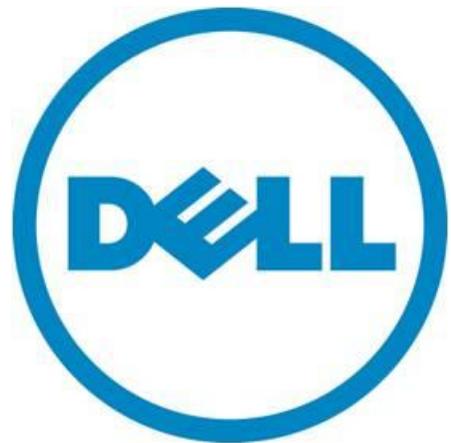


Dell DX Object Storage Platform- Advanced Storage Node Configuration

A Dell Technical White Paper

Dell

Storage Engineering



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.

© 2010 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

Dell, the *DELL* logo, and the *DELL* badge are trademarks of Dell Inc. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

May 2010

Contents

A Dell Technical White Paper i
Dell i
Storage Engineering i
Introduction 2
Storage Node Configuration File 2
Sub-Cluster Overview 3
Assigning Sub-Cluster Membership..... 4
Archive Storage Node Overview 6
Enabling Archive Mode 7
Additional References 8

Figures

Figure 1: Sub-cluster Configuration..... 4
Figure 2: DX Storage Admin Console..... 5
Figure 3: Node Unassigned to a Sub-Cluster..... 6
Figure 4: Archive Storage Node Configuration 7
Figure 5: DX Storage Admin Console Node Status Page 8

Introduction

The Dell™ DX Object Storage Platform provides configuration of cluster-wide parameters through the web-based management console on the DX Cluster Services Node (CSN). There are additional node level configurations that are not supported through the web-based management console and require advanced knowledge of the node configuration files located on the DX Cluster Services Node. This document provides instructions on configuring the following advanced configurations through the node configuration files:

- Configuration of Sub-clusters
- Configuration of Storage Nodes in archive mode

The target audience for this document is advanced DX Storage administrators and support teams.

Storage Node Configuration File

All cluster-wide Storage Node configurations should be done through the CSN Console. The configuration parameters are stored in a single configuration files (`/var/opt/caring/NetBootMGR/config/base/common.cfg`) that is automatically generated. DX Storage also allows Storage Node level configuration. Node level configuration is not supported through the CSN console. It requires manual editing of configuration files located on the CSN. There is an individual Storage Node configuration file for every SN in the cluster that can be modified to define custom configuration for each storage node.

The configuration file can be manually developed and stored on the CSN prior to booting the Storage Node, or the file will be created automatically during the network boot process by default. The file associated with an individual Storage Node will be given the name of the Storage Node machine (MAC) address (`snMACAddress.cfg`). They are located in the `/var/opt/caring/NetBootMGR/config/custom/` directory.

Note: Based on the SN IP address, the associated MAC address can be identified via iDRAC or through the `/var/lib/dhcpd/dhcpd.leases` file on the CSN after Netboot is complete.

The steps for the manually recreating the configuration files prior to booting the Storage Node are as follows:

1. Get the MAC address via iDRAC.
2. Create the `snMACAddress.cfg` file
3. PXE boot the SN.

The steps for the updating the configuration file after booting up the Storage Node are as follows:

1. PXE boot the SN.
2. Identify the SN via the IP to MAC mapping from `dhcpd.leases` file
3. Modify the `snMACAddress.cfg` file for that SN.
4. Reboot the SN.

All configuration files contain name/value pairs for setting configuration options. The format of the name/value pair lines is `Name = value`. Any white spaces within the name/value pair are ignored.

DX Object Storage Platform: Advanced Node Configuration

Blank lines and lines beginning with # are ignored. All fields are case-sensitive. Below is a sample of a node configuration file that was created automatically during the Netboot process.

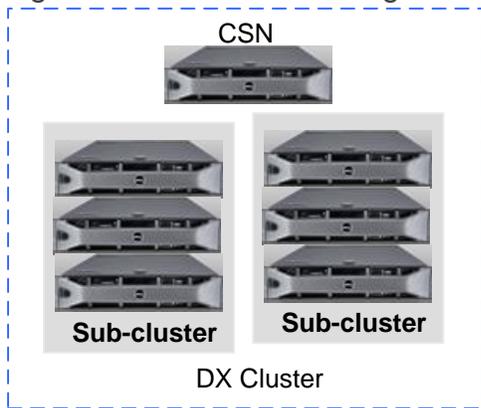
```
#  
# CAStor node Configurataion File  
#  
# Node: sn0026b9348cda  
# MAC: 00-26-b9-34-8c-da  
#  
# ipaddress =  
# netmask =  
# gateway =
```

Sub-Cluster Overview

The DX Object Storage Platform allows a single Storage Cluster to be logically split into two or more separate entities, or sub-clusters. in order to define storage distribution strategies. While the Storage Nodes in a sub-cluster are physically part of the same cluster, managed by a common CSN and communicate on the same IP subnet, they are recognized as logically separate entities by the DX Object Storage software. The existence of a sub-cluster changes the behavior of the cluster so that DX Storage attempts to create the greatest logical spread between an object's replicas by moving them into separate sub-clusters. A sub-cluster configuration is primarily used to enable local area replication. Examples of where sub-clusters are useful are:

- Splitting a cluster based on location (data cabinet, building wing, etc.)
- Grouping nodes based on common infrastructure (network, power)

Figure 1: Sub-cluster Configuration



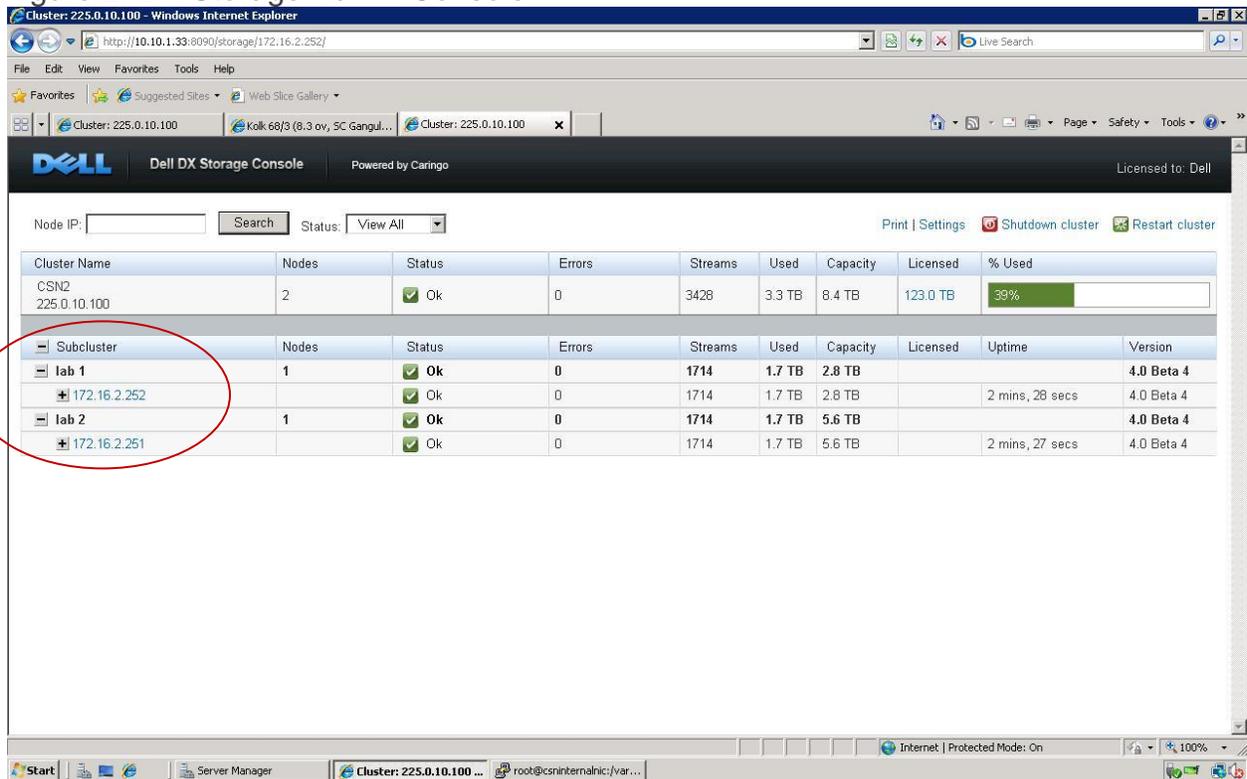
Assigning Sub-Cluster Membership

To configure sub-clusters, the node configuration file for each node must be modified to identify the sub-cluster of which the node is a member. This is done by adding the name/value pairing for the sub-cluster option in the node configuration file for each storage node. The name is “subcluster” and the value is the name of the sub-cluster of which the node is a member. Sub-cluster names cannot contain special characters like quotes and can have a maximum of 16 characters. Below is a sample of a node configuration file that has been modified to assign sub-cluster membership to the node.

```
#  
# CASstor node Configurataion File  
#  
# Node: sn0026b9348cda  
# MAC: 00-26-b9-34-8c-da  
#  
# ipaddress =  
# netmask =  
# gateway =  
subcluster = lab 1
```

After updating all SN configuration files to assign sub-cluster membership, restarting the Storage Nodes via the DX Storage Admin Console is required in order for the sub-cluster membership to take effect. Once the cluster restart is complete, the DX Storage Admin Console displays the defined sub-clusters and the Storage Nodes that have membership to those sub-clusters.

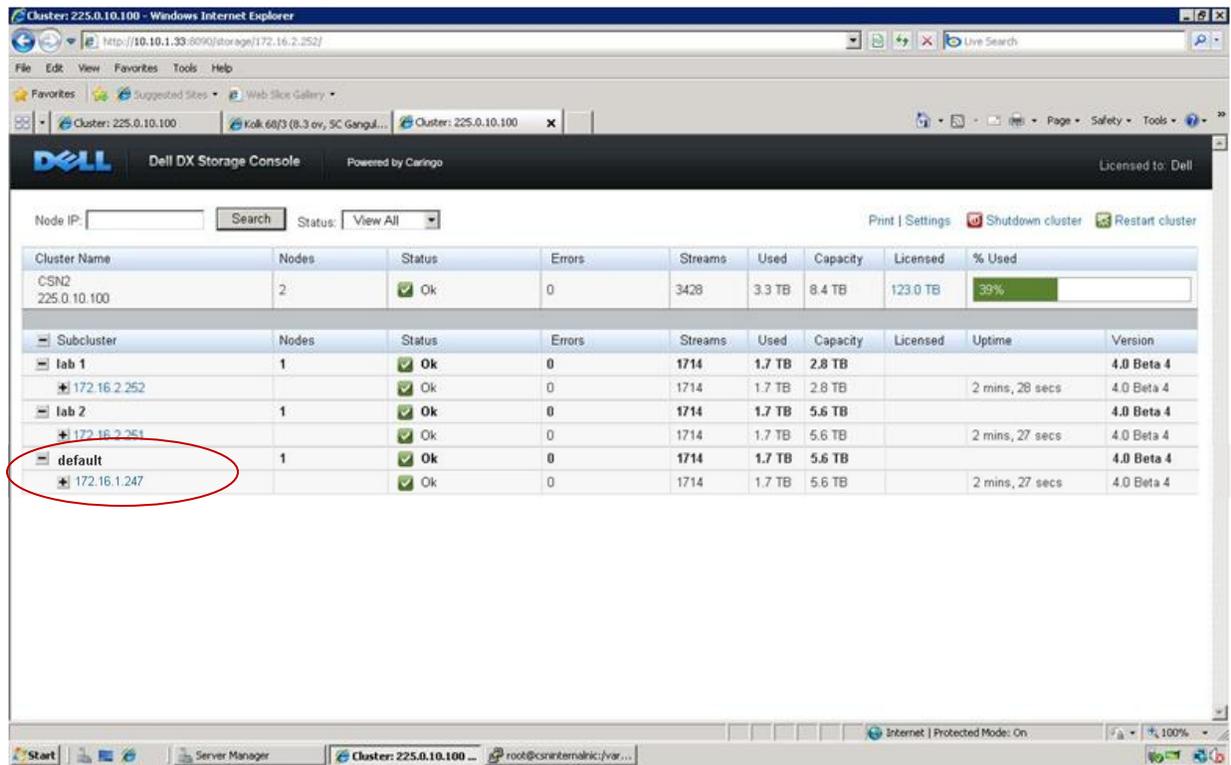
Figure 2: DX Storage Admin Console



When adding new Storage Nodes to a cluster with sub-clusters configured, it is important that the administrator assign sub-cluster membership to the new Storage Node by updating its configuration file. If the SN is allowed to network boot prior to the configuration file being modified, the Storage Node will need to be restarted to pick up the sub-cluster membership information.

If sub-clusters are defined for a DX Storage cluster, any Storage Node that joins the cluster that does not have a sub-cluster membership defined in its configuration file is automatically assigned to a 'default' sub-cluster. In other words, if a new Storage Node is added to a cluster where sub-clusters are configured and the administrator fails to update its configuration file to assign it to a defined sub-cluster, the new node will be added to the default cluster.

Figure 3: Node Unassigned to a Sub-Cluster

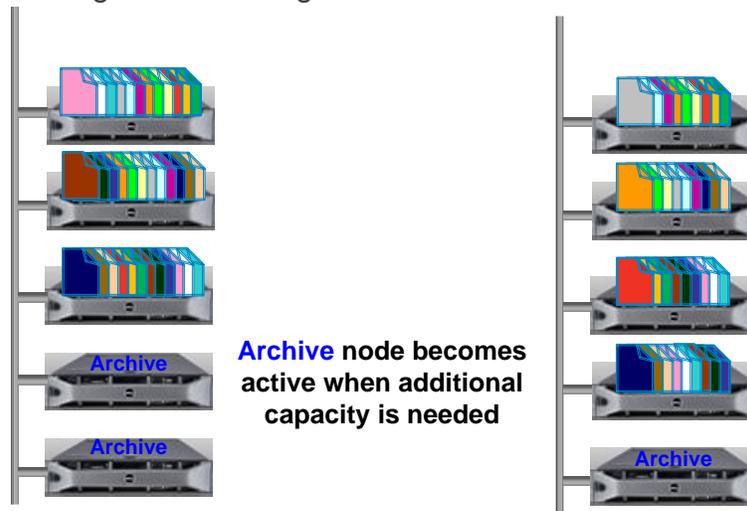


Archive Storage Node Overview

The DX Object Storage Platform allows an administrator to designate a new or empty Storage Node as an archive Storage Node that remains in idle / low-power state without participating in cluster activity until its capacity is needed. This enables administrators to have additional capacity available without paying for the associated power costs of the additional SN.

DX Object Storage automatically load balances across all Storage Nodes in the cluster in order to maximize capacity and throughput. In adaptive power savings mode, a Storage Node does not go idle as long as it is servicing incoming application or cluster requests, depriving an administrator of potential power savings of a Storage Node that is not really needed by the cluster yet. Enabling archive mode for a new Storage Node will prevent objects from being stored on the SN until the additional capacity is needed. This is accomplished by raising the bidding on a Storage Node in archive mode so that it does not win a bid until other nodes become full and the archived Storage Node starts winning the bidding process and stores new content.

Figure 4: Archive Storage Node Configuration



Enabling Archive Mode

To configure a Storage Node for archive mode, the Storage Node configuration file must be updated to enable archive mode. This is done by adding the name/value pairing for the archive mode option in the SN configuration file. The option name is “archiveMode” and the value is 1 to enable and 0 to disable. If the “archiveMode” name field is not specified in the configuration file, then the Storage Node operates in normal mode and objects are balanced across the Storage Node. Below is a sample of a SN configuration file that has been modified to enable archive mode.

```
#  
# CASstor node Configurataion File  
#  
# Node: sn0026b9348cda  
# MAC: 00-26-b9-34-8c-da  
#  
# ipaddress =  
# netmask =  
# gateway =  
archiveMode = 1
```

If the Storage Node configuration file is updated to enable archiveMode after the SN has boot up, the Storage Node must be restarted for the configuration to take effect. The DX Storage Admin Console Node Status Page indicates the Storage Node configuration for archive mode by displaying “archiveMode True” in the expandable Node Configuration section when the SN is configured to enable archive mode. The archive mode setting is a permanent state. The node will continue to adjust its bidding so that other Storage Nodes in the cluster that are not a full capacity will win the bids and take on new content.

