



Technical Report

Sybase IQ Data Warehouse on NetApp Storage: Proof of Concept

Suresh Vundru, NetApp
March 2011 | TR-3905

TABLE OF CONTENTS

1	INTRODUCTION	3
2	BUSINESS CHALLENGE	3
3	THE SOLUTION	4
3.1	PRODUCTION DATA LAYOUT	5
3.2	DATA FLOW MODEL	5
3.3	BACKUP CONFIGURATION	10
3.4	POST BACKUP ACTIONS	12
4	BUSINESS BENEFITS	15
5	REFERENCES	16

LIST OF TABLES

Table 1)	Overview of deployed structure.	6
Table 2)	Volume use and locations.	6
Table 3)	Business benefits.	15

LIST OF FIGURES

Figure 1)	Layout before and after POC.	3
Figure 2)	Primary and DR site infrastructure.	4
Figure 3)	Production data layout.	5
Figure 4)	Data flow between storage systems.	7
Figure 5)	Storage data layout.	8
Figure 6)	Development data layout.	9
Figure 7)	SATA data layout.	10
Figure 8)	Snap Creator workflow.	11
Figure 9)	Process to perform primary backup.	12
Figure 10)	Post backup actions.	13
Figure 11)	Update SnapMirror.	13
Figure 12)	Process to update SnapVault.	14
Figure 13)	Cycle primary Snapshot copies.	15
Figure 14)	Post Snap Creator action.	15

1 INTRODUCTION

This document captures the configuration for a solution that was developed during a customer proof of concept (POC). The solution integrates a Sybase IQ data warehouse with a NetApp® storage system. This document covers the following procedures:

- Define the high-level design for placing and managing Sybase IQ on NetApp storage.
- Perform backup and recovery with NetApp Snapshot® and SnapRestore®.
- Deploy NetApp SnapMirror® to create the disaster recovery site.
- Leverage the NetApp FlexClone® feature to create a dev/test environment based on production data.

2 BUSINESS CHALLENGE

The customer has several databases that support different applications. The largest of these databases is a Sybase IQ data warehouse. The size of this database (~4TB) creates three main issues:

- The amount of time required to back up to or recover from tape takes over 24 hours.
- Development databases are hard to refresh and are smaller than the production database because the configuration is old and the production data is growing rapidly, which results in high storage costs.
- Disaster recovery for this environment is estimated to take 3 to 5 days.

These business drivers combined with the technological benefits that NetApp provides drove the decision to refresh this environment onto new servers and a new storage infrastructure.

Figure 1) Layout before and after POC.

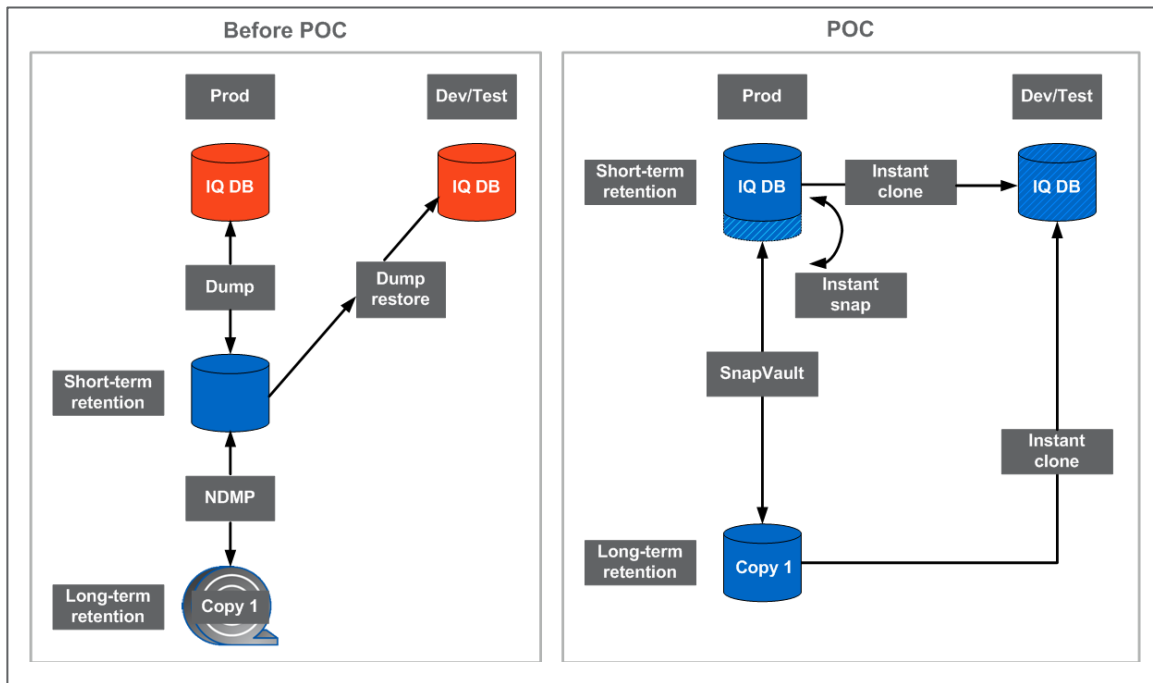


Figure 1 shows the layout of the environment before and after the POC. This report details the benefits and the implementation of the POC; the solution used for placing and managing Sybase IQ on NetApp storage; and the design, the technologies used, and the process for managing the environment.

3 THE SOLUTION

The general strategy for managing the Sybase IQ environment is to leverage the primary storage device to support core IT support functions (backup, restores, and refreshes). From a business support viewpoint, this solution provides the following functionality:

- Supports the primary (production) environment through business as usual (BAU) functions such as data growth
- Provides a disaster recovery (DR) capability with a two-hour recovery point objective (RPO)
- Provides seven years of daily backups
- Allows any of the backup copies to be used as a full-site copy for development or testing
- Provides rapid restores for any of the long-term backups
- Provides storage efficiency with improved performance
- Provides these benefits at a lower cost to the business by using NetApp storage and data management features

The ability to generate consistent Snapshot copies, which is a core NetApp technology, is the key to this project. These Snapshot copies constitute a backup and enable the customer to perform DR, long-term data archival, and a dev/test data refresh.

Figure 2) Primary and DR site infrastructure.

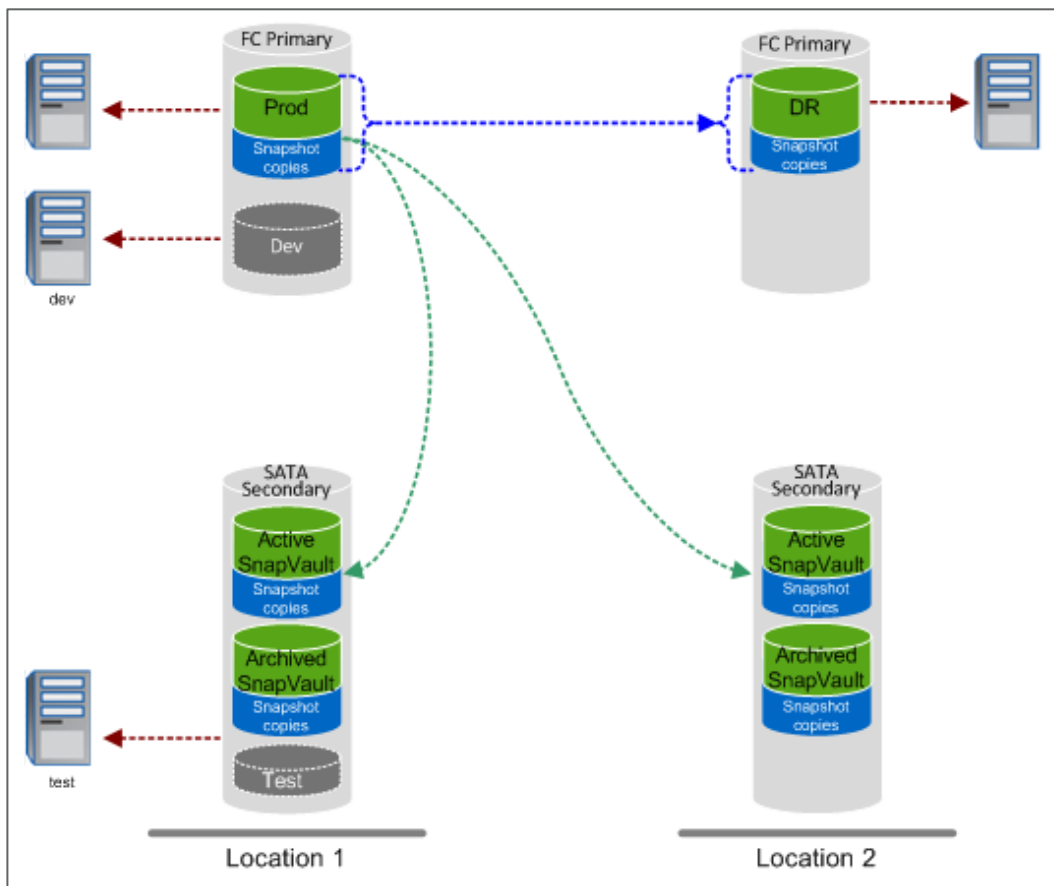


Figure 2 details the primary and DR site infrastructure. The primary site (location 1) consists of production, development, and test environments. The DR site (location 2) has copies of the production

environment datasets available in case of a disaster. The data flow, DR procedure, and technologies used to perform these tasks are described in the following sections.

3.1 PRODUCTION DATA LAYOUT

A very simple volume layout is selected because this is the minimum needed to perform consistent backups as well as to segregate high-block change data that is not required to restore or refresh.

Figure 3) Production data layout.




Figure 3 shows a single large aggregate dedicated for Sybase IQ. Three production volumes are contained within this aggregate (we omitted the development volumes for the sake of simplicity).

- **Volume A.** This is the largest of the volumes in which the database, binaries, configuration files, and data files reside. This volume must be protected because it is required to perform restores. This volume contains multiple 250GB data logical unit numbers (LUNs) that are presented to the host and managed by Veritas™ Volume Manager (VxVM).
- **Volume B.** This volume houses a small dump file, created as part of the backup process. The dump file contains database metadata information as well as transactions that were processed during the backup of volume A. This is required to bring data files within the data volume back to a consistent state.
- **Volume C.** This volume houses data files used for the temporary tablespaces. The data within these files is not needed for recovery; however, Sybase IQ requires the data files to exist and be of the correct size. We keep track of the data files within the dump directory so that they can be recreated in the event of a restore. This volume is not backed up, and Snapshot backups are not created. This volume consists of two 64GB LUNs used by VxVM to create a small disk group to support the re-creation of temporary tablespaces.

This information describes how the three production volumes are placed. Volume A and volume B are protected with Snapshot copies and can be restored when needed.

3.2 DATA FLOW MODEL

The storage allocation is designed around three different volumes. Each volume maps to a specific function and in turn has different backup sequencing requirements. Table 1 provides an overview of the structure being deployed. This table omits cloned volumes that are static allocated storage.

Table 1) Overview of deployed structure.

Environment	File System	VxVM Disk Group	Volume	Mirrored	Vaulted
Production	/sybhm	prod_dg_01	vol10_01	Yes	Yes
Production	/sybhm/eng/svr1				
Production	/svr1_dev				
Production	/svr1_tmp_dev	prod_dg_02	vol10_03	Yes	No
Production	/sybhm_dumps	nfs	vol10_02	Yes	Yes
DR	/sybhm	dr-dg-01	voldr10_04	No	No
DR	/sybhm_dumps	nfs	voldr10_02	No	No
Dev	/sybhm	dev_dg_01	voldev10_04	No	No
Dev	/dev_tmp	dev_dg_02	voldev10_03	No	No
Test	/sybhm	test_dg_01	voltest10_04	No	No
Test	/test_tmp	test_dg_02	voltest10_03	No	No

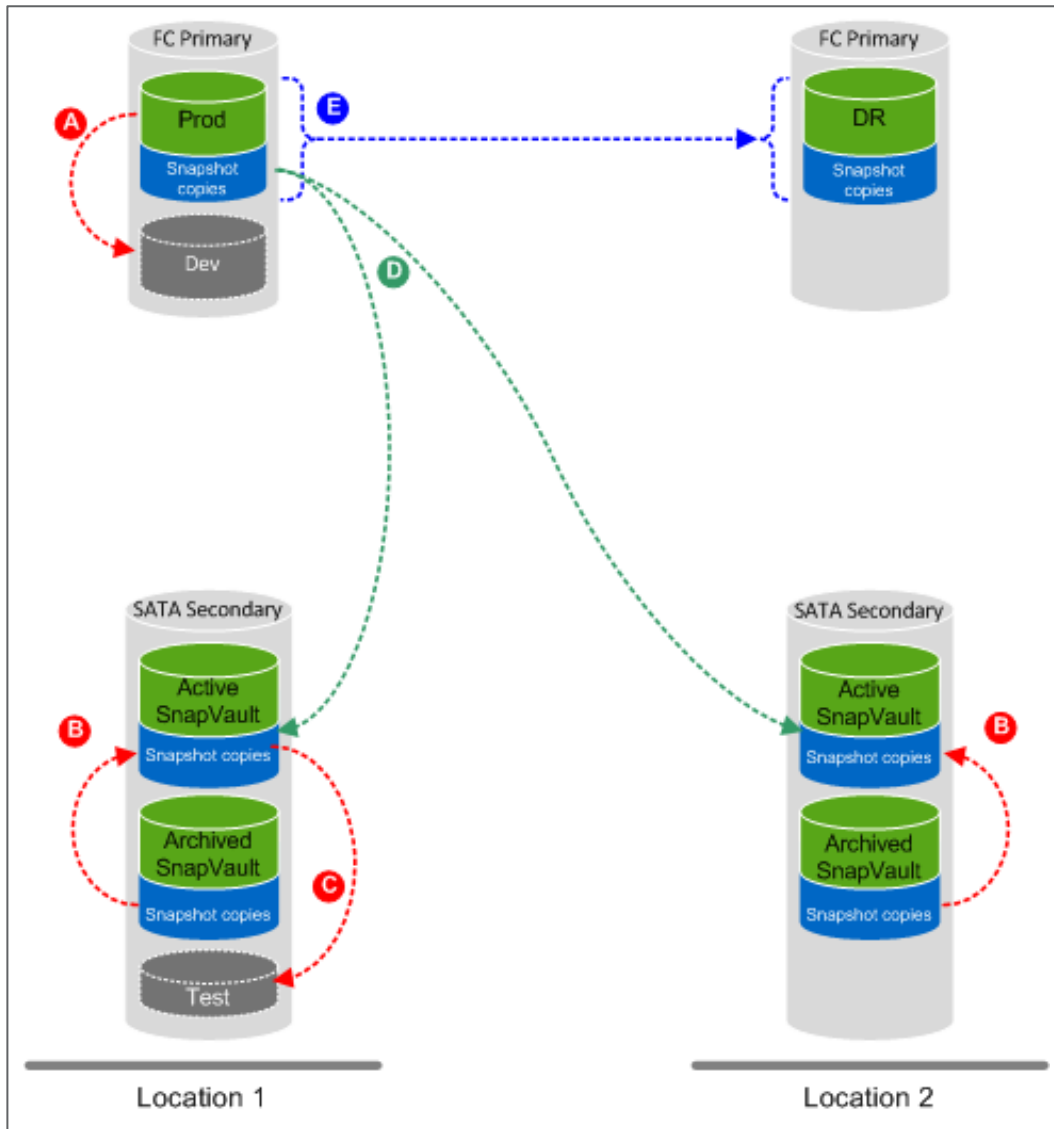
Table 2 details the volumes that are generated from production copies by using FlexClone. These volumes are based on the latest Snapshot backup copy.

Table 2) Volume use and locations.

Environment	File System	VxVM Disk Group	Volume	Clone Source
Dev	/sybhm/eng/dev_svr1	dev_dg_03	voldev10_01	Production; vol10_01
Dev	/svr1_dev			
Dev	/sybhm_dumps	nfs	voldev10_02	Production; vol10_02
Test	/sybhm/eng/test_svr1	test_dg_03	voltest10_01	Backup; vol10_svXX
Test	/svr1_test			
Test	/sybhm_dumps			

Figure 4 shows the data flow between the storage systems.

Figure 4) Data flow between storage systems.



FlexClone, SnapVault[®], and SnapMirror are the three main technologies used. Each technology performs a different function, but all of them use the same source: Snapshot copies created on the primary disk. The following letters, designated in Figure 4, describe the functions and the reason for their use.

- A. A requirement is that one of the development environments must run off fast Fibre Channel (FC) disk. To provide a copy of this data, we use FlexClone based on the latest Snapshot backup, and then we split the clone. Splitting the clone removes the dependency on the parent volume, which addresses the locked Snapshot issue, but it does cause this copy to be full size.
- B. As mentioned previously, retaining seven years of daily backups is a requirement; however, NetApp Data ONTAP[®] has a 255 Snapshot limit per volume. To bypass this issue, when the Snapshot count exceeds 240 Snapshot backups, we create a new FlexClone volume on the SnapVault destination volume. We then continue running the SnapVault updates to the new FlexClone volume. This step avoids the Snapshot limitation of 255 copies in a volume. The use of clones avoids the need to rebaseline the SnapVault relationship each time a new volume is created.
- C. Although one of the development environments is based off the primary FC disk, the other one makes use of the long-term backup to obtain its copy of the database. We create a FlexClone copy from one

of the backup volumes and present this to the development host. Unlike the primary case, we do not split the clone. This allows disk blocks to be shared between the backup and the development environments, but it does not affect the performance in a development environment with small backup windows of less than 10 minutes.

- D. To capture remote disk-based backups, two SnapVault relationships are updated at the end of each backup.
- E. To support DR capability, we use SnapMirror to mirror production volume to the DR volume every hour. This translates to 25 updates every day. Twenty-four of these are crash consistent; that is, no attempt to quiesce the database is made. The 25th transfer is based on a consistent Snapshot copy.

Note: The agreed-on RPO is two hours. Therefore, the frequency of updates is configured as one update per hour, which exceeds the required RPO. Remember that SnapMirror only keeps the latest copy.

DISASTER RECOVERY DATA LAYOUT

The bulk of the data layout for the DR environment is inherited from the production environment by the use of the volume SnapMirror. However, one important change is the addition of a separate volume that is used to provide a preprovisioned and preconfigured `/sybhm` file system. The NetApp view of the storage layout is shown in Figure 5.

Figure 5) Storage data layout.




Figure 5 shows an aggregate containing four volumes. Three of these volumes (translucent) are production copies managed by SnapMirror. The fourth volume, indicated by the letter A in Figure 5, is the standalone volume used to host a DR copy of the `/sybhm` file system. The customer's DBA team requested a separate copy of `/sybhm` to make it possible to preconfigure the environment.

Note: The original production copy of the `/sybhm` file system is mirrored using SnapMirror and is available if needed.

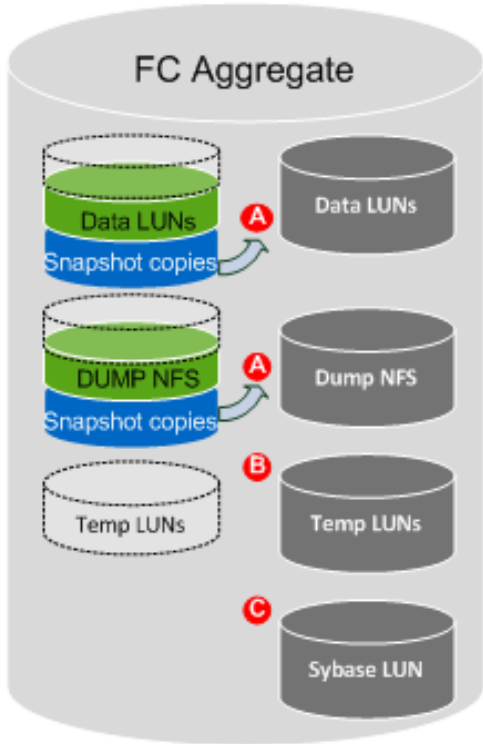
We mounted the NFS Sybase IQ dump volume within the DR environment. The volumes at the primary data center are mirrored to a secondary location by using SnapMirror software; however, this is a read-only copy. When a DR is performed, this volume becomes a read/write copy. Having it premounted makes the DR process slightly faster.

DEVELOPMENT DATA LAYOUT

Similar to the DR data layout, the bulk of the development environment data layout is inherited from the production data layout. The exception is a local `/sybhm` that is provisioned for the local configuration required to support the refresh of the development environment with minimal effort and a local volume, which is used to house the temp data files used by Sybase IQ.

Figure 6 shows how the volumes are used in the development environment. The production volumes are translucent volumes and are only included to provide context for how the development volumes are created.

Figure 6) Development data layout.



As shown in Figure 6, four volumes are presented to the development environment:

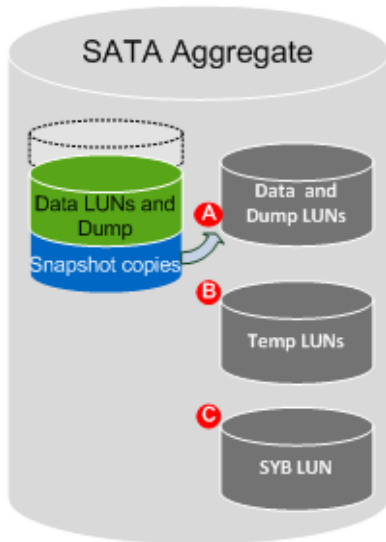
- The data LUN and NFS dump volume are cloned (A) from the production volumes and then split to provide a full copy of the data without locking production volume Snapshot copies.
- The volume containing the temp LUNs (B) is a separate standalone volume that is not refreshed.
- The last volume (C) is a local volume that contains a custom version of the `/sybhm` file system. This version enables the DBA team to maintain an environment that can be rapidly refreshed.

To perform a refresh, we remove storage, destroy the clone, reclone, and re-present. The temp and Sybase volumes (B) and (C) are static and are not refreshed.

TEST DATA LAYOUT

The test environment is refreshed from the production SnapVault backups. A single Snapshot copy on a FAS system with a NearStore[®] license contains all the data required to perform a restore and removes the risk of data loss when performing a refresh. This means that all data is contained within a single volume as opposed to the split volume model present in all other environments. In addition, as shown in Figure 7, the transient temp file system (B) is hosted locally, as is a test version of the `/sybhm` file system (C).

Figure 7) SATA data layout.



The process to refresh the test environment is similar to the one used to refresh the development environment. The exception is that the data and dump reside within the same volume, which in turn means that only a single clone (A) must be created to provide datasets for testing.

Note: The original production copy of the `/sybhm` file system is backed up by using SnapVault and is available if needed.

TOOLSETS USED

A number of tools are installed. However, only two NetApp data management tools are used:

- **SnapDrive[®] for UNIX[®].** This tool provides the ability to create Snapshot copies of the NFS volumes.
- **Snap Creator.** This tool provides the data Snapshot copy and movement framework that are the core of the backup process.

A number of supporting toolsets are also used such as host attach kits; DataFabric[®] Manager (DFM) agents; and the Symantec[™] volume manager, Array Support Libraries.

3.3 BACKUP CONFIGURATION

SNAP CREATOR FRAMEWORK CONFIGURATION

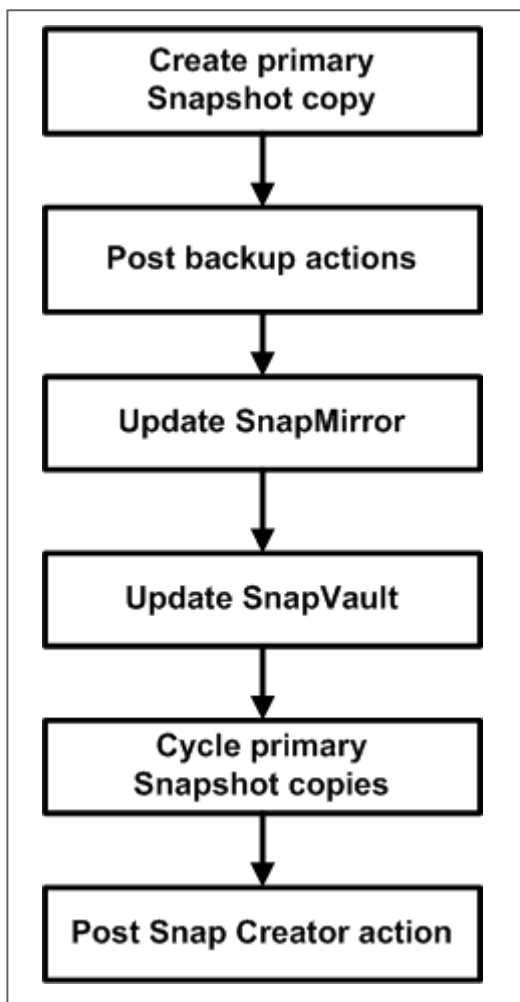
The Snap Creator framework can be used to integrate applications, processes, and Snapshot copies into a single coherent service offering. We use this framework to perform Snapshot management and task scheduling.

The key to Snap Creator is its configuration file. A number of items are defined within this file:

- Storage system access
- Volumes to manage
- Number of Snapshot copies to maintain
- Hooks at each stage of the process to allow integration.

The workflow, which is very linear, is described in Figure 8.

Figure 8) Snap Creator workflow.

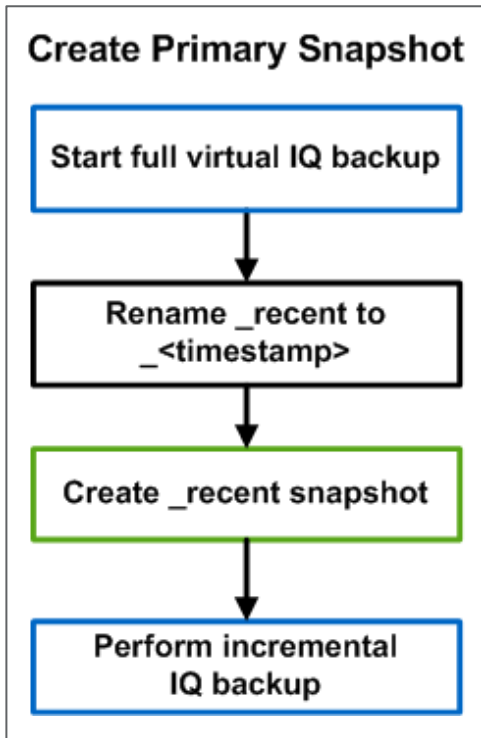


Each of the main functions contains a number of sub-events (or tasks), which are described in the following sections. Tasks in blue boxes are custom items; we have made some modification to the default behavior of the product. Tasks in green boxes are the default behavior of the product and are included for the sake of completeness.

PERFORM PRIMARY BACKUP

As shown in Figure 9, four main tasks are performed at this stage of the process.

Figure 9) Process to perform primary backup.



The first step in the process is to start the virtual full backup of the IQ database. After the virtual full backup has been performed, Snap Creator manages the creation of the Snapshot copies within the two defined volumes.

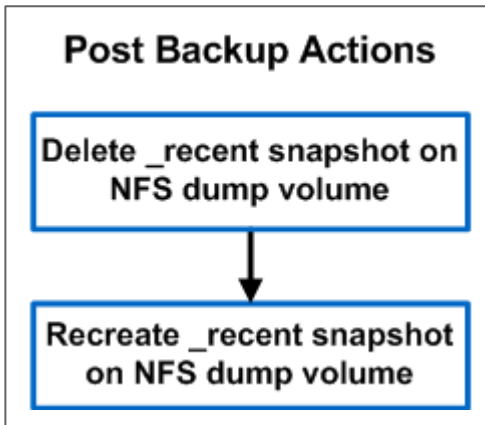
The management of Snapshot copies is a simple matter of renaming the old Snapshot copy (which is always suffixed with `recent`) and creating a `new_recent` Snapshot copy. The actual Snapshot copy name contains the configuration and policy that create a Snapshot copy name as follows `sybaseIQ_<configuration>_<policy>_recent`. When the Snapshot copy is renamed, the `_recent` is replaced with a timestamp that provides a unique Snapshot name within the volumes.

3.4 POST BACKUP ACTIONS

One limitation of Snap Creator is that all defined volumes have their Snapshot copies created at the same time. This presents a problem when we want the NFS dump volume Snapshot copy created after the second incremental because the full backup has completed. As a workaround, we delete the Snapshot copy of the NFS dump area and recreate it as part of the post backup action. As long as the same Snapshot name is retained, it preserves the ability for the rest of the process (notably the SnapVault update) to complete.

The volume in question is NFS based; therefore, we can use SnapDrive to delete and recreate the Snapshot copy. This also makes the callouts defined within the section callouts to the SnapDrive executable. Figure 10 defines the two-step process that we follow.

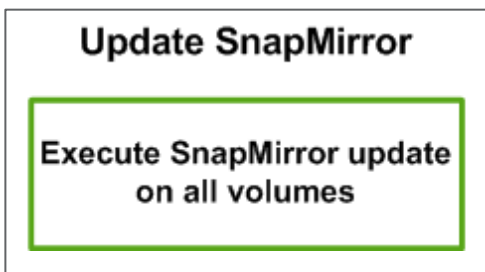
Figure 10) Post backup actions.



PERFORM SNAPMIRROR UPDATE

A single-step process polls the SnapMirror status and updates the destination file system to reflect the incremental changes that occurred to the source. Figure 11 defines the step.

Figure 11) Update SnapMirror.

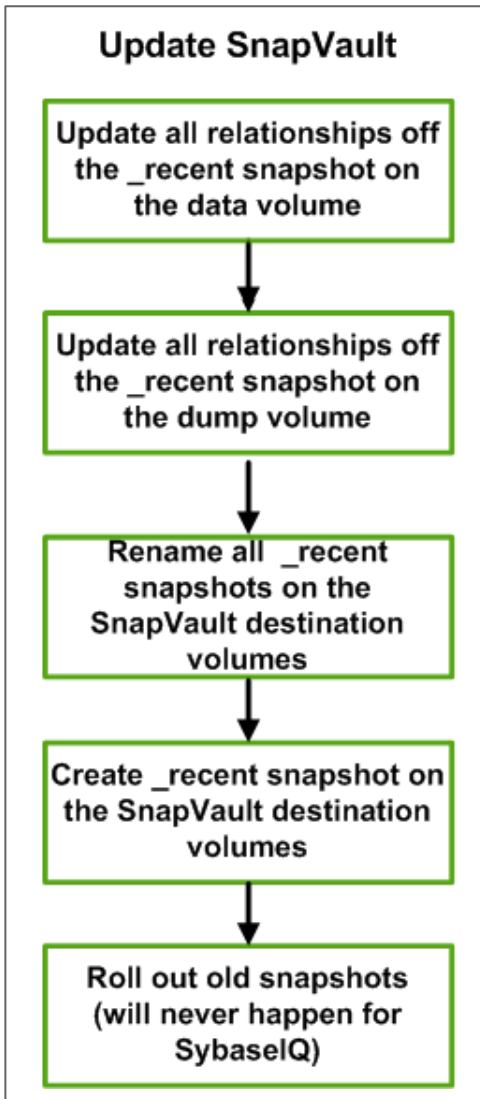


Snap Creator interrogates the source volumes (named within the configuration file) to determine the remote storage system and volumes that will be updated. Therefore, the SnapMirror relationships must be defined before performing the backup. For more information on NetApp SnapMirror, refer to <http://media.netapp.com/documents/tr-3446.pdf>.

PERFORM SNAPVAULT UPDATE

As shown in Figure 12, Snap Creator entirely manages a multistage process.

Figure 12) Process to update SnapVault.




Similar to the SnapMirror part of the process, Snap Creator determines the SnapVault relationships based on the source volume. This process has a similar requirement that the SnapVault relationships must exist before performing the backup. For more information on NetApp SnapVault, refer to <http://media.netapp.com/documents/tr-3487.pdf>.

CYCLE PRIMARY SNAPSHOT COPIES

At this stage, the backup is all but complete. The SnapMirror and SnapVault destinations have been updated, and the Snapshot copies of the SnapVault destinations have been performed. However, we have not enforced any Snapshot retention policies. This enforcement is done within the single-step phase shown in Figure 13.

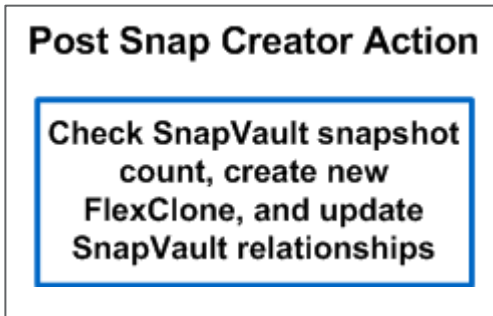
Figure 13) Cycle primary Snapshot copies.



POST SNAP CREATOR ACTIONS

The final task is the management of the SnapVault destination volumes to allow seven years of Snapshot copies to be retained. The simple process is to count the number of Snapshot copies. If the amount is greater than 240, use FlexClone to reestablish the SnapVault relationships with the FlexClone volume. Figure 14 shows the completion of the Snap Creator process.

Figure 14) Post Snap Creator action.



4 BUSINESS BENEFITS

Table 3 summarizes the amount of time saved during the backups and restores for this customer.

Table 3) Business benefits.

	Before	After	Improvement
Creating new DB instance			
Copy from short term	12 hours	10 minutes	99%
Copy from long term	24 hours	10 minutes	99%
Amount of time required for backup and restore			
Full backup to short term	12 hours	45 seconds	99%
Full backup to long term	12 hours	4 minutes 15 seconds	99%
Full restore from short term	12 hours	10 minutes	99%
Full restore from long term	24 hours	8 hours	67%

As shown in Table 3, by using NetApp Snapshot copies, the backup time was reduced from 12–24 hours to a few minutes, which is a 99% reduction in the amount of time required to back up and restore a Sybase IQ database. The recovery time objective (RTO) also improved.

The business benefits of this solution include a fast and efficient way of preparing development and testing environments. In addition, it provides the business with a disaster recovery (DR) capability within a two-hour recovery point objective (RPO).

5 REFERENCES

- SnapMirror Async Overview and Best Practices Guide
<http://media.netapp.com/documents/tr-3446.pdf>
- SnapVault Best Practices Guide
<http://media.netapp.com/documents/tr-3487.pdf>

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.



www.netapp.com

© 2011 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, DataFabric, Data ONTAP, FlexClone, NearStore, SnapDrive, SnapMirror, Snapshot, SnapRestore, and SnapVault are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Symantec and Veritas are trademarks of Symantec Corporation. UNIX is a registered trademark of The Open Group. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. TR-3905