

Contents

Replication Introduction	1
Replication Process Overview	2
Configuring Replication	2
Step 1: Create Source, Target and Blockpool File Systems.....	3
Step 2: Setting up the Blockpool.....	5
Step 3: Creating Replication Targets..	6
Step 4: Create a Replication Storage Policy	7
Configuration Steps Summary.....	11

StorNext 4.1.1 Replication Quick Reference

This document summarizes the steps involved in configuring replication. For more information about the specific steps and replication in general, see the Replication chapter and the Replication appendix in the *StorNext User's Guide*.

Replication Introduction

StorNext Replication makes a copy of a source directory and sends the information to one or more target directories. The target directories may be on other host machines, or may be on the same host as the source directory.

Replication behavior is defined by a *Replication/Deduplication Policy*. (The other type of StorNext policy is a *Storage Manager Policy*, which governs how StorNext Storage Manager works).

Here are some important facts about StorNext Replication/Deduplication policies.

- A replication/deduplication policy exists on only one SNFS file system. For example, a policy in a file system called `/stornext/sn1` can be used only to replicate directories in that file system. A separate policy would be needed to replicate directories from file system `/stornext/sn2`.
- If a replication/deduplication policy will be used in any file system on a machine, you must configure a *blockpool* for that machine. The blockpool for a machine contains data (called blocklets) if the Deduplication feature is used, but the blockpool must be configured for replication use even if you do not use deduplication.
- A policy may be applied to one directory or more than one directory in the file system.

Replication Process Overview

The actual replication process occurs in two stages:

- 1 **Data Movement Stage:** In this stage StorNext moves the data for files from the source file system to the target file system. Data movement occurs continuously as files are created or modified in a source directory.

Data movement occurs in one of the following two ways.

- **Deduplicated Data:** If deduplication has been enabled for the policy that controls the source directory, deduplicated data moves from the source host machine to the target host. With deduplication enabled there may be less data moved than if the entire file were copied. This is because for deduplicated replication, only the unique deduplicated segments need to be copied.
- **Non-deduplicated Data:** If deduplication is not enabled for the policy that controls the source directory, the entire file is copied from the source directory to the target host. The entire file is copied whenever a file is created or modified.

When data movement is in progress or even after it has just completed, the replicated files may not be visible yet in the target file system's directories. Replicated files become visible in stage 2.

- 2 **File System Namespace Realization Stage:** In this stage StorNext enumerates all the files in the source directory and recreates the file name and subdirectory structure (the *namespace*) on the target file system. Unlike in the Data Movement Stage, this stage happens only at scheduled times, or when namespace realization is manually initiated by the administrator.

Configuring Replication

This section describes how to configure simple one-to-one replication from one source directory on one file system to one target file system. The source and target StorNext server computers can be the same machine, standalone servers, or High Availability (HA) redundant servers. When replication-target file systems are on an HA Cluster, it is best to convert the cluster to HA before configuring replication source policies that point to them. This allows the use of the virtual IP (VIP), which is required for HA configurations.

Before you begin configuring, make sure you have the Replication and/or Deduplication licenses required for these features. If you are using an HA configuration, basic StorNext single-server or HA Clusters should already be set up. (For more information, see the Configuration Wizard chapter in the *StorNext User's Guide*.)

These instructions assume you are using the StorNext Configuration Wizard and have already completed its first three steps: **Welcome**, **Licenses**, and **Name Servers**.

Step 1: Create Source, Target and Blockpool File Systems


After you complete the first three Configuration Wizard steps, the first replication step is to create file systems: the source and target file systems and the blockpool file system(s), you plan to use.

Note: Although StorNext supports replicating from multiple source hosts and file systems to multiple target hosts and file systems, for simplicity this procedure describes how to replicate between one source and one target file system on the same host.

Creating Source File Systems


- 1 If you have not already done so, launch the StorNext Configuration Wizard and proceed through **Welcome**, **License** and **Name Servers** steps to the **File System** step.
- 2 The **Configuration > File Systems** screen appears.
- 3 On the **Configuration > File System** screen, click **New**. The **Configuration > File Systems > New Screen** appears.

Figure 1 Configuration > File Systems > New Screen



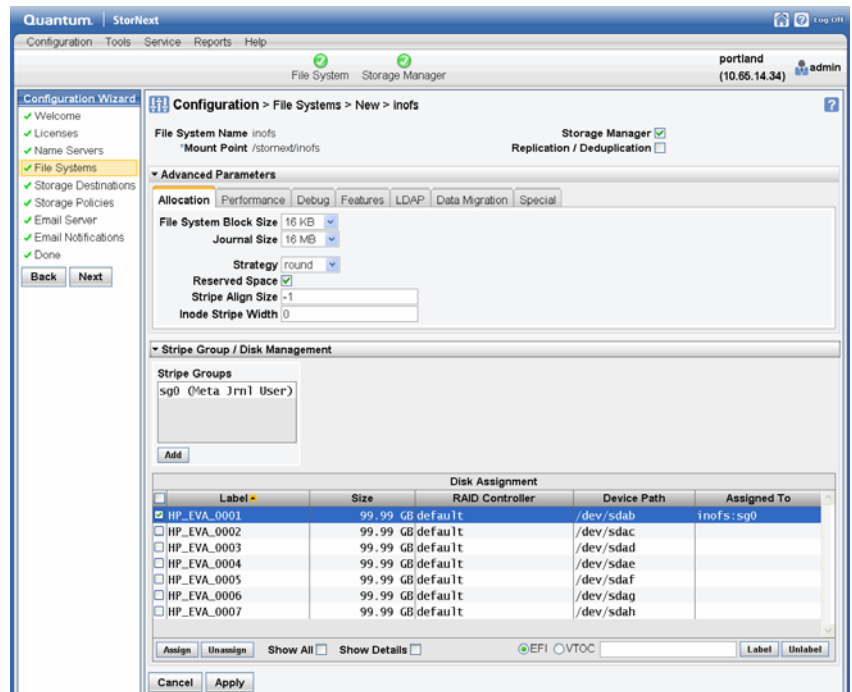
- 4 At the **File System Name** field enter the name of a file system to be used as a replication source. A default mount-point path automatically appears but you can change this mount point if you wish.
- 5 Choose the **Replication/Deduplication** option. A warning message alerts you that "A blockpool has not been created." Disregard this message for now because you will create the blockpool file system in [Step 2: Setting up the Blockpool](#).
- 6 Make sure the **Generate** option is selected, and then click **Continue** to proceed.

Figure 2 Configuration > File System > New Screen 2



- 7 Select a set of LUNs for the file system, and then click **Assign**.
- 8 Click **Continue**.

Figure 3 Configuration > File System > New Screen 3



- 9 If desired, click the arrows beside the **Advanced Parameters** and **Stripe Group/Disk Management** headings to view information.
- 10 Click **Apply** to save the new file system. (For more information about creating file systems, see the Configuration Wizard chapter in the *StorNext User's Guide*.)

Creating a Target File System and Blockpool File System

- 1 Repeat the process (steps 1 - 10) and create the file system you intend to use as a target for replication on this same server.
- 2 Configure another file system for the Blockpool that has neither Data Migration nor Replication/Deduplication enabled.

Step 2: Setting up the Blockpool

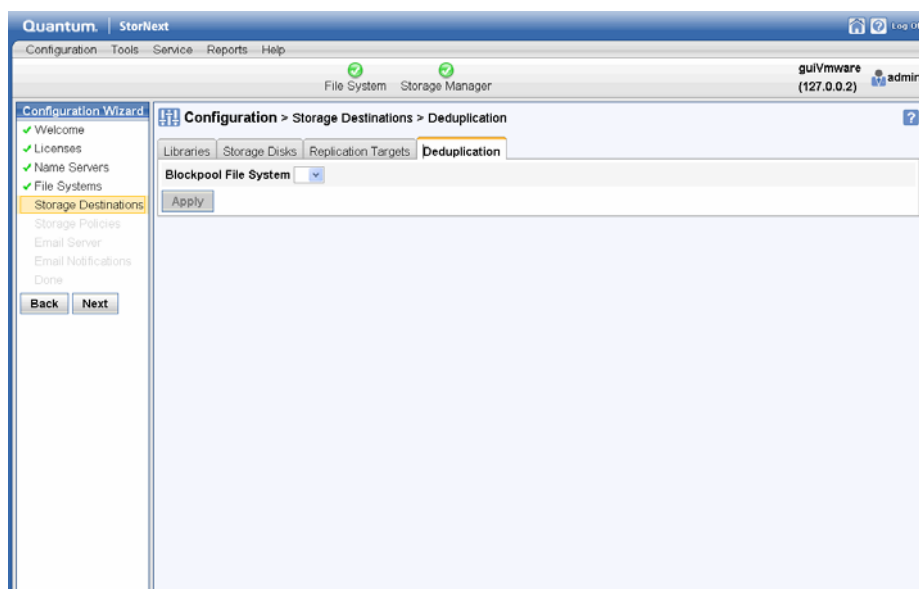
In this step you will set up the blockpool on the blockpool file system you just created in the previous step.

- 1 Choose the StorNext Configuration Wizard's **Storage Destinations** task. The **Configuration > Storage Destinations** screen appears.

There are four tabs on this screen: **Library**, **Storage Disk**, **Replication Targets**, and **Deduplication**. When configuring replication we are concerned with the **Replication Targets** and **Deduplication** tabs. (The deduplication infrastructure is used to handle the transfer of file data for the replication feature, so it must be configured even when the deduplication feature is not used.)

- 2 Click the **Deduplication** tab. The **Configuration > Storage Destinations > Deduplication** Screen appears.

Figure 4 Configuration > Storage Destinations > Deduplication Screen (Blockpool)



- 3 Click the **Deduplication** tab. This tab has only one field called **Blockpool Host File System**. At this field select from the dropdown list the file system to use for the blockpool. (This is the file system you created in the previous step.)

Note: Once applied, the blockpool location cannot be moved to another file system. Be certain of the blockpool location before you continue.


- 4 After you select the blockpool file system, click **Apply**. A background job is started to create the blockpool. This job typically finishes in a few minutes depending on the size parameter of your deduplication license.

Step 3: Creating Replication Targets

In this step you will specify the actual targets to which you want replicated data sent. (Namespace realization will also occur on these targets.)

- 1 Click the **Replication Targets** tab. The **Configuration > Storage Destinations > Replication Targets** Screen appears.
- 2 Click **Add Host**.

Figure 5 Storage Destinations > Replication Targets Screen



- 3 At the **Hostname or IP** field, enter the host name or its IP address. If the target is an HA cluster, the address should be the VIP for that cluster.
- 4 Click **Scan Host** to populate the **Mount Point** box with appropriate file systems which are configured for replication/deduplication.
- 5 Select the file system you created for use as the target in [Step 1: Create Source, Target and Blockpool File Systems](#), and then click **Add**.
- 6 Click **Apply**. At this point you should see your file system listed as a replication target.

Note: If you were adding additional replication targets, you would repeat steps 3 - 6 to add additional hosts and file systems.


Step 4: Create a Replication Storage Policy

The next step in configuring replication is to create a replication storage policy. This policy contains the replication "rules" specific to your replication source and target file systems. You must create a replication policy for the source directory and enable inbound replication for the target file system.

Creating the Source Directory Replication Policy

- 1 Choose the StorNext Configuration Wizard's **Storage Policy** task. The **Configuration > Storage Policies** Screen appears.
- 2 Click **New**. The **Configuration > Storage Policies > New** Screen appears.

Figure 6 Configuration > Storage Policies > New Screen

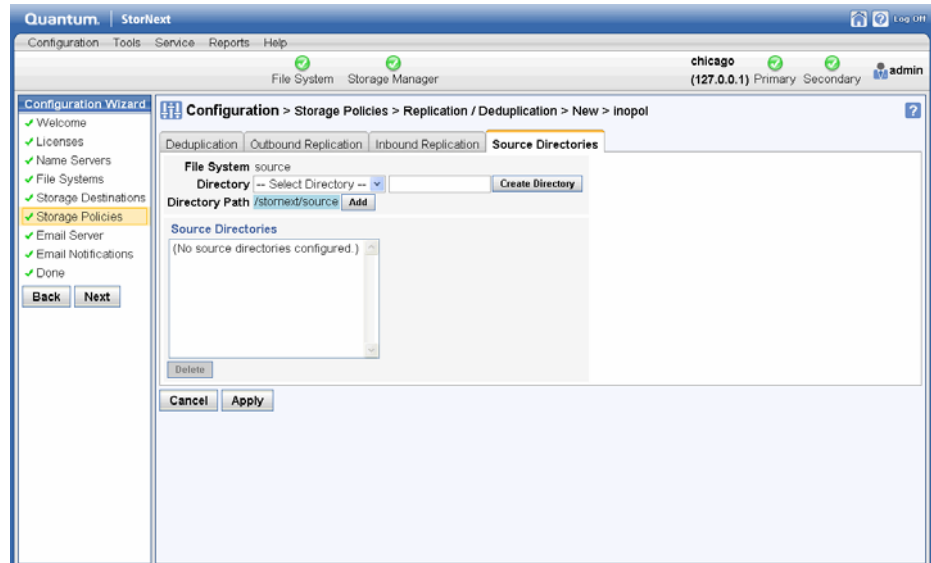


- 3 Enter the following fields:
 - **Policy Name:** The name of the new policy you are creating
 - **Policy Type:** Choose **Replication/Deduplication** to create a replication storage policy.
 - **File System:** Choose the source file system from the dropdown list.
 - Click **Configure**.

Choose the Source File System

- 1 After you click **Configure**, the screen for configuring a replication/deduplication storage policy appears.
- 2 Click the **Source Directories** tab.

Figure 7 Configuration > Storage Policies > New / Source Directories Screen



- 3 At the **Directory** field, enter the name of a directory you want to create for the replication source directory.
- 4 Click **Create Directory**, and then click **Add** to add the new directory to the policy.

Enter Outbound Replication Information

- 1 Click the **Outbound Replication** tab.


Figure 8 Storage Policies > New > Outbound Replication Tab



- 2 At the **Outbound Replication** field, enable outbound replication (going out from the source) by clicking the area to the right of the field so that **On** is displayed.

- 3 Click **Inherit** to the right of the **Copies to Keep on Target** field, and then select **1** from the dropdown list.
- 4 (Optional) To create a new replication schedule, in the **Replication Schedules** box, click **New**. Additional fields appear where you can create a replication schedule.

Figure 9 Outbound Replication
Tab > Replication Schedule




- 5 Under the heading **Select Available Targets**, select the target file system on the target server.
- 6 Create a schedule by making a selection in every column. If you select none of the schedule columns, this creates an unscheduled policy that must be run manually.
- 7 Click **Continue** to complete the schedule and target selections.
- 8 Click **Apply** to finish creating the policy with the options from all of the tabs.
- 9 After a message informs you that the policy was created successfully, click **OK**.

Enter Inbound Replication Information


- 1 On the **Configuration > Storage Policies** screen, select the replication/deduplication policy named "target" for the replication target file system, and then click **Edit**.

Figure 10 Configuration > Storage Policies Screen (Select "target")



- 2 When the **Configuration > Storage Policies > Edit > target** screen appears, click the **Inbound Replication** tab.

Figure 11 Storage Policies > Edit > target > Inbound Replication Tab



- 3 At the **Inbound Replication** field, select **On**.
- 4 Click **Apply** to finish editing the policy with your selected options.

Configuration Steps Summary

The preceding four configuration steps accomplished the following:

- Created a source replication policy and associated a source directory with it.
- Selected a target file system on a target host machine and left the target directory unspecified, which uses the directory name of the source.
- Set a replication schedule that runs every day at midnight.
- Enabled inbound in the target policy.
- Enabled outbound replication in the source policy.

The contents of the source directory (additions and deletions) will now be replicated to the target directory every night. You can test this by running the policy manually at any time on the **Configuration > Storage Policies** screen. (Select the policy you want to test and then click **Run**.)



For assistance, contact the Quantum Customer Support Center:
USA: **800-284-5101 (toll free) or 949-725-2100**
EMEA: **00800-4-782-6886 (toll free) or +49 6131 3241 1164**
APAC: **+800 7826 8887 (toll free) or +603 7953 3010**
Worldwide: <http://www.quantum.com/ServiceandSupport>

Quantum®

Backup. Recovery. Archive. It's What We Do.

©2011 Quantum Corporation. All rights reserved. Quantum, the Quantum logo, and all other logos are registered trademarks of Quantum Corporation or of their respective owners. Protected by Pending and Issued U.S. and Foreign Patents, including U.S. Patent No. 5,990,810.

About Quantum

Quantum Corp. (NYSE:QTM) is the leading global storage company specializing in backup, recovery and archive. Combining focused expertise, customer-driven innovation, and platform independence, Quantum provides a comprehensive range of disk, tape, media and software solutions supported by a world-class sales and service organization. This includes the DXi™-Series, the first disk backup solutions to extend the power of data deduplication and replication across the distributed enterprise. As a long-standing and trusted partner, the company works closely with a broad network of resellers, OEMs and other suppliers to meet customers' evolving data protection needs.