

Dell DX Object Storage Platform-Remote Replication 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
Asynchronous Replication Overview	2
Configuration of the Publisher	4
Configuration of the Replicator	5
Additional References	7

Figures

Figure 1: Asynchronous Replication between Clusters	2
Figure 2: Disaster Recovery Configuration	3
Figure 3: Mirrored Configuration	3
Figure 4: Multi-site DR Roll-up	3
Figure 5: Distributed Configuration	4
Figure 6: Publisher Console	5
Figure 7: Replicator Console	6

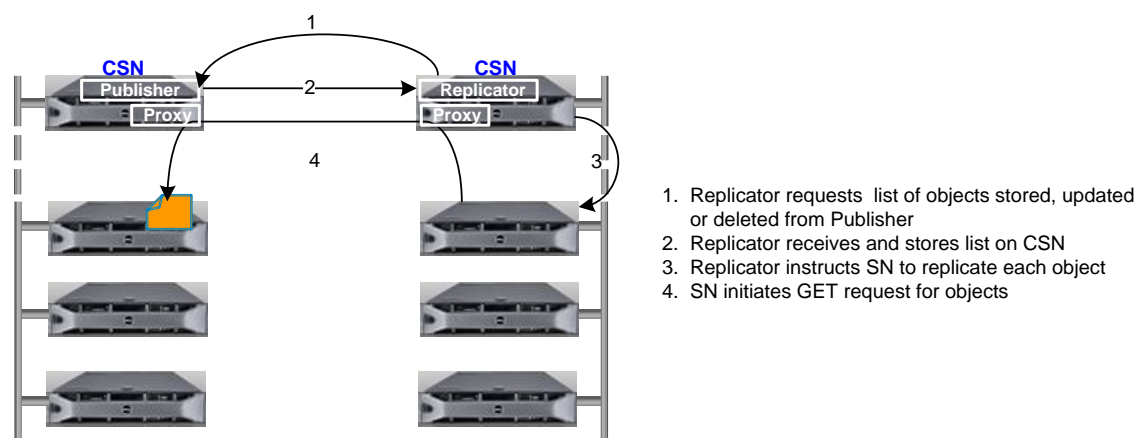
Introduction

The Dell™ DX Object Storage Platform allows data to be asynchronously replicated from one DX Object Storage Cluster to another over a standard IP network. Replication is managed by the DX Content Router (CR) application located on the DX Cluster Services Node (CSN). The CR identifies data to be replicated based on filter rules defined by the DX Administrator. The rules are compared to meta-data stored with the objects to determine which objects are to be replicated to a remote cluster. These rules allow customized data distribution between multiple clusters. Most of the common replication topologies are achievable through the web-based management console located on the CSN. However, more complex replication topologies require advanced configuration through configuration files located on the CSN. The target audience for this document is advanced DX Storage administrators and support teams. It is meant to provide instructions on advanced configuration of asynchronous remote replication through these configuration files.

Asynchronous Replication Overview

The CSN provides replication services through the CR application. It also provides a SCSP proxy interface to the DX Storage Nodes (SN) that allows communication between DX Storage Clusters. This interface accepts operations from applications and re-directs them to the DX Storage Nodes in the designated cluster. Asynchronous replication between clusters utilizes the proxy interface to transport the objects between clusters. The CR application provides two services. The first service is called Publisher and it is responsible for identifying objects to be replicated based on replication rules set by DX Administrator and publishing a list of those objects for replication. The second service provided by CR is called Replicator. The Replicator subscribes to a Publisher, retrieves the list of objects identified for replication and replicates the objects between clusters. The cluster where data is originated runs the publisher service. The cluster where the data is to be replicated runs the replicator service.

Figure 1: Asynchronous Replication between Clusters



Replication can be configured to support many customized configurations. Replication can be one-way or two-way. A configuration using one-way replication, from primary to secondary, recognizes a disaster recovery configuration. Each CSN can run both Publisher and Replicator services enabling two-way replication for a mirrored configuration. The Publisher service supports multiple channels that correspond to distinct filter rules, each of which can each be subscribed to by multiple replicators. A single Replicator can subscribe to multiple Publishers to support a DR roll-up scenario or Replicators

can be configured on multiple CSNs across different clusters to support multi-site disaster recovery. Topologies can also be configured to enable a unique set of objects to be replicated to multiple remote clusters to support a distributed replication model.

Figure 2: Disaster Recovery Configuration




Figure 3: Mirrored Configuration




Figure 4: Multi-site DR Roll-up


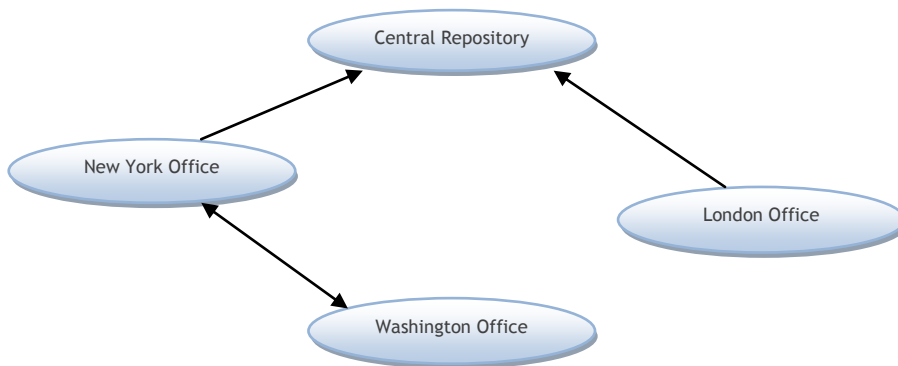


Figure 5: Distributed Configuration

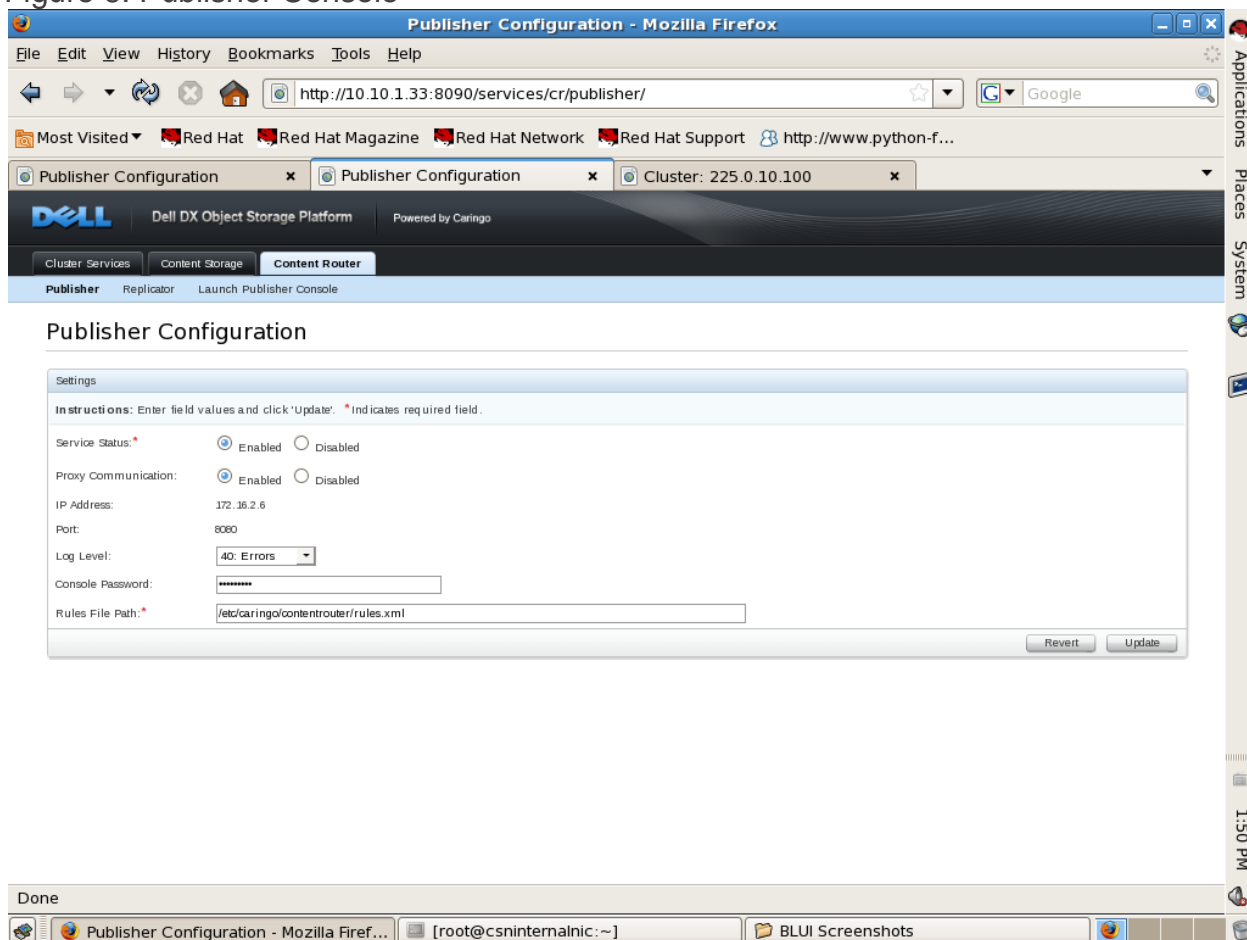


Replication between clusters can be easily customized to support customer needs. The proxy on the CSN is configured during the bootstrap process and requires no additional configuration to support remote replication. In order to configure remote replication, Publisher and Replicator services must be configured.

Configuration of the Publisher

The publisher can be configured through the DX Cluster Services Node web-based management console under the Content Router->Publisher Tab. The set of rules used to control what content is published to each subscriber / replicator, is defined in a rules file (rules.xml) located on the CSN. The rules file can define a unique rules set for multiple subscribers. A default rules file, /etc/caringo/contentrouter/rules.xml, is set to replicate all objects in the cluster by default. This default configuration enables one-way replication of all objects to a remote publisher. This file can be modified to support more advanced filtering. Refer to the DX Content Router Setup and Configuration Guide for detailed instructions on configuring customized rule sets.

Figure 6: Publisher Console



Configuration of the Replicator

The Replicator can be configured through the DX Cluster Services Node web-based management console under the Content Router->Replicator Tab. The set of rules used to control what content filtered for publishing to subscriber/replicators, is defined in a rules file (rules.xml) located on the source cluster where the Publisher is running. The following settings can be configured:

Service Status: Enables / Disables replication services

Log Level: The level of logging verbosity

Subscription Check Interval: Time in seconds between check for new UUIDs. Value can be between 5 and 3600 seconds.

Online After: Time in seconds before the Publisher displays replicator as offline in the Publisher console. Value is a minimum of 60 seconds.

Error Offline After: Time in seconds before the publisher logs a critical error message if a subscriber has not been heard from. Value is a minimum of 60 seconds.

DX Object Storage Platform: Remote Replication Configuration

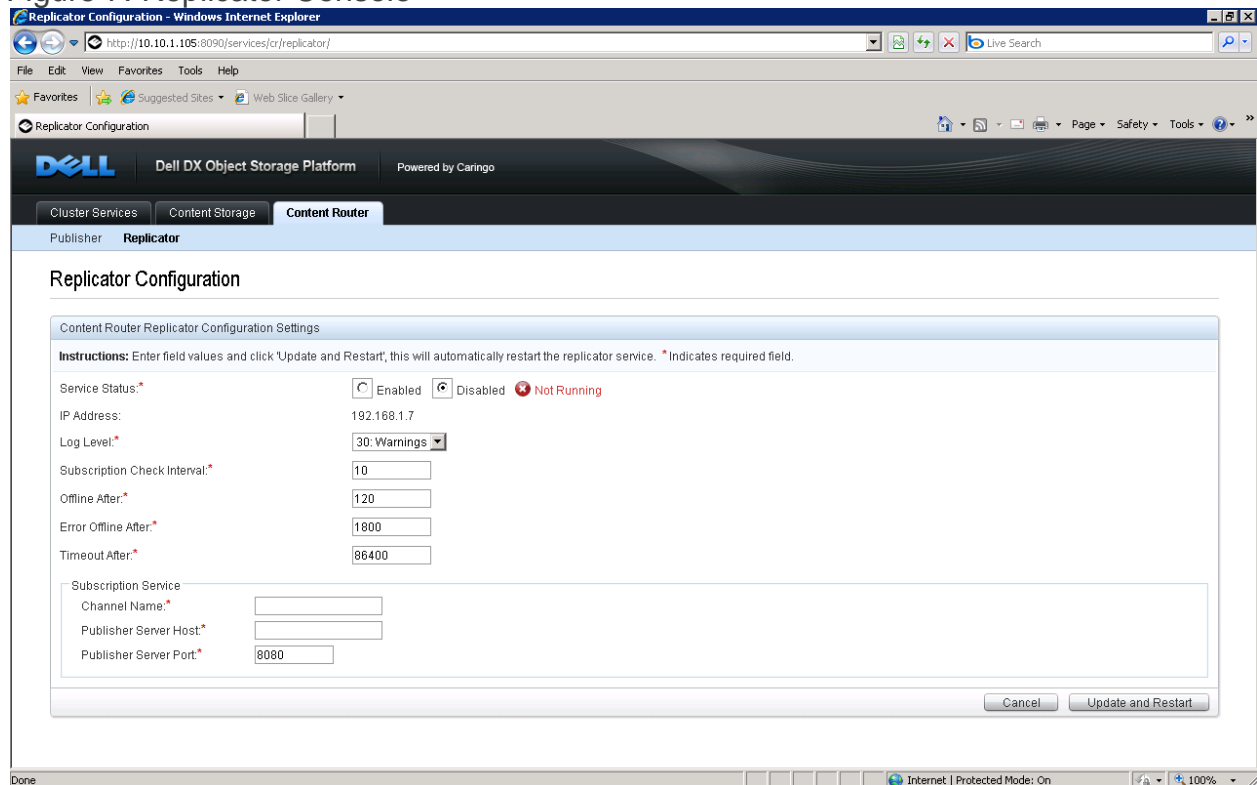
Timeout After: Time in seconds before the publisher will terminate a subscriber if it has not been heard from. Value is a minimum of 60 seconds.

Subscription Service Channel Name: The name of the channel the replicator subscribes to. This must match the subscription name in the remote publisher rules.xml file. The default subscription name in the rules file is “ReplicateAll”, which causes all objects written to the local cluster to be replicated to the remote cluster.

Subscription Service Publisher Server Host: The IP address or hostname of the Publisher to which the Replicator subscribes. This is the external IP address for the remote CSN.

Subscription Service Publisher Server Port: The port on which the publisher publishes UUIDs. Default is port number is 8080.

Figure 7: Replicator Console



Configuration of multiple subscribers for a replicator, as required for the multi-site disaster recovery configuration, requires advanced configuration through the replicator configuration file, /etc/caringo/contentrouter/replicator.cfg. This is done by adding the name/value pairing for the subscription option in the replicator configuration file. The name field is “subscribeTo” and the value is one or more publisher IPs with a channel name and a port number of the publisher service. Below is a sample of a replicator configuration file that has been modified to replicate from 2 source clusters / publishers.

DX Object Storage Platform: Remote Replication Configuration

```
###  
# Replicator settings  
# Parameter set written at 08-Apr-10 18:27  
#  
consolePort = 7091  
# publisher generates syslog error message after this many seconds  
errOfflineAfter = 1800  
group = 225.0.10.100  
ipaddress = replicatornic  
loghost = localhost  
# Log level  
loglevel = 40  
# publisher considers replicator offline after this many seconds  
offlineAfter = 120  
storageDir = /u01/var/opt/caringo/contentrouter/replicator  
# subscriptions to query. format is name@host:port,name@host:port...  
subscribeTo = ReplicateAll@192.168.1.47:8080,ReplicateAll@192.169.1.43:8080  
# interval between subscription checks, in seconds  
subscriptionCheckInterval = 10  
# publisher will terminate enumerator after this many seconds  
timeoutInterval = 86400
```

NOTE: The multi-site replication is not supported via the CSN Replicator Configuration Console and hence upon manual modification of replicator.cfg file the CSN Replicator Configuration Console may show invalid settings.

Additional References

The following documents can be found on www.support.dell.com and can be referenced for additional information on advanced remote replication configurations.

Document	Description
Content Router Setup and Configuration Guide	Covers the steps required to configure Content Router and the administration actions necessary to monitor remote cluster replication.
Cluster Services Node Installation and Configuration Guide	Provides information on configuration and administration tasks for deploying and monitoring a CSN through the CSN Management Console.