for replication with the DR Series backup appliance. As always, the information provided below are guidelines and best practices and are meant to be supplemental to the information provided in the DR Series system administration guide.

## Replication considerations

Refer to the following important considerations before setting up replication.

- The number of concurrent replications supported depend on the platform of the appliance used. 32/64/128 on a DR4000, DR4100, DR4300e/DR6000, DR4300/DR63000 appliances respectively.
- Replicated data is already compressed and deduplicated prior to its transfer to the destination DR appliance. This results in approximately an 85%- 90% reduction in data being transferred from the source to the target device.
- Bandwidth throttling works between pairs of devices.
- Replication supports none, 128bit, 256bit and encryption options. 128-bit encryption is recommended.
- Replication uses a 10MB TCP window by default. Contact Quest support if this is needed to be adjusted higher for high latency/low bandwidth links.
- Replication can be scheduled on a per container basis.
- Container names should match on each DR Series system to simplify disaster recovery.
- Replication also replicates CIFS and NFS security bits. For CIFS shares the target DR Series system also needs to be in the same domain or forest as the source DR Series system for ACLs to be applied correctly.

# Replicating containers

The maximum amount of containers supported is 32/64/128 on a DR4000, DR4100, DR4300e/DR6000, DR4300/DR63000 appliances respectively . For optimum replication performance, it is recommended that the number of replication containers be kept at a minimum.

When planning larger deployments, the following recommendations should be considered to maintain an acceptable level of performance:

- In larger environments, it can be easy to quickly approach the container limit. A simple method to reduce the number of replicated containers is to leverage container directories.
- In some scenarios it may become necessary to replicate more than the allowed containers to a single physical site. In such a situation it is required to utilize two separate head units and fewer expansion shelves in order to provide the necessary resources for improved performance.

## Calculate replication interval

Calculating the required bandwidth for replication will assist in properly sizing the infrastructure for maximum performance.

To calculate the time required to replicate a given container the following two points of data are required:

- Identify the amount of data that will be replicated. A common method is to view current backup jobs and log files to determine the amount of data being backed up each day. The more precise this number is, the more accurate the bandwidth calculation will be. Since any data transferred during replication has already been compressed and deduplicated to roughly 85%-90% of the original size, start by multiplying the original data size by 15% to determine the amount of data to be replicated. For example, to transfer two terabytes of data, break it down into megabytes by multiplying the value by 1048576. To convert 2TB to MB the formula would be:  $2TB * 1048576 = 2097152 MB$ . Now, reduce this number by 85% and assign it to the variable:

replica\_data:

```
replica_data = 2097152 MB * .15 = 314572 MB
```

- Determine the effective bandwidth. A common method to determine bandwidth between two sites utilizes a freeware product called iPerf (<https://code.google.com/p/ipperf/>). The following steps provide the command line instructions to calculate bandwidth:
  - a Run `iperf -s -w 5M` on a server at the central site.
  - b Run `iperf -c <IP of server at central site> -w 5M`
  - b Capture the resulting bandwidth reported from the server at the central site, showing the calculated bandwidth between sites.

c Assign the acquired bandwidth value to the variable `effective_bandwidth`. This value may be obtained from other methods of your choosing, but should be specified in MB for further calculations. A bandwidth value of 10MB will be used for the following example.

- Determine the acceptable amount of time allowed for replication. Convert the allowed replication time to seconds.(i.e. 10 hours converted is  $10*60*60 = 36000$  seconds)

With this information, the following formula can be used to calculate the time required to replicate the data.

$$\text{replication\_interval} = \text{replica\_data} / \text{effective\_bandwidth}$$

Using the example of 2TB and a 10MB effective bandwidth, time can be calculated as follows:

$$\text{replication\_interval} = 314572 \text{ MB} / 10 \text{ MB} = 31457 \text{ seconds (8.7 hours)}$$

## Bandwidth optimization

When utilizing replication, it might be necessary to enforce bandwidth limitations, which can lower the impact of replication traffic on the network. For example, if the WAN link also is required to support Video Conferencing, limits can be applied to the DR traffic to allow that application the required bandwidth.

Keep in mind, when setting bandwidth throttling policies on a container replicating between two physical DR Series system appliances, the policy will be applied for all containers between the replicated devices. For example, containers 1 and 2 on a DR Series system appliance are set to replicate to a secondary DR Series system appliance. When bandwidth throttling is set on container 1, the bandwidth on the second container will also be throttled.

Replication schedules can be set to ensure one container gets more time replicating (for high priority data) or to schedule replication to occur in off-peak hours, although in most situations it is recommended to leave the default settings and let replication transfer data as needed over the network.

When multiple containers are being replicated between the same DR Series system appliances, the replication engine round-robins the requests across the containers. In this situation the containers may not be replicated in synchronicity if there are large amounts of data waiting to be transferred.

## Replication encryption

For best performance vs. security, encryption settings should be set at 128 bit encryption, providing a good balance for most environments. For situations where the replicated data is being transferred across the open Internet, it is strongly recommended to tunnel the replication traffic through an encrypted VPN tunnel.

# Addressing packet Loss

Occasionally a given link between two sites may introduce packet loss, resulting in slow data transfer or replication processes that terminate due to errors or a simple timeout. If this is happening, factors such as packet loss, link speed, jitter or latency should be checked. Once these have been checked and replication still presents issues, please contact Quest Support so they can assist in determining the root cause(s).

## Replication tuning

In some instances it may be beneficial to prevent replication from running all the time. Consider the following scenarios when deciding whether to replicate on a per-container basis.

### High bandwidth conditions

In situations where bandwidth is very high between a source and destination DR Series system appliance, or many sources are replicating to a single target appliance, the backup window may be negatively impacted when backups occur during the replication process. If the backup window is unacceptable in such a scenario, the replication should be rescheduled to occur outside the backup window, or the available bandwidth should be limited to < 50MB/s.

A high bandwidth scenario is when the WAN Bandwidth or Bandwidth between the source and destination DR Series systems is > 100MB/s (800Mbits).

### Reducing bandwidth concerns

A successful replication configuration interacts with its environment positively. In situations where replication traffic is negatively impacting other services, fine tuning becomes necessary to limit the impact.

The following suggestions will help make sure replication is not placing a burden on other services.

- Make certain to schedule the replication process outside the hours where it could impact business. If video conferencing or other business critical applications are experiencing the effect of a low-bandwidth situation, verify the replication schedule and make sure it happens during non-business hours. Also, make sure to account for the backup window when the daily backups are being stored to the DR Series system appliance.
- In situations where the bandwidth of the impacted services is well known, the bandwidth available to the replication process can be reduced. Throttling the replication bandwidth requires more time for the replication process to complete, but leaves room for the impacted services to do their thing. For example, if video

conferencing requires 1MB/s on a 10MB/s link, scale back the available replication bandwidth to 9MB/s, providing bandwidth for everyone to play nice together.

## Domain access

In addition to data, NFS and CIFS security information is also replicated between DR Series system appliances. This allows access to each DR Series system appliance joined to the same domain / forest from user or group accounts with appropriate permissions. Access to a given DR Series system appliance will be denied when not joined to the same domain/forest. Make certain that all DR Series system appliances are joined to the same domain, which provides access to all devices without concern for entering different permissions for each appliance.



**NOTE:** CIFS and NFS protocol access is not synced between source and target DR Series systems. It is not known what is at the remote site so these settings are not transferred. Before failing over, make sure the target DR Series system is setup with the appropriate protocol access as required.

# Networking

## Networking capabilities

The DR Series appliance provides many networking capabilities, designed to further improve the ingest and recovery speeds in any environment. One such feature is secure separation, allowing network optimization by preventing unnecessary traffic on the production network via routing the backup, management, and replication traffic to separate network interfaces. The DR Series system appliance also supports a multiple network interface cards, including 10GbE, providing features such as adaptive load balancing and dynamic link aggregation. The following sections define the various options and configurations associated with the DR Series system appliance, including steps to optimize the appliance to a given environment. As always, the information provided below are guidelines and best practices that are meant to be supplemental to the information provided in the DR Series system administration guide.

## Supported network cards

The DR Series appliance supports various network card configurations. Table 1 outlines the different combinations of network cards with which it is possible to configure the DR Series system appliance. Both the 1GB and 10GB BASE-T network cards support auto-negotiation which is enabled by default. The Broadcom 10GB SPF interface has the ability to auto-negotiate when using the R8H2F transceiver.

When using the Broadcom 10GB SPF network interface card, ensure that the 10GB SR SPF+ Transceiver (R8H2F) is used. Failure to use the proper transceiver will result in errors on the login screen.

*Table 1 Supported Network Card Configurations for the DR Series Appliance*

Combo	Internal (NDC)	Add-On (NIC)	
1	Broadcom (FM487) 4 x 1GB BASE-T	Intel Dual (XP0NY) 2 x 1GB BASE-T	Broadcom (57M9) 2 x 1GB BASE-T

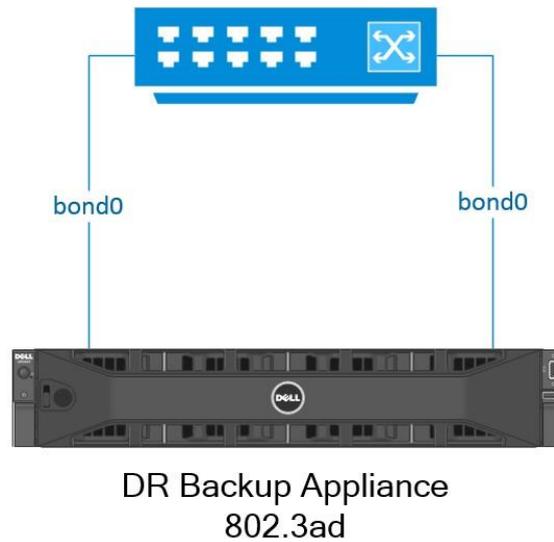
Combo	Internal (NDC)	Add-On (NIC)	
2	Intel Quad R1XFC 4 x 1GB BASE-T	Intel Dual (XP0NY) 2 x 1GB BASE-T	Broadcom (57M9) 2 x 1GB BASE-T
3	Intel Dual (P71JP) 2 x 10GB BASE-T 2x 1GB BASE-T	Intel Dual (XP0NY) 2 x 1GB BASE-T	Broadcom (57M9) 2 x 1GB BASE-T
4	Broadcom (MT09V) 2 x 10GB BASE-T + 2 1GB BASE-T	Intel Dual (XP0NY) 2 x 1GB BASE-T	Broadcom (57M9) 2 x 1GB BASE-T
5	Broadcom (Y36FR) 2 x 10GB SFP + 2 x 1GB BASE-T	Intel Dual (XP0NY) 2 x 1GB BASE-T	Broadcom (57M9) 2 x 1GB BASE-T

## Network interface card bonding

Network interface card (NIC) bonding provides additional throughput and/or failover functionality in the event a link is lost. The DR Series appliance supports two bonding modes: dynamic link aggregation and adaptive load balancing (802.3ad and ALB). Each of these modes has their own advantages and disadvantages that should be considered before choosing a mode.

Dynamic link aggregation (Mode 4 or 802.3ad) creates aggregation groups that utilize the same speed and duplex (i.e. 10GB and 10GB full-duplex links). Mode 4 (See Figure 1) is highly beneficial in increasing speed and bandwidth capability for multiple data streams, but will not increase the speed or bandwidth capability of a single data stream. Slave selection for outgoing traffic is executed according to a simple XOR policy. When utilizing mode 4 it is important to note that the maximum bandwidth available is not always equal to the sum of each link in the bond. Also, always ensure that the switch(s) being used support 802.3ad Dynamic link.

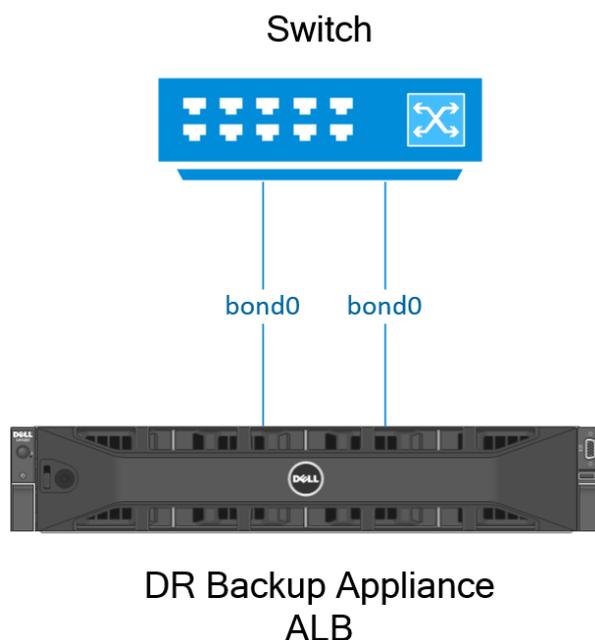
Switch (Switch must be configured to use LACP)



*Figure 1 Dynamic Link Aggregation*

Adaptive load balancing (Mode 6 or ALB), the default load balancing mode, is transmit load balancing with the addition of receive load balancing (See figure 2). The receive load balancing uses address resolution protocol (ARP) to intercept packets and reassign destination MAC addresses. This means that traffic is distributed across slave NICs according to current load. In this mode, it is not possible to mix interfaces of different speeds. (i.e. 10GB and 1GB links) and no specialized switch support is required. When utilizing mode 6, the total available bandwidth of the bond is equal to the bandwidth of a single physical connection.

**i** **NOTE:** Always ensure that data source systems (that is, systems that send data to the DR Series system) are located on the same subnet. Failure to do so will result in all traffic being sent to the first interface in the bond. This is because adaptive load balancing cannot properly load balance when data sources are located on a remote subnet. This is a result of ALB use of (ARP), which is subnet- specific and a router's inability to send ARP broadcast and updates.



*Figure 2 Adaptive Load Balancing*

By default, the DR Series appliance sets up bonding on its fastest interfaces. If the appliance is configured with NIC combo 3 (see Table 1) the 2x 10GB interfaces will become the new bond0. If the DR appliance is configured with combo 1, the 4x 1GB interfaces will become bond0. As of firmware 2.1 the default bond behavior can be changed by utilizing the `auto_bonding_speed` parameter. See below for usage example.

To reset the default bond run the command:

```
network --factory_reset [--auto_bonding_speed <1G|10G>]
```

After resetting the interface the DR appliance will prompt you to enter a reboot command (`System -reboot`).

**i** **NOTE:** This command has to be issued from iDRAC, KVM or a local interface because it will disconnect any existing network connections.

**Example:**

```
administrator@DR1 > network -- factory_reset --auto_bonding_speed 1G
```

Warning: This will stop all system operation and will reset the network configuration to factory settings and will require a system robbon. Existing configuration will be lost.

Password required to proceed. Please enter the administrator password:

```
administrator@DR1 > System -reboot
```

**i** **NOTE:** When creating a bond ensure that the interfaces to bond have active links. This will ensure that the system will be operational after the bonding of the interfaces is complete.

# Secure network separation

Secure network separation is advanced networking functionality that enables administrators to segment traffic to different NIC's and subnets (see Table 2). Advanced networking makes it possible to assign backup traffic to one bond, management traffic to a second bond and replication traffic to a third bond.

This section will cover the following advanced networking scenarios in detail:

- Scenario 1: Leverage separate interfaces for management, replication and backup traffic Scenario 2: Leverage one bonded interface for management, replication and OST traffic and another for backup traffic
- Scenario 3: Replication between sites with dedicated interfaces
- Scenario 4: Multiple appliance replication
- Scenario 5: Backup to different IPs on a single DR Series system appliance

*Table 2 Traffic Segmentation*

<b>Segment</b>	<b>Traffic</b>
Management	GUI /Telnet/ SSH
Replication	Replication/OST Op-dupe
Backup	CIFS/NFS/RDA/OST

## Scenario 1

### **Leverage separate interfaces for management, replication and backup traffic**

Sarah, a network administrator, has an office located in Boston and another located in New York (See Figure 3). At each office she has a DR Series system appliance and wishes to configure separate interfaces for management, backup and replication traffic.

Her desired configuration is as follows:

- Use bond0 for management.
- Use dedicated 1GB interface for replication traffic
- Use dedicated 2x 10GB bonded interfaces for backup traffic.

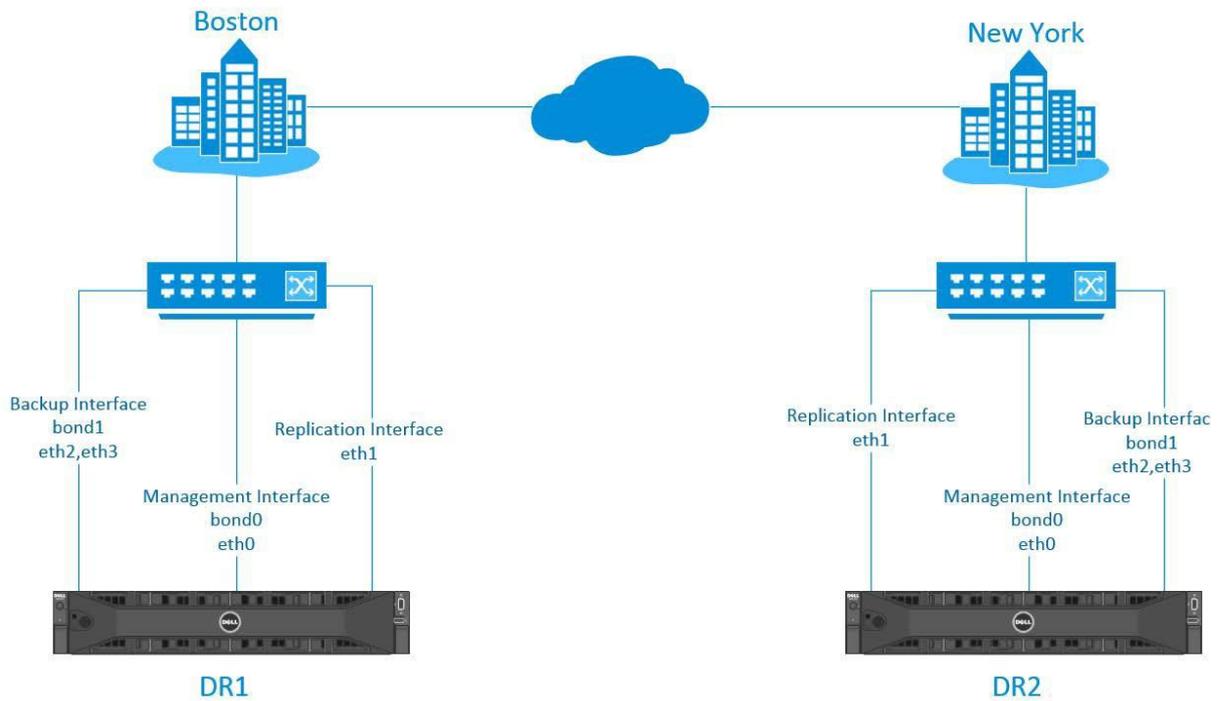


Figure 3 Scenario 1 Topology



**NOTE:** Bond0 should be used for management traffic as it contains the default route and will allow for management of the device from all accessible subnets. For Backup and Replication Traffic, static routes should be added to allow communication between the devices.

#### Hardware configuration:

- DR1– 2x 1GB interfaces and bond0 as 2x 10GB interfaces
- DR2 – 2x 1GB interfaces and 2x 10GB interfaces

## Configure DR1

- 1 Display the original configuration of DR1 using the following command:

```
network -show
```

#### Example:

```
administrator@DR1 > network --show
Automatic bonding speed : 10G
Device                  : bond0
Enabled                 : yes
Link                    : yes
Boot protocol           : dhcp
```

```

IP Addr           : 10.250.243.132
Netmask          : 255.255.252.0
Gateway          : 10.250.240.1
MAC Addr         : BC:30:5B:F3:22:70
MTU              : 9000
Bonding options  : "mode=balance-alb miimon=100 xmit_hash_policy=2"
Member Interfaces : eth2,eth3
Interface name   : swsys-78
eth2 MAC         : 78:2B:CB:1B:4A:28
eth2 Max Speed   : 10000baseT/Full
eth2 Speed       : 10000Mb/s
eth2 Duplex      : Full
eth3 MAC         : 78:2B:CB:1B:4A:29
eth3 Max Speed   : 10000baseT/Full
eth3 Speed       : 10000Mb/s
eth3 Duplex      : Full

```

- 2 Configure the DR Series system to use the 1GB interfaces as bond0 members and free the 10GB interfaces for backup traffic.

- a Perform a factory reset of the network and then reboot the appliance:

```

network --factory_reset --auto_bonding_speed_1G
system --reboot

```

Example:

```

administrator@DR1 > network --factory_reset --auto_bonding_speed 1G

```

```

WARNING: This will stop all system operation and will reset the network
configuration to factory settings and will require a system reboot.
Existing configuration will be lost.

```

```

Password required to proceed. Please enter the administrator password:
Resetting network configuration, please wait....

```

```

Reboot the system using the command 'system --reboot' to complete the
network factory

```

```

reset

```

```

administrator@DR1 > system --reboot

```

```

Please enter administrator password:

```

```

Broadcast message from root (pts/2) (Tue Nov 19 17:43:07 2013): The system
is going down for reboot NOW!

```

- b Use the following command to verify that bond0 now consists of two 1GB interfaces:

```

network --show

```

Example:

```

administrator@DR1 > network --show
Automatic bonding speed: 1G
Device      : bond0
Enabled     : yes
Link       : yes
Boot protocol : dhcp
IP Addr    : 10.250.242.173
Netmask    : 255.255.252.0
Gateway    : 10.250.240.1
MAC Addr   : BC:30:5B:F3:22:74
MTU       : 9000
Bonding options : "mode=balance-alb miimon=100 xmit_hash_policy=2"
Member Interfaces : eth0,eth1
eth0 MAC      : BC:30:5B:F3:22:74
eth0 Max Speed :1000baseT/Full
eth0 Speed    :1000Mb/s
eth0 Duplex   :Full
eth1 MAC      : BC:30:5B:F3:22:75
eth1 Max Speed :1000baseT/Full
eth1 Speed    :1000Mb/s

```

### 3 Break bond0 to create the following configuration:

- Bond0 with single 1GB interface to be used for management.
- A single 1GB interface to be used for replication traffic.
- A bonded 2x 10GB interface to be used for backup traffic.

a Delete eth1 from bond0 using the following command:

```
network -- delete --member <ethN>
```

Example:

```

administrator@DR1 > network -- delete -- member eth1
Interface delete successful. Please restart networking for the changes to take effect.

```

b Create bond1 consisting of 2x 10GB interfaces using the following command:

```
network --create_bond --bondif<bond_name> --nwif<eth1,eth2,..ethN> --static ip<ip_address> --netmask<network_mask> --restart
```

Example:

```

administrator@DR1 > network --create_bond --bondif bond1 --nwif eth2,eth3 -
--static --ip
10.250.242.221 --netmask 255.255.252.0 --restart

```

```

WARNING: During network restart a loss of connection may occur and a
relogin may be necessary. Password required to proceed.

Please enter the administrator password: Restarting network...

Shutting down interface bond0: [ OK ] Shutting down interface bond1: [
OK ]

Shut Bringing up loopback interface: [ OK ]

Bringing up interface bond0:Determining IP information for bond0... done.
[ OK ]

Bringing up interface bond1: [ OK ]

Updating DNS entry for swsys-231.ocarina.local to 10.250.242.173 ..
Skipping DNS Update 10.250.240.4: IP already updated.

Starting the filesystem...doneting down loopback interface: [ OK ]

```

c Create a single 1GB interface using the following command:

```
network --create_eth --nwif< ehtN> --static --ip<ip address> --netmask<netmask> --
name <name> --restart
```

**Example:**

```

administrator@DR1 > network --create_eth --nwif eth1 --static --ip
10.250.243.222 -- netmask 255.255.252.0 -- name DR1-replication --
restart

```

```

WARNING: During network restart a loss of connection may occur and a
relogin may be necessary. Password required to proceed.

```

```

Please enter the administrator password:

```

```

Interface operation successful. Network restart will now be done.

```

```

Restarting network...

```

```

Shutting down interface bond0: [ OK ]

```

```

Shutting down interface bond1: [ OK ] Shutting down loopback interface:
[ OK ] Bringing up loopback interface: [ OK ]

```

```

Bringing up interface bond0:Determining IP information for bond0... done.
[ OK ] Bringing up interface bond1: [ OK ]

```

```

Bringing up interface eth1: [ OK ]

```

```

Updating DNS entry for swsys-231.ocarina.local to 10.250.242.173 ..

```

```

Skipping DNS Update 10.250.240.4: IP already updated.

```

d Enter the following command to verify all changes have been properly made:

```
network -show
```

4 Dedicate traffic to interfaces in the following manner:

- Management - bond0
- Backup Traffic -bond1
- Replication- eth1

- a Dedicate bond0 to management traffic using the following commands:

```
system --mgmt_traffic --add --type < Webserver|Telnet| > --interface  
<bond(0-N) | eth(0-N) | lo#>
```

**Example:**

```
administrator@DR1 > system --mgmt_traffic --add --type Webserver --  
interface bond0  
  
Successfully added application webserver. Restarting webserver service ...  
done.  
  
administrator@DR1 > system --mgmt_traffic --add --type Telnet --interface  
bond0  
  
Successfully added application telnet. Restarting telnet service ... done.
```

- b Dedicate bond1 to CIFS and NFS backup traffic using the following commands:

```
system --backup_traffic --add --type < NFS|CIFS|OST|RDS> --interface <bond(0-N) |  
eth(0- N) | lo#>
```

**Example:**

```
administrator@DR1 > system --backup_traffic --add --type CIFS --interface  
bond1  
  
WARNING: This operation requires Windows access server restart. Do you want  
to continue (yes/no) [n]? y  
  
Successfully added application CIFS. Restarting Windows Access Server...  
Done.  
  
administrator@DR1 > system --backup_traffic --add --type NFS --interface  
bond1  
  
Do you want to continue (yes/no) [n]? y Successfully added application NFS.  
Restarting file system ... done.
```

- c Dedicate eth1 to replication traffic using the following commands:

```
system --replication_traffic --add --interface <bond(0-N) | eth(0-N) | lo#>
```

**Example:**

```
administrator@DR1 > system --replication_traffic --add --interface eth1  
Successfully added application replication
```

## Configure DR2

- 5 Follow steps 1-4 to configure the network interfaces for DR2.
- 6 Setup static replication between DR1 and DR2 by running the following CLI command:

```
replication --add --name backup --role source --peer DR2-replication --peer_name  
backup-from-DR1
```

**Example:**

```
administrator@DR2 > replication --add --name backup --role source --peer DR2-  
replication --peer_name backup-from-DR1
```

Enter password for administrator@DR2-replication: Replication entry created successfully.

```
Replication Container      : backup
Replication Role          : Source
Replication Target        : DR2-replication.ocarina.local
Replication Target IP     : 10.250.243.220
Replication Target Mgmt Name : DR2-replication.ocarina.local
Replication Target Mgmt IP  : 10.250.243.220
Replication Local Data Name : 10.250.243.222
Replication Local Data IP  : 10.250.243.222
Replication Target Container : backup-from-DR1
Replication Enabled       : Yes
Replication Compression Enabled : Yes
Replication Encryption    : Not Enabled
```

- 7 On DR-1 add a route from DR1-eth1 to DR2-eth1 using the following command:

```
network --route --add --network <ip_address>--netmask <netmask> --gateway <gateway_address> --interface <eth(0-N)>
```

**Example:**

```
administrator@DR1 > network --route --add --network 10.250.243.220 --netmask 255.255.255.0 --gateway 10.250.243.100 --interface eth1
```

- 8 On DR-2 add a route from DR1-eth1 to DR2-eth1 using the following command:

```
network --route --add --network <ip_address>--netmask <netmask> --gateway <gateway_address> --interface <eth(0-N)>
```

**Example:**

```
administrator@DR2 > network --route --add --network 10.250.243.222 --netmask 255.255.255.0 --gateway 10.250.243.100 --interface eth1
```

## Scenario 2

### Leverage one bonded interface for management, replication and OST traffic and another for backup traffic

Robert has a DR Series appliance with firmware 2.1, 2x 10GB interfaces and 2x 1GB interfaces. He wishes to use the 1GB interfaces for replication, management and OST traffic. He wants to use the 10GB interfaces for backup traffic only. Robert will need to do the following to accomplish his goals:

- Verify that all interfaces are detected by the DR appliance.
- Issue a factory reset with default bonding set to the 1GB interfaces. This will ensure that the faster interfaces are not used for bond0.
- Create a bond for the 10GB interfaces.

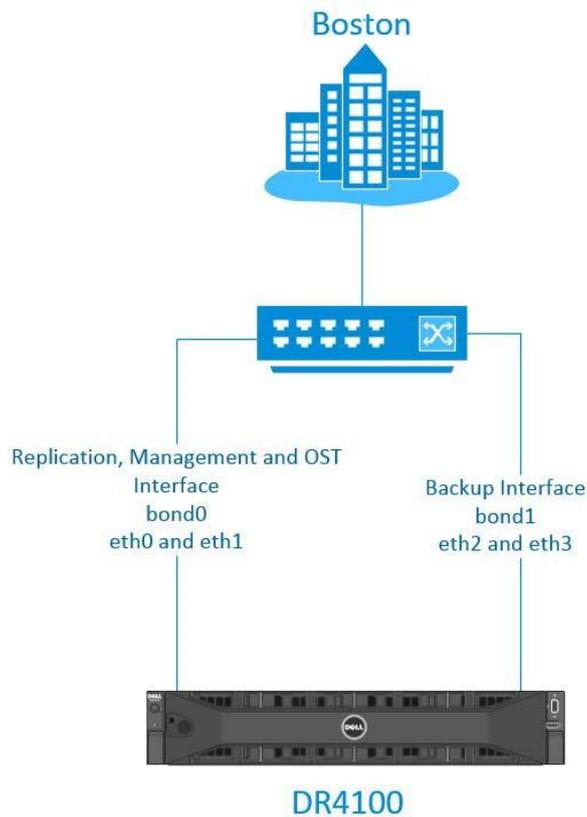


Figure 4 Scenario 2 Topology

**i** **NOTE:** Make sure all 1G interfaces are plugged in and can have an IP address via DHCP. If there is no DHCP server available on the 1G network, then once the system comes up after step-2, a static IP has to be assigned to the 1G bonded interface and a 'network --restart' command should be executed for the system to be back in 'Operational mode'. Bond0 should always have an IP address for the system to be in operational state.

- 1 Set bond0 to 1GB using the following command:

```
network --factory_reset [--auto_bonding_speed <1G|10G>]
```

**Example:**

```
administrator@DR1 > network -- factory_reset -auto_bonding_speed 1G
Warning: This will stop all system operation and will reset the network
configuration to factory settings and will require a system robbon. Existing
configuration will be lost.

One or more of these interfaces 'eth0, eth1' are in use by an application.
Factory reset cannot be done while interfaces are in use by an application.
Factory reset will remove all application interface settings.

Password required to proceed. Please enter the administrator password:
administrator@DR1 > System --reboot
```

- 2 Use the following command to view the new configuration:

**network --show**

**Example:**

```
administrator@DR1 > network -- show
Automatic bonding speed: 1G
Device      : bond0
Enabled     : yes
Link        : yes
Boot protocol : dhcp
IP Addr     : 10.250.242.173
Netmask     : 255.255.252.0
Gateway     : 10.250.240.1
MAC Addr    : BC:30:5B:F3:22:74
MTU         : 9000
Bonding options : "mode=balance-alb miimon=100 xmit_hash_policy=2"
Member Interfaces : eth0,eth1
eth0 MAC      : BC:30:5B:F3:22:74
eth0 Max Speed :1000baseT/Full
eth0 Speed     :1000Mb/s
eth0 Duplex    :Full
eth1 MAC      : BC:30:5B:F3:22:75
eth1 Max Speed :1000baseT/Full
eth1 Speed     :1000Mb/s
eth1 Duplex    :Full
```

- 3 Create a bond for the 2x 10GB interfaces using the following command:

```
network --create_bond --bondif<bondN> --dhcp --nwif <eth2, eth3, ethN> -- mode  
<ALB |802.3ad> --restart
```

**Example:**

```
administrator@DR1 > network --create_bond -bondif bond1 --dhcp --nwif  
eth2,eth3, -- mode ALB --restart
Shutting down interface bond0: [ OK ] Shutting down interface bond1: [ OK ]
Shutting down loopback interface: [ OK ] Bringing up loopback interface: [ OK ]
Bringing up interface bond0:Determining IP information for bond0... done. [ OK ]
Bringing up interface bond1:Determining IP information for bond1... done. [ OK ]
Updating DNS entry for DR1.local to 10.250.xxx.x ..
Skipping DNS Update 10.250.xxx.x: IP already updated.
```

- 4 Use the instructions in Step # 4 in Scenario #1 to configure Management, Replication, and Backup OST traffic for the respective interfaces.

# Scenario 3

## Replication between sites with dedicated interfaces

Daniel has two sites, each with a DR appliance that he wishes to configure as a replication pair over dedicated links.

Each DR appliance will have the following configuration:

- DR1 has 2x 10GB interfaces in bond0 with 2 x1GB interfaces not bonded dedicated to replication
- DR2 has 2x 10GB interfaces in bond0 with 2 x10GB interfaces not bonded dedicated to replication

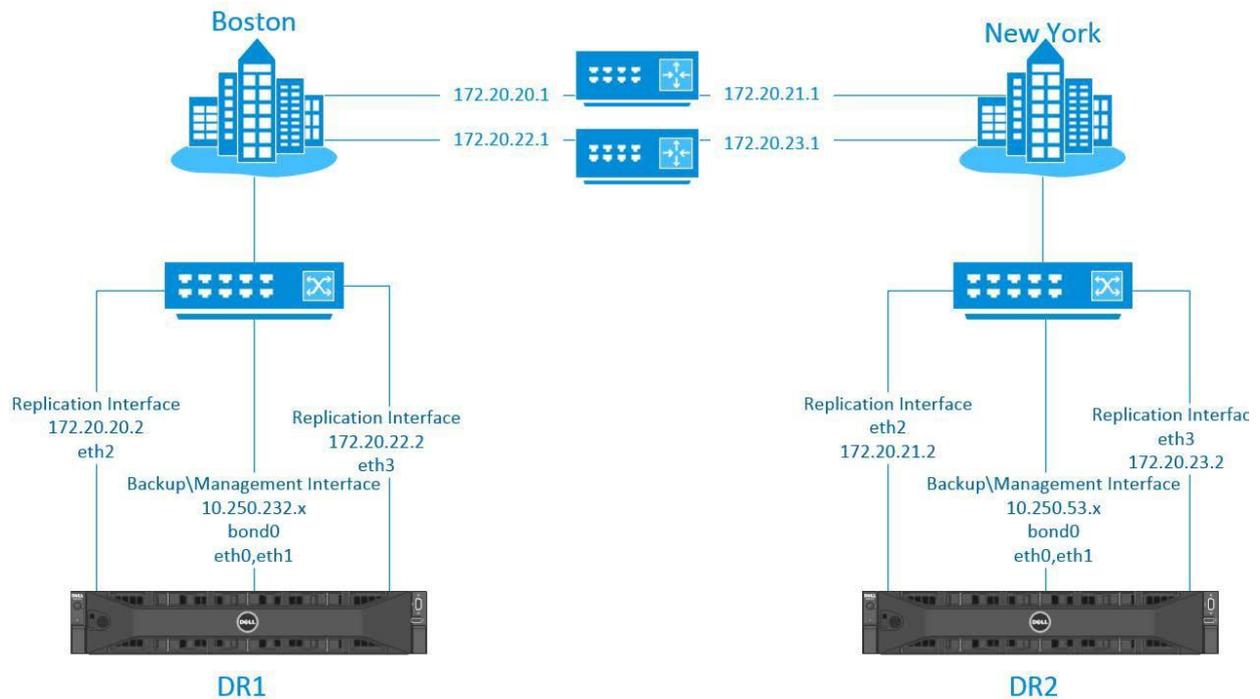


Figure 5 Scenario 3 Topology

In the default configuration of the DR Series system, the default route for the system is always set on the network which belongs to bond0. In this example for both the DR Series systems the default route points to the gateway on the 10.250.232.x network. This allows for administration of the DR Series system from anywhere on the network.

Example route:

```
[root@DR1 ~]# route
Kernel IP routing table
Destination Gateway Genmask Flags Metric Ref Use Iface
10.250.232.0 * 255.255.248.0 U 0 0 0 bond0
default 10.250.232.1 0.0.0.0 UG 0 0 0 bond0
```

In order for the two DR Series systems to replicate over the dedicated links, static routes should be provisioned in such a way that the replication flows over those paths. This is done in the steps below.

- 1 Use the following command to assign a static IP to eth2 on DR1:

```
network --create_eth --nwif <eth0,eth2,ethN> --static --ip <ip_address> --  
netmask <netmask>
```

**Example:**

```
administrator@DR1 > network --create_eth --nwif eth2 --static --ip 172.20.20.2 -  
netmask  
255.255.255.0
```

- 2 Use the following command to assign a static IP to eth3 on DR1:

```
network --create_eth --nwif <eth0,eth2,ethN> --static --ip <ip_address> --  
netmask <netmask>
```

**Example:**

```
administrator@DR1 > network --create_eth --nwif eth3 --static --ip 172.20.22.2 -  
-netmask  
255.255.255.0
```

- 3 Ensure connectivity of DR1 eth2 and eth3 by pinging the gateway using the following command:

```
network --ping --destination <destination_ip_address> --interface <ethN>
```

**Example:**

```
administrator@DR1 > network --ping --destination 172.20.20.1 --interface eth2  
administrator@DR1 > network --ping --destination 172.20.22.1 --interface eth3
```

- 4 On DR2 use the following commands to release 1GB interfaces eth2 and eth3 from bond0:

```
network --delete --member <eth0,eth1,ethN>  
network -restart network --show
```

**Example:**

```
administrator@DR2> network --delete --member eth2,eth3  
administrator@DR2 > network --restart  
administrator@DR2 > network --show
```

- 5 Use the following commands to assign IP address to eth2 and eth3 on DR2:

```
network --create_eth --nwif <eth0,eth2,ethN> --static --ip <ip_address> --  
netmask <netmask>  
  
network --show
```

**Example:**

```
administrator@DR2> network --create_eth --nwif eth2 --static --ip 172.20.21.2 --  
netmask 255.255.255.0  
  
administrator@DR2 > network --create_eth --nwif eth3 --static --ip 172.20.23.2 -  
-netmask  
255.255.255.0 --restart
```

6 Ensure connectivity of DR2 eth2 and eth3 by pinging the gateway using the following command:

```
network --ping --destination <destination_ip_address> --interface <ethN>
```

Example:

```
administrator@DR2 > network --ping --destination 172.20.21.1 --interface eth2
administrator@DR2 > network --ping --destination 172.20.23.1 --interface eth3
```

7 Add route from DR2 eth2 to DR1 eth2 and from DR2 eth3 to DR1 eth3 using the following commands:

```
network --route --add --network <ip_address> --netmask <mask> --gateway
<gateway_ip> -- interface <eth0,eth1,ethN>
network --show -routes
```

Example:

```
administrator@DR2 > network --route --add --network 172.20.20.2 --netmask
255.255.255.0
--gateway 172.20.21.1 --interface eth2
administrator@DR2 > network --route --add --network 172.20.22.2 --netmask
255.255.255.0 --gateway 172.20.23.1 --interface eth3
administrator@DR2 > network --show --routes
Destination      Gateway          Mask             Interface
172.20.20.0      172.20.21.1     255.255.255.0   eth2
172.20.22.0      172.20.23.1     255.255.255.0   eth3
```

8 Add route from DR1 eth2 to DR2 eth2 and from DR1 eth3 to DR2 eth3 using the following commands:

```
network --route --add --network <ip_address> --netmask <mask> --gateway
<gateway_ip> -- interface <eth0,eth1,ethN>
network --show -routes
```

Example:

```
administrator@DR1 > network --route --add --network 172.20.21.2 --netmask
255.255.255.0
--gateway 172.20.21.1 --interface eth2
administrator@DR1 > network --route --add --network 172.20.23.2 --netmask
255.255.255.0
--gateway 172.20.23.1 --interface eth3
administrator@DR1 > network --show --routes
Destination      Gateway          Mask             Interface
172.20.21.0      172.20.20.1     255.255.255.0   eth2
172.20.23.0      172.20.22.1     255.255.255.0   eth3
```

9 From either DR1 or DR2 verify network connectivity using the following command:

```
network -- ping -- destination <ip_address> --interface<eth0,eth1,ethN>
```

Example:

```
administrator@DR1 > network --ping --destination 172.20.23.2 --interface eth3 --tries 1
```

- 10 Assuming that DR1 and DR2 already had replication established, stop replication and update it to use the newly configured interfaces using the commands below.

```
replication --stop --name <replication_container> --role <source|target>
replication --update --name <replication_container> --peer <target_ip_address> -
- replication_traffic <source_ip_address> --role source
```

**Example:**

```
administrator@DR1 > replication --stop --name cv-replicated-maglib --role source
administrator@DR1 > replication --update --name cv-replicated-maglib --peer
172.20.21.2 --
replication_traffic 172.20.20.2 --role source
```



**NOTE:** If a replication pair does not already exist between DR1 and DR2, use the steps outlined in scenario 1 for assigning replication to designated interfaces and to create a replication pair.

## Scenario 4

### Multiple appliance replication

Jose has one DR Series appliance located at his Seattle site and two DR appliances located at his Lansing site. He wants to replicate data from his Seattle site (Seattle1) to his Lansing site (Lansing1). He would also like to replicate data backed up at Lansing1 to another site, Lansing2.

Jose will need to do the following to accomplish his goals:

- Create bonds on the appropriate interfaces on all three appliances.
- Add replication to the designated interfaces.
- Establish network routes between the DR appliances. The appliances will be configured as follows:
- Seattle1: 2x 10GB interfaces in bond0 and 2x 1GB interfaces in bond1
- Lansing1: 2x 10GB interfaces in bond0 and 2x 1GB interfaces in bond1
- Lansing2: 2x 10GB interfaces in bond0

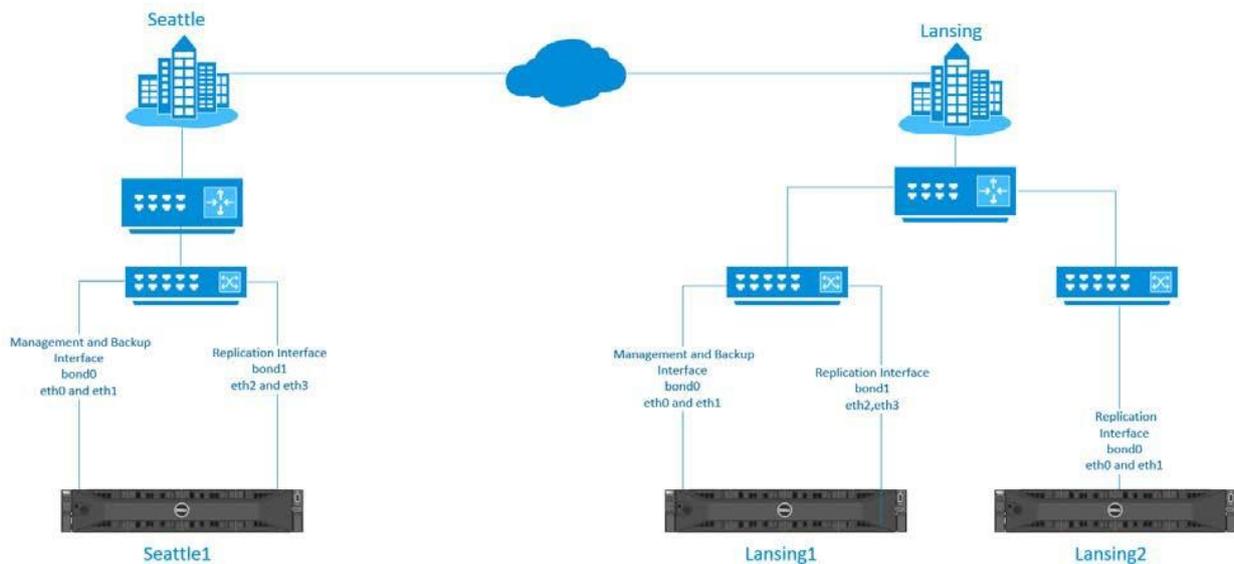


Figure 6 Scenario 4 Topology

- 1 On Lansing1 create a bond on the 1GB ports using the following command:

```
network --create_bond --bondif <bondN> --static -nwif <eth0,eth1,ethN> --mode <ALB | 802.3ad> -mtu <512-9000> --ip <ipaddress> --netmask -restart
```

**Example**

```
administrator@Lansing1 > network --create_bond --bondif bond1 --static -nwif eth2,eth3 -- mode ALB --mtu <512-9000> --ip <ipaddress> --netmask --restart
```

- 2 Add replication to Lansing1 bond1 using the following command:

```
system --replication_traffic --add --interface <bond(0-N) | eth(0-N) | lo#>
```

**Example**

```
administrator@Lansing1 > system --replication_traffic --add --interface bond1
```

- 3 On Seattle1 create a bond on the 1GB ports using the following command:

```
network --create_bond --bondif <bondN> --static -nwif <eth0,eth1,ethN> --mode <ALB | 802.3ad> -mtu <512-9000> --ip <ipaddress> --netmask -restart
```

**Example**

```
administrator@Seattle1 > network --create_bond --bondif bond1 --static -nwif eth2,eth3 -- mode ALB --mtu <512-9000> --ip <ipaddress> --netmask --restart
```

- 4 Create a route from Seattle1 to Lansing1 using the following command:

```
network --route --add --network <destination_network> --gateway <gateway addresses> -- interface <bond(0-N) | eth(0-N) | lo#>
```

**Example**

```
administrator@Seattle1> network --route --add --network <Lansing1's 1G network> -- gateway <gateway addresses> --interface bond1
```

**5 Create a route from Lansing1 to Seattle1 using the following command:**

```
network --route --add --network <destination_network> --gateway <gateway addresses> -- interface <bond(0-N) | eth(0-N) | lo#>
```

**Example**

```
administrator@Lansing1 > network --route --add --network <Seattle's 1G network> -- gateway <gateway addresses> --interface bond1
```

**6 Establish replication from Seattle1 to Lansing1 using the following command:**

```
replication --add --name < container-name> --role <source | target> --peer <IP address> -- replication_traffic <ip address | hostname> --encryption <none | aes128 | aes256>
```

**Example**

```
administrator@Seattle1 > replication --add --name backup --role source --peer <IP of 1G Bond of Lansing1> --replication_traffic <ip address of local 1G interface to be used for replication| hostname> --encryption aes256
```

**7 On Lansing2 add replication to the 10GB bond using the following command:**

```
system -replication_traffic --add --interface <bond(0-N) | eth(0-N) | lo#>
```

**Example**

```
administrator@Lansing2 > system --replication_traffic --add --interface bond0
```

**8 Create a route from Lansing2 to Lansing1 using the following command:**

```
network --route --add --network <destination_network> --gateway <gateway addresses> -- interface <bond(0-N) | eth(0-N) | lo#>
```

**Example**

```
administrator@Lansing2 > network -route --add --network <Lansing1's 10G network> -- gateway <gateway addresses> --interface bond0
```

**9 Create a route from Lansing1 to Lansing2 using the following commands:**

```
network --route --add --network <destination_network> --gateway <gateway addresses> -- interface <bond(0-N) | eth(0-N) | lo#>
```

**Example**

```
administrator@Lansing1 > network --route --add --network <Lansing2's 10G network> -- gateway <gateway addresses> --<IP of 10Gbond of Lansing1>
```

**10 Establish replication from Lansing1 to Lansing2 using the following command:**

```
replication --add --name < container-name> --role <source | target> --peer <IP address> -- replication_traffic <ip address | hostname> --encryption <none | aes128 | aes256>
```

**Example**

```
administrator@Lansing1 > replication --update --name <source-container-name> -- role source --peer <IP of 10G Bond of Lansing2> --replication_traffic <IP of 10G bond of Lansing1>
```

# Scenario 5

## Backup to different IP's on a single DR appliance

Michelle is in Los Angeles and has a Netbackup media server that she wishes to have backed up to her DR appliance via different IP addresses. Her appliance is running an older firmware and needs to be upgraded to the 2.1.

Michelle will need to do the following to accomplish her goals:

- Upgrade the DR Series system to firmware to 2.1.
- Set the media server to use the two interfaces of the DR Series appliance. The DR Series appliance will be configured as follows:
  - 1G interfaces are not exposed pre 2.1 firmware upgrade
  - Bond0 will contain 2x10G interfaces
  - 2x1G interfaces not bonded

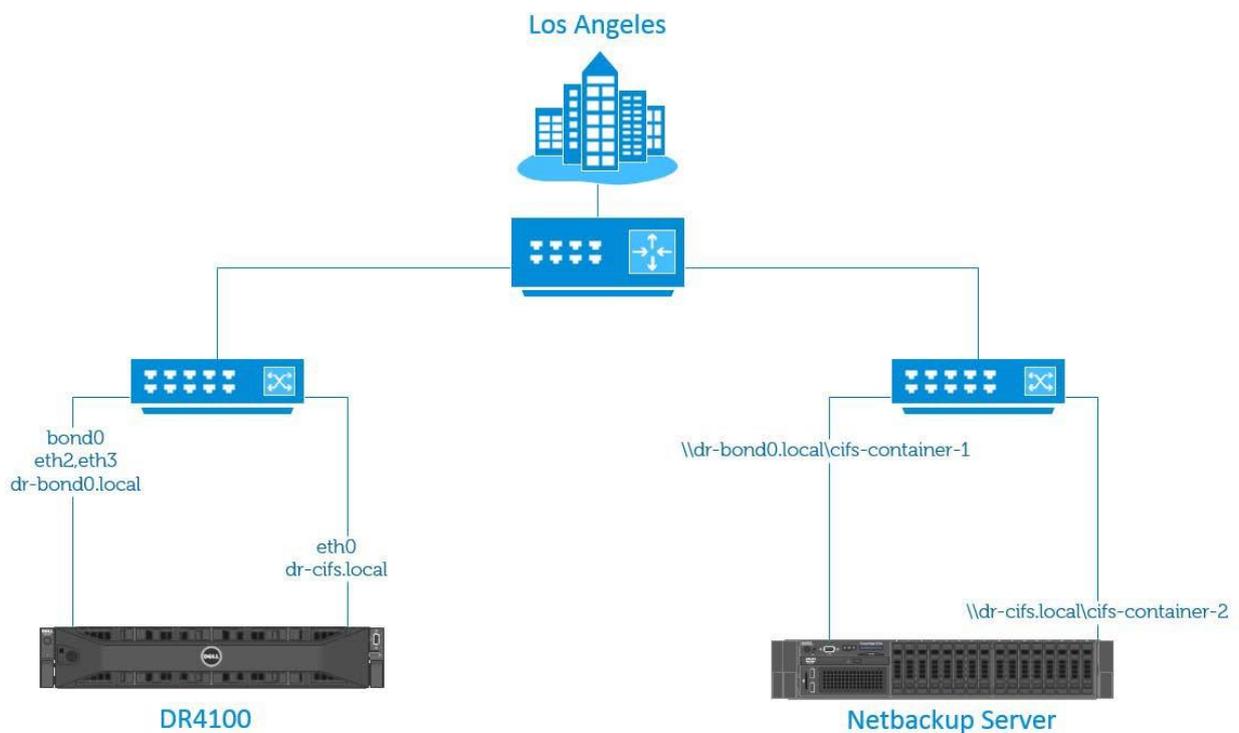


Figure 7 Scenario 5 Topology

- 1 Upgrade the DR Series system appliance to the latest 2.1 firmware.
- 2 After the upgrade, two additional 1 GB interfaces will appear. The 10GB interfaces will be bonded as bond0. Use the following command to view the available network interfaces:

```
network --show
```

- 3 Break bond0 and release eth0 from bond0 using the following command:

```
network -- delete --member <eth(0-N)>
```

```
network -- restart
```

- 4 Use the following command to view the released interface eth0:

```
network --show
```

- 5 Assign an IP address and DNS name to the released interfaces using the following command:

```
network --create_interface --nwif< ehtN> --static --ip<ip address> --  
netmask<netmask> -- name<dns name> --restart
```

**i** | **NOTE:** Bond0 will maintain the default name given to it at time of setup.

- 6 Use the following command to verify network settings:

```
network --show
```

**i** | **NOTE:** Take note of the DNS names for bond0 and for the configured 1GB interface. These DNS names are required to force traffic to the specified interface.

- 7 Create two CIFS containers on the DR Series system appliance. One container called cifs-container-1 and the other cifs-container-2.

On the Netbackup Media server:

- 1 Check connectivity between the Media server and the newly created configured interface on the DR Series system using the interface's DNS name (dr-cifs.local).
- 2 Create a storage unit on the Media server with the UNC path to the DR Series system's CIFS (or NFS if using an NFS Media server) container. Use the DNS name of the newly created interface in the UNC path.

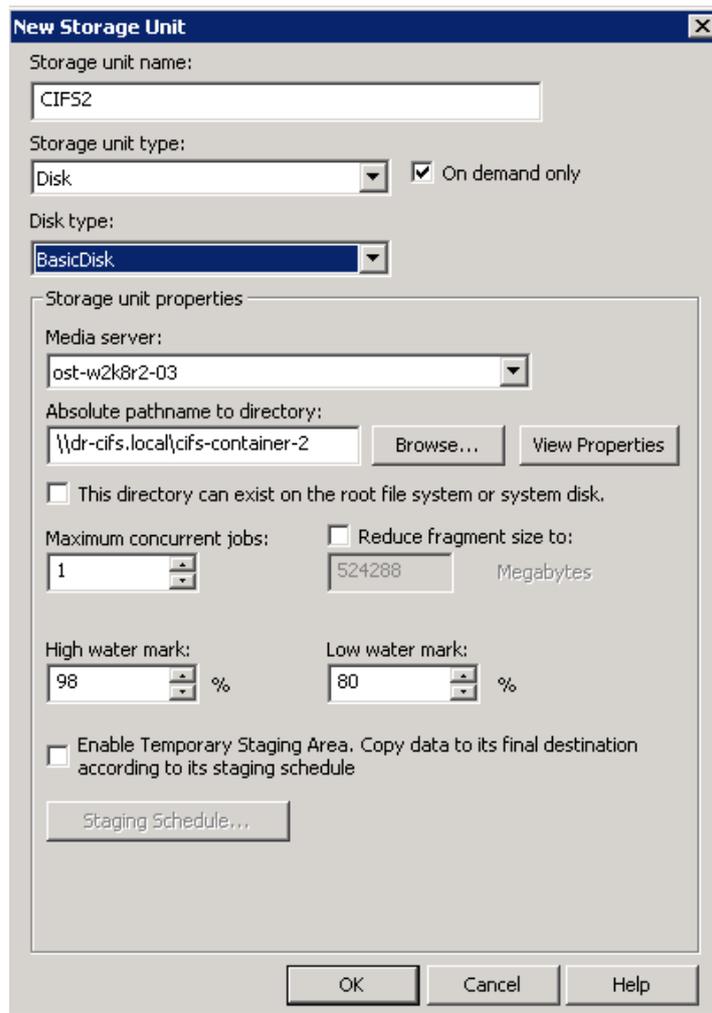


Figure 8 Creating a new storage unit

- 3 Create a second storage unit with the UNC path to the DR Series system's second container. This time use the bond0's name in the UNC path

```
\\dr-bond0.local\cifs-container-1
```

**i** | **NOTE:** Note that the two storage units do not have to point to separate containers.

## Troubleshooting

Follow the steps below to troubleshoot connectivity problems between source and target DR Series systems.

- 1 Enter the following command to the target command to troubleshoot connectivity:

```
replication -- troubleshoot --peer <ip_address>
```

#### Example

```
administrator@DR1 > replication --troubleshoot --peer 10.250.243.222
```

```
Testing connection to port 9904... Connected! Testing connection to port 9911...  
Connected! Testing connection to port 9915... Connected! Testing connection to  
port 9916... Connected!
```

```
Replication troubleshooting completed successfully - Connection to all ports is  
OK!
```

- 2 Make sure the management interface is up and reachable by using the following command:

```
telnet<managment_ip>
```

# How to achieve peak performance with the DR6000

## Guidelines

As part of peak performance tests on DR6000, a maximum throughput of roughly 27TB/hr for RDA Deduplication protocol was achieved. 90% dupe data was used to achieve peak performance. The DR6000 used for testing was implemented with 2 x 10G Intel NIC cards. The following steps were used to achieve 27TB/hr performance.

Prerequisites needed for peak performance testing:

- 1 A DR6000 with 4 x 18TB expansion shelves was configured. Ensure the background initialization is completed before starting tests. Use Default fast compression.
- 2 Configure the MTU interface setting to 9000 (jumbo frames) for both the DR Series system and the clients.
- 3 Testing was done with a clean slate system with no pre-existing data. Reboot the DR Series system in order to clear buffers, and caches.
- 4 Both the DR Series system appliance and clients should be on 10G network, and connected back-to- back. Check network speed between DR Series system appliance and media server or clients using iperf (Refer CLI guide for DR Series system commands: `network --iperf_server`, `network --iperf_client`)
- 5 Ensure that all clients are installed with the RDA plugin compatible to the OS that is running on each.
- 6 Schedule backups, replication and cleaner to run at different times.

Maximum write (ingest) throughput was achieved using 9X clients (1 stream each). Ensure each client has 9 or more cores. Also, the backup speed depends on the speed at which data can be read from source / primary

storage. To ensure the reading of disk is not a bottleneck, confirm the disk read speed is equal to or greater than 100MB/s x the number of cores in the client.

The testing steps to achieve 27TB/hr are as follows:

- 1 Create a RDA container on the DR Series system appliance with connection type as RDS.
- 2 Update the clients with dedupe mode on the DR Series system appliance.  

```
rda -update_client --name <linux_client> --mode dedupe
```
- 3 Install the RDA plugin on all clients.
- 4 Create 90% of deduplicated data on each client that will be ingested into the DR Series system.
- 5 During data ingestion, keep at least 1TB free space on the DR Series system appliance so the cleaner doesn't need to run while backups are taking place.
  - a Start parallel data ingestion from all clients (in our case we started 9 parallel streams i.e., 1 stream from each client).
  - b Once ingest is started, (i.e., 9 streams) one of the thread will be finishing first. Record the "Current Bytes" and "time" when the first thread finished. These will be the end bytes and end time which will be used in calculations.

**i** | **NOTE:** Note the stats at the end of first thread because only during this time all parallel ingests are going on.

- c At this point, the Start and End bytes along with the time elapsed during this process are found.
- d Use the following formulae to calculate the peak:

$$\text{RDA Peak} = (\text{Start bytes} - \text{End bytes}) / (\text{Start time} - \text{End time})$$

This equation will show MiBps. Convert it to MBps then to TB/hr. During Quest Labs testing, iterations of adding clients and repeated testing are taken place until a point where the RDA peak number obtained is saturated. Additionally, if a client has less than 32 cores, RDA will require 256MB of memory per stream. If a client has 32 or more cores, RDA will require 256MB + (4MB \* (# of cores - 32)) of memory per stream. The minimum CPU and system memory requirement for RDA to support client side deduplication is 4 cores and 2GB of memory in a client.

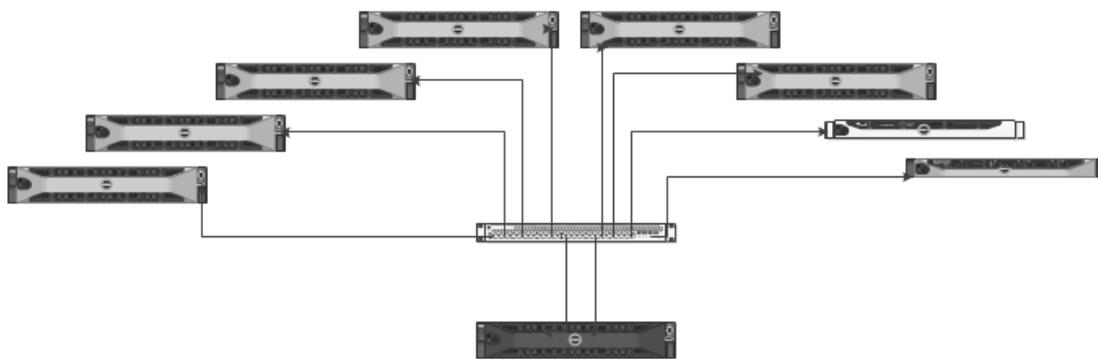
Quest Testing labs used the following list of clients for peak performance testing:

Hostname	Type	HW	NIC Speed	OS Type
Client-1	Linux	R310	10G	CentOS 6.6
Client-2	Linux	R620	10G	RHEL 5.8
Client-3	Linux	R720xd	2 x 10G	CentOS 5.7

Client-4	Linux	R720xd	2 x 10G	CentOS 5.7
Client-5	Linux	R720xd	2 x 10G	CentOS 5.7
Client-6	Linux	R720xd	2 x 10G	CentOS 5.7
Client-7	Linux	R720xd	2 x 10G	CentOS 5.7
Client-8	Linux	R720xd	2 x 10G	CentOS 5.7
Client-9	Linux	R720xd	2 x 10G	CentOS 5.7

PowerConnect Switch      8024F      24 x 10G

- 168 client CPU cores were used to achieve 27TB/hr.



**i** | **NOTE:** The number of hops from any client to DR Series system will be 'one'

```
[root@r310-perf-04 ~]# traceroute dr6000-05
traceroute to dr6000-05
(10.250.212.94), 30 hops max, 60 byte packets
1 dr6000-05.ocarina.local (10.250.212.94) 0.953 ms 0.957 ms 0.951 ms
[root@r310-perf-04 ~]#
```

# How to achieve peak performance with the DR6300

## Guidelines

As part of peak performance tests on DR6300, a maximum throughput of roughly 33TB/hr for RDA Deduplication protocol was achieved. 90% dupe data was used to achieve peak performance. The DR6300 used for testing was implemented with 2 x 10G Intel NIC cards. The following steps were used to achieve 33TB/hr performance.

Prerequisites needed for peak performance testing:

1. A DR6300 with 4 x 72TB expansion shelves was configured. Ensure the background initialization is completed before starting tests. Use Default fast compression.
2. Configure the MTU interface setting to 9000 (jumbo frames) for both the DR Series system and the clients.
3. Testing was done with a clean slate system with no pre-existing data. Reboot the DR Series system in order to clear buffers, and caches.
4. Both the DR Series system appliance and clients should be on 10G network, and connected back-to-back. Check network speed between DR Series system appliance and media server or clients using iperf (Refer CLI guide for DR Series system commands: network --iperf\_server, network --iperf\_client)
5. Ensure that all clients are installed with the RDA plugin compatible to the OS that is running on each.
6. Schedule backups, replication and cleaner to run at different times.

Maximum write (ingest) throughput was achieved using 8X clients (1 stream each). Ensure each client has 8 or more cores. Also, the backup speed depends on the speed at which data can be read from source / primary storage. To ensure the reading of disk is not a bottleneck, confirm the disk read speed is equal to or greater than 100MB/s x the number of cores in the client.

The testing steps to achieve 33TB/hr are as follows:

- Create a RDA container on the DR Series system appliance with connection type as RDS.
- Update the clients with dedupe mode on the DR Series system appliance.
  - o `rda -update_client --name <linux_client> --mode dedupe`
- Install the RDA plugin on all clients.
- Create 90% of deduplicated data on each client that will be ingested into the DR Series system.
- During data ingestion, keep at least 1TB free space on the DR Series system appliance so the cleaner doesn't need to run while backups are taking place.

- Start parallel data ingestion from all clients (in our case we started 8 parallel streams i.e., 1 stream from each client).
  - e Once ingest is started, (i.e., 8 streams) one of the thread will be finishing first. Record the “Current Bytes” and “time” when the first thread finished. These will be the end bytes and end time which will be used in calculations.

**i** | **NOTE:** Note the stats at the end of first thread because only during this time all parallel ingests are going on.

f At this point, the Start and End bytes along with the time elapsed during this process are found.

g Use the following formulae to calculate the peak:

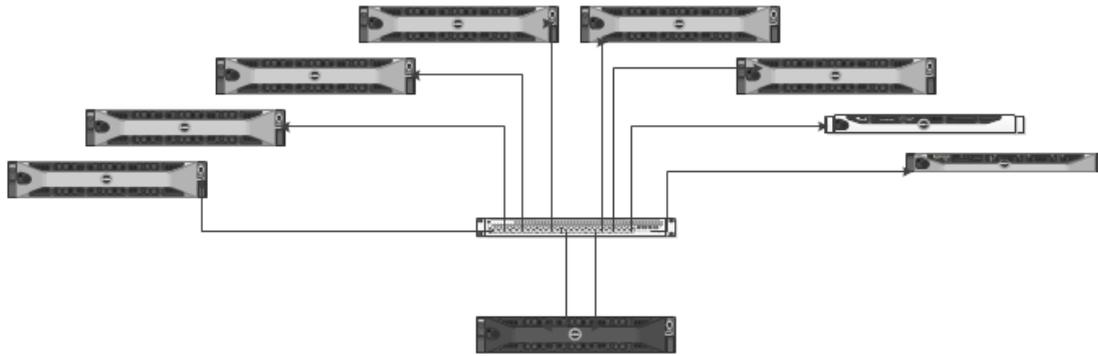
$$\text{RDA Peak} = (\text{Start bytes} - \text{End bytes}) / (\text{Start time} - \text{End time})$$

This equation will show MiBps. Convert it to MBps then to TB/hr. During Quest Labs testing, iterations of adding clients and repeated testing are taken place until a point where the RDA peak number obtained is saturated. Additionally, if a client has less than 32 cores, RDA will require 256MB of memory per stream. If a client has 32 or more cores, RDA will require 256MB + (4MB \* (# of cores – 32)) of memory per stream. The minimum CPU and system memory requirement for RDA to support client side deduplication is 4 cores and 2GB of memory in a client.

Quest Testing labs used the following list of clients for peak performance testing:

Hostname	Type	HW	NIC Speed	OS Type
Client-1	Linux	R310	10G	CentOS 6.6
Client-2	Linux	R620	10G	RHEL 5.8
Client-3	Linux	R720xd	2 x 10G	CentOS 5.7
Client-4	Linux	R720xd	2 x 10G	CentOS 5.7
Client-5	Linux	R720xd	2 x 10G	CentOS 5.7
Client-6	Linux	R720xd	2 x 10G	CentOS 5.7
Client-7	Linux	R720xd	2 x 10G	CentOS 5.7
Client-8	Linux	R720xd	2 x 10G	CentOS 5.7
PowerConnect Switch		8024F	24 x 10G	

- 168 client CPU cores were used to achieve 33TB/hr.



**i** | **NOTE:** The number of hops from any client to DR Series system will be 'one'

```
[root@R310-perf-03 ~]# traceroute dr6300-12
traceroute to dr6300-12 (10.250.213.49), 30 hops max, 60 byte packets
 1 dr6300-12.ocarina.local (10.250.213.49)  2.063 ms  2.075 ms  2.072 ms
[root@R310-perf-03 ~]#
```

## Resources

### DR Series manuals

- [DR4300e Series appliance manuals](#)
- [DR4300 Series appliance manuals](#)
- [DR6300 Series appliance manuals](#)
- [DR4100 Series appliance manuals](#)
- [DR6000 Series appliance manuals](#)

### Quest support

<http://support.quest.com>