DR Series - Best Practice Guide

Part 1: Setup, Replication and Networking

Dell Data Protection Group

July 2016 | Version 2.0

Contents

Introduction

This document contains some of the best practices for deploying, configuring, and maintaining a Dell DR Series backup and deduplication appliance in a production environment. Following these best practices can help ensure the product is optimally configured for a given environment. This document applies to all DR series models including the DR Virtual appliance.

Note, this guide is not a replacement for the administrator's guide. In most cases, this guide provides only a high level description of the problem and solution. Please refer to the administrator guide for the specific options and steps required to implement the recommended actions.

Please consider the administrator's guide as a pre-requisite for this paper, or at least as a reference guide on how to complete the tasks listed here in this guide. The administrator's guide can be found in the list of links below.

In addition to this document, we recommend the following documents for the advanced reader.

DR Series Manuals

Dell DR Series System Administrator Guide

Dell DR Series System Interoperability Guide

Dell DR Series System Command Line Reference Guide

Other DR Series Documentation

DR6300

DR4300

DR4300e

DR2000v

Case Studies

Monrif Group

Durham Region Police Service

Nissan

DCIG – Dell Encryption and VTL for DR Series

ESG DR4X00 Lab Review

Wholesale Electric Supply

System Setup

Hardware

When the DR 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 product has the ability to notify administrators of recent updates. Simply register the DR 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 product has built in monitoring through Dell 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.

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

Password Reset

The DR 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 is to be joined to an Active Directory domain it is recommend 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 to Active Directory

Logon to the DR GUI, Click on Active Directory in the side panel, click on join in the upper right hand corner and enter in the name of the domain and credentials to add the DR 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 DR product 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 is not joined to an Active Directory domain, it is best to configure the DR to use NTP. If the DR is joined to an Active Directory, it will automatically sync its time.

Container Creation

The DR appliance uses containers to store backup data. These containers are segmented folders that have individual protocols, security permissions, marker types and connection types. Whatever number of containers are created, the DR Series deduplicates across all containers.

Below are considerations to take into account when creating containers:

1. Depending on what model of DR appliance is being used, there is a maximum limit to the number of containers that can be created. Refer to the DR Interoperability Matrix for maximum container limits.
2. Access protocols are set at a container level (NFS, CIFS, OST, RDS, iSCSI, NDMP)
3. Security is set at a container level (Locking down Via IP, Unix/Windows ACLs, etc.)

4. Markers are set at the container level
5. OST quotas are set on a container level and are applied at non-deduplicated capacities.
6. Container names cannot contain spaces

Container limit

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.

Recommendation: Robert should create a forth 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 Dell 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, NDMP Only, iSCSI Only

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 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, 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, TSM, ARCserve or HP Data Protector, the marker type setting should be set to Auto. The remaining supported DMA's will be set to it's matching Marker Type.

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 DR4100 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 administration guide.

Replication considerations

- Different DR models, older & newer, larger & smaller, virtual and physical can replicate to and from each other if they are running the same OS version (major.minor). For example. A DR4000 can replicate to a DR6300 if they are both running version 3.2.X version.
- Replicated data is already compressed and deduplicated prior to its transfer to the destination DR appliance. This results in approximately 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 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 to simplify disaster recovery.
- Replication also replicates CIFS and NFS security bits. For CIFS shares the target DR also needs to be in the same domain or forest as the source DR for ACLs to be applied correctly.
- The DR replication target container is read only until the replication relationship with the primary DR is removed.

Replicating Containers

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:

1. In larger environments it can be easy to quickly approach a DR container limit. A simple method to reduce the number of replicated containers is to leverage container directories.
2. In some scenarios it may become necessary to replicate more than a DR container limit 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. In order to calculate the time required to replicate a given container the following two points of data are required:

1. 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.

2. Determine the effective bandwidth. A common method to determine bandwidth between two sites utilizes a freeware product called iPerf (https://code.google.com/p/iperf/). 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
   c. Capture the resulting bandwidth reported from the server at the central site, showing the calculated bandwidth between sites.

   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.

3. 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.

replication_interval = replica_data / effective_bandwidth

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

replication_interval = 314572 MB / 10 MB = 31457 seconds (8.7 hours)

Bandwidth Optimization

When utilizing replication, it may be necessary to enforce bandwidth limitations. This 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 appliances, the policy will be applied for all containers between the replicated devices. For example, containers 1 and 2 on a DR appliance are set to replicate to a secondary DR 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 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. When packet loss arises due to slow links, higher speed links, jitter or high latency, the window size may need to be adjusted to account for this. In such instances please contact the Dell Support department, and they will assist in custom tuning the window size for the particular link.

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 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 units 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.

1. 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 appliance.

2. 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 appliances. This allows access to each DR appliance joined to the same domain / forest from user or group accounts with appropriate permissions. Access to a given DR appliance will be denied when not joined to the same domain/forest. Make certain that all DR 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 DRs. It is not known what is at the remote site so these settings are not transferred. Before failing over, make sure the target DR is setup with the appropriate protocol access as required.

Networking

The DR 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 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 DR4100, 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 administration guide.

Supported Network Cards

The DR4100 supports various network card configurations.. 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 DR Series login screen.

The list of supported network card configurations for each DR model can be found within the Dell DR Series System Interoperability Guide.

Network Interface Card Bonding

Network interface card (NIC) bonding provides additional throughput and/or failover functionality in the event a link is lost. The DR4100 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.

Figure 1. Dynamic link aggregation

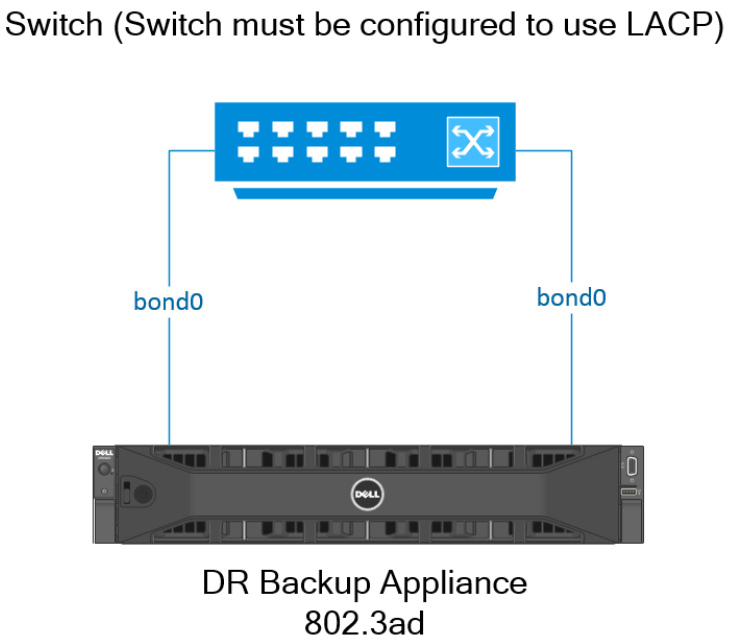

Adaptive load balancing (Mode 6 or ALB), the default load balancing mode, is transmit load balancing with the addition of receive load balancing (*See figure 2*). The receive load balancing uses address resolution protocol (ARP) to intercept packets and reassign destination MAC addresses. This means 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.

When ALB is used with multiple switches, spanning tree issues can arise. Consideration of one of the following may be needed:

1. Disable Spanning Tree
2. Connect to a single switch using ALB
3. Enable 802.3ad

Note: Always ensure that data source system (i.e. systems that send data to the DR4100) 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

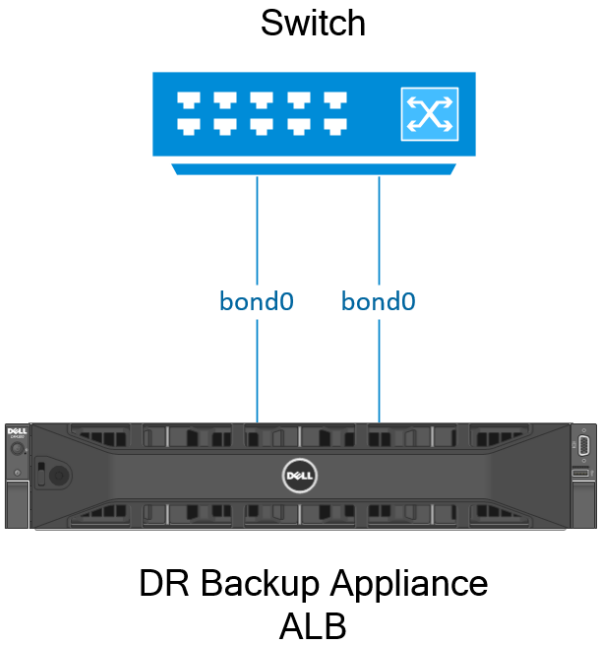

By default, the DR4100 sets up bonding on its fastest interfaces. If the appliance is configured with NIC combo 3 (*see Table 1*) the 2x 10GB interfaces will become the new bond0. If the DR4100 is configured with combo 1, the 4x 1GB interfaces will become bond0. As of firmware 2.1 the default bond behavior 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 DR4100 will prompt you to enter a reboot command**.**

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

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 robbon. Existing configuration will be lost.

Password required to proceed. Please enter the administrator password :

administrator@DR1 > System --reboot
```

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 1*). 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 appliance

Table 1. Traffic segmentation

Segment	Traffic
Management	GUI /Telnet/ SSH
Replication	Replication/OST & RDA Op-dupe
Backup	CIFS/NFS/RDA/OST/iSCSI/NDMP

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 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

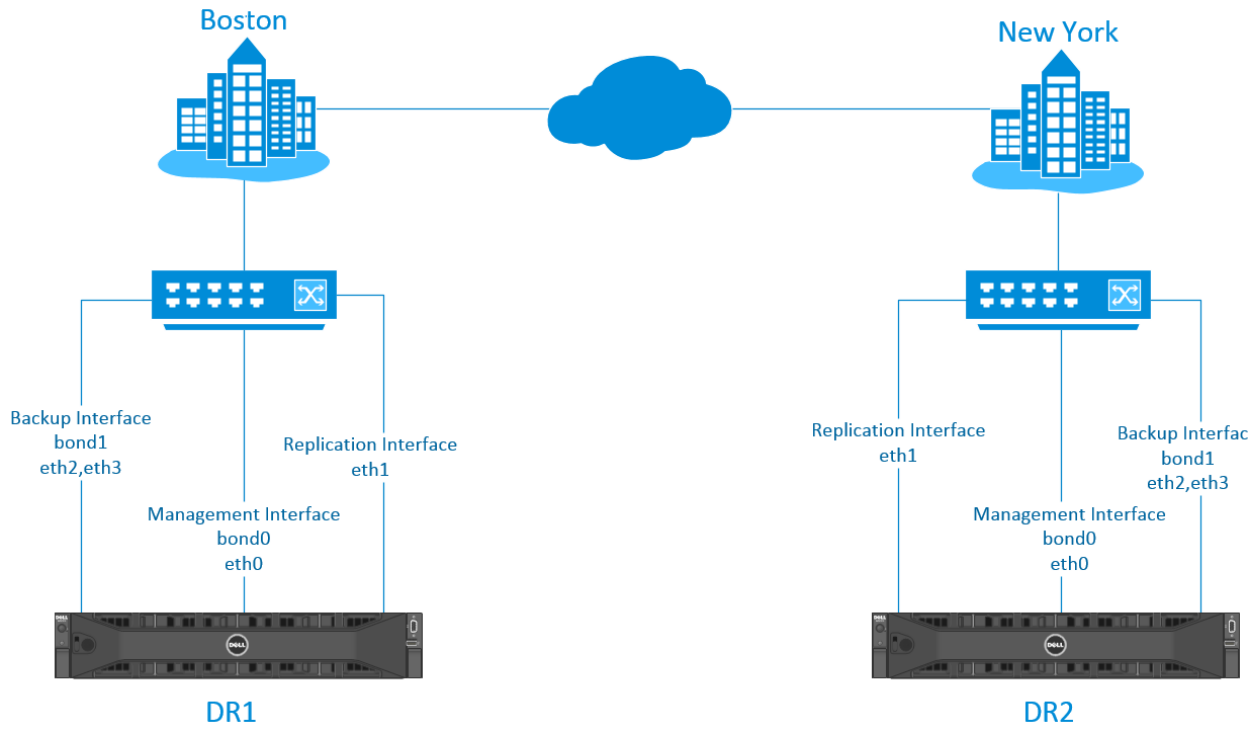

Note: Bond0 should be used for management traffic as it contains the default route and will allow for management of the device from all accessible subnets. For Backup and Replication Traffic. static routes should be added to allow communication between the devices.

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 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
eth1 Duplex             :Full
```

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) Issue 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|NDMP|ISCSI|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 DR4100 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

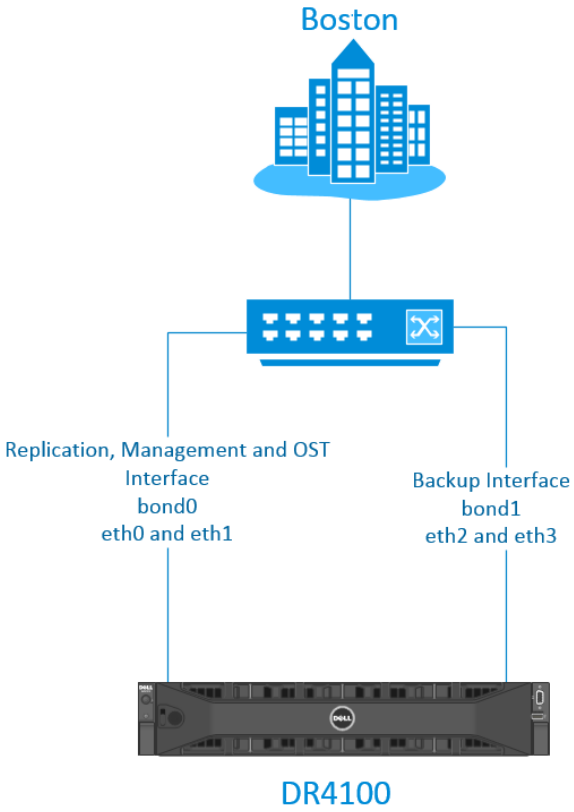

Note: Make sure all 1G interfaces are plugged in and can have an IP address via DHCP. If there is no DHCP server available on the 1G network, then once the system comes up after step-2, a static IP has to be assigned to the 1G bonded interface and a 'network --restart' command should be executed for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

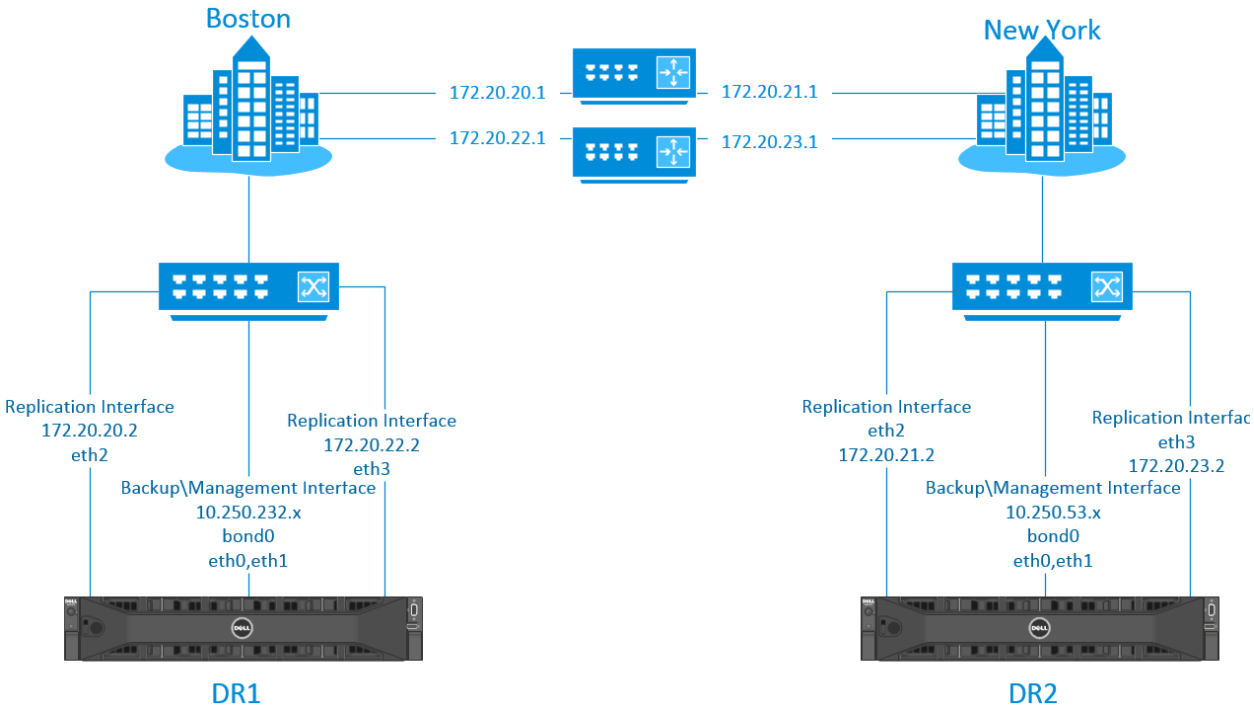

In the default configuration of the DR, the default route for the system is always is set on the network which belongs to bond0. In this example for both the DRs the default route points to the gateway on the 10.250.232.x network. This allows for administration of the DR from anywhere on the network.

Example route:

```
[xxxxx@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
```

Note: The above is an example of a default route. The command above is not accessible to customers.

In order for the two DRs 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 eth 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
administrator@DR2 > network --show
```

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 DR4100 appliance located at his Seattle site and two DR4100s 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
- Lannsing1: 2x 10GB interfaces in bond0 and 2x 1GB interfaces in bond1
- Lansing2: 2x 10GB interfaces in bond0

Figure 6. Scenario 4 topology

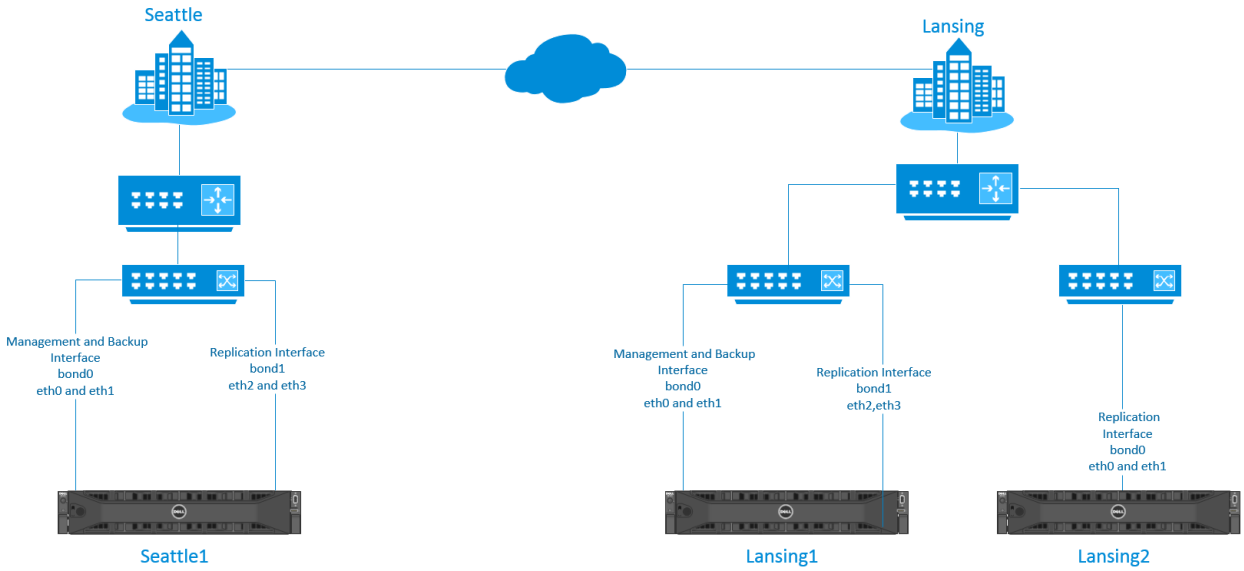

1. On Lansing1 create a bond on the 1GB ports using the following command:

   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 flowing 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 upgrade to the 3.2.

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

- Upgrade the DR to firmware to 3.2.
- Set the media server to use the two interfaces of the DR appliance.

The DR4100 will be configured as follows:

- 1G interfaces are not exposed pre 3.2 firmware upgrade
- Bond0 will contain 2x10G interfaces
- 2x1G interfaces not bonded

Figure 7. Scenario 5 topology

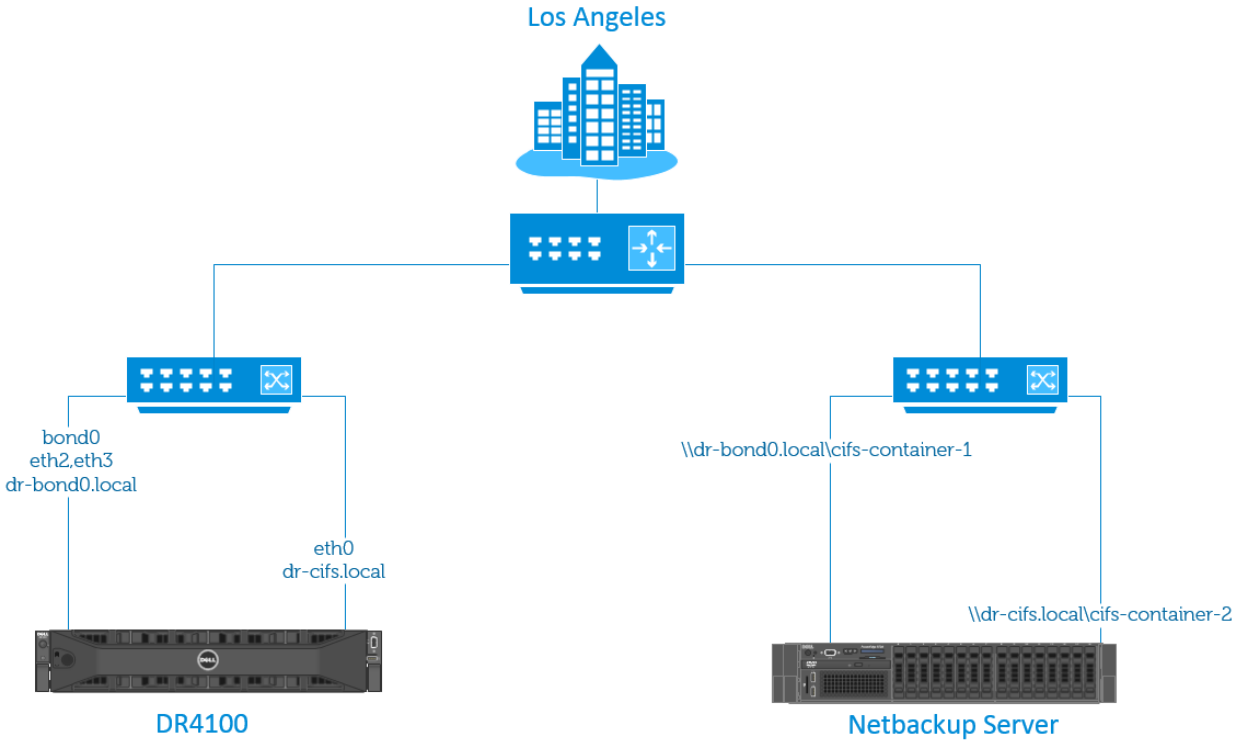

1. Upgrade the DR appliance to the latest 3.2 firmware.
2. After the upgrade, two additional 1 GB interfaces will appear. The 10GB interfaces will be bonded as bond0. Use the following command to view the available network interfaces:

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

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

   network -- restart

   Example

   ```
   administrator@DR> network --delete -- member eth0
   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

   Example

   ```
   administrator@DR> network --create_interface  --nwif eth0  --static_ip x.x.x.x  --netmask
   255.255.255.0 --name  dr-cifs.local  --restart
   ```

   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

   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 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 using the interface's DNS name (dr-cifs.local).

> **Note**: If complete Netbackup setup and configurations steps are needed, refer to the following document: NetBackup Setup Guide.

2. Create a storage unit on the Media server with the UNC path to the DR's CIFS (or NFS if using an NFS Media server) container. Use the DNS name of the newly created interface in the UNC path.

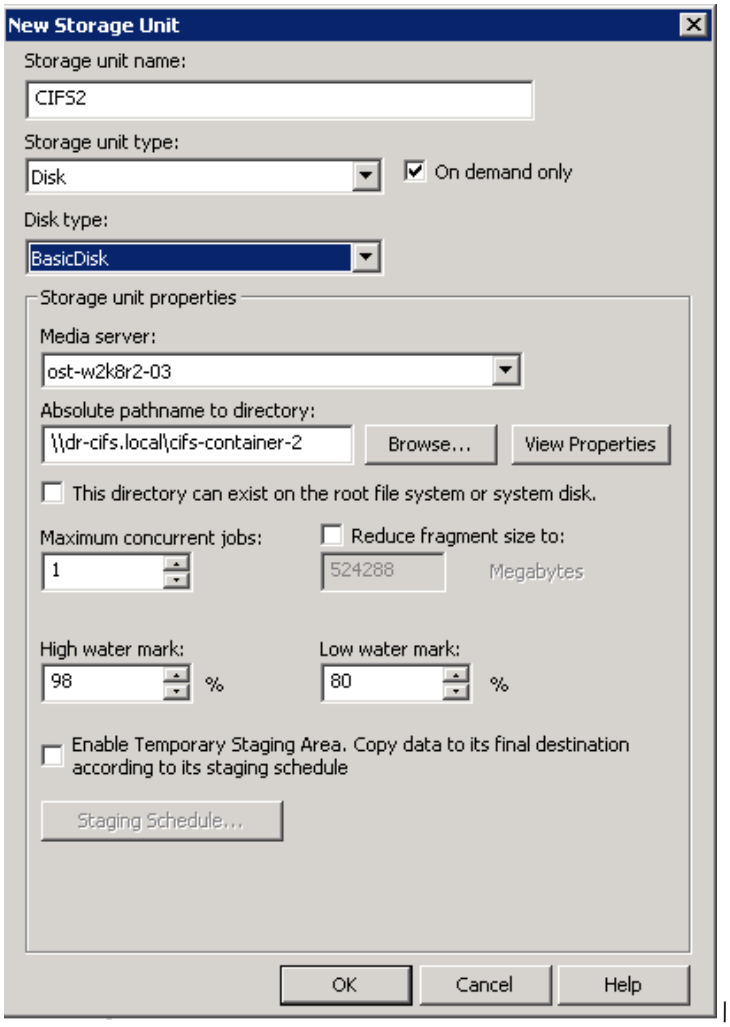

3. Create a second storage unit with the UNC path to the DR's second container. This time use the bond0's name in the UNC path

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

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

Troubleshooting

Follow the steps bellow to troubleshoot connectivity problems between source and target DRs.

a) Issue 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!
```

b) Make sure the management interface is up and reachable by using the following command:

telnet<manamgent_ip>

Cleaner

The DR appliance cleaner is process that is configured to run efficiently and effectively out of the box with no tuning or adjusting required. Only in extreme cases will the DR cleaner possibly may need calibration. This document is intended only for DR appliances which are exposed to the following extreme cases:

- Ingesting or deleting 100TB's or more per week
- Ingest occurs 24x7
- When poor savings is experienced or reported
- Full cleaner pass not finishing once per week

The DR Cleaner best practices guide can be found here.

Other Resources

Dell Support

http://support.dell.com

Dell TecCenter

http://en.community.dell.com/techcenter/