



Technical Report

Integrating Sybase IQ Server Enterprise for UNIX with a NetApp Storage System

Suresh Vundru, NetApp
June 2010 | TR-3851

TABLE OF CONTENTS

1	INTRODUCTION	3
2	INSTALLATION PREREQUISITES	3
2.1	SYSTEM REQUIREMENTS	3
2.2	SYBASE IQ SERVER REQUIREMENTS	3
2.3	NETAPP STORAGE SYSTEM REQUIREMENTS	3
2.4	NETWORK REQUIREMENTS	4
2.5	SYBASE IQ/SERVER USER ACCOUNT	4
2.6	CREATING AN NFS VOLUME	4
2.7	CREATING THE MOUNT POINT	5
2.8	ADDITIONAL INFORMATION	6
3	SYBASE IQ SERVER INSTALLATION	6
3.1	INSTALLATION PROCEDURE	6
3.2	TESTING THE SYBASE IQ DATABASE	7
3.3	CREATING A DATABASE	8
3.4	VERIFYING THE DATABASE	10
4	DATABASE AND DISK RESIZING BY USING FLEXIBLE VOLUMES	10
4.1	CREATING A FLEXIBLE VOLUME WITH A SPACE GUARANTEE	11
4.2	CREATING A FLEXIBLE VOLUME WITH THIN PROVISIONING	12
4.3	DISADVANTAGES OF USING TRADITIONAL VOLUMES	14
5	DATA PROTECTION USING SNAPSHOT TECHNOLOGY	14
5.1	CREATING A SNAPSHOT COPY	14
5.2	PERFORMING SNAPRESTORE	16
6	DEVELOPMENT AND TEST ENVIRONMENT	18
6.1	CREATING A VOLUME CLONE	18
6.2	MOUNTING A VOLUME CLONE ON THE TEST SERVER	18
	ACKNOWLEDGEMENT	19

1 INTRODUCTION

This document describes the procedure to integrate Sybase IQ Server for UNIX[®] with a NetApp[®] storage system. It offers Sybase IQ and Storage administrators the ability to effectively deploy Sybase IQ based data warehouse environments with the space efficient NetApp Snapshot[™] feature to perform online backup and recovery. It describes how Sybase IQ and storage administrators can leverage the space-efficient NetApp FlexClone[®] feature to create a dev/test environment based on the production data.

This document assumes familiarity with the following:

- Sybase IQ Server
- The operation of NetApp storage systems
- Knowledge of the UNIX operating system. The supported UNIX platforms are Solaris[™], HP/UX, IBM AIX, and Linux[®]

2 INSTALLATION PREREQUISITES

2.1 SYSTEM REQUIREMENTS

Before beginning installation, perform the following steps:

- Check for the platform and operating system release levels, required RAM size, and supported network protocols; for example, Linux on x86 or x86-64 must have at least 1GB of RAM space dedicated for Sybase IQ.
- Check for the latest Sybase IQ software versions and certification at <http://media.netapp.com/documents/sybase-certification-matrix.pdf>
- Refer to the installation guide for instructions on installing Sybase IQ on your platform. Sybase supports the TLI/TCP and TLI/SPX network protocols.
- Make sure that the following hardware and software prerequisites are met before installing Sybase IQ on UNIX for use in conjunction with a storage system:
 - A server machine in which Sybase IQ will be configured (running an OS version of UNIX supported by Sybase)
 - An installed and configured storage system
 - A network with supported protocols
 - A UNIX user account named “sybase”

2.2 SYBASE IQ SERVER REQUIREMENTS

Sybase IQ Server software is required for the installation. Complete the entries on the preinstallation requirement sheet provided in the Sybase IQ installation guide. Follow procedure described in the “Sybase IQ Server Enterprise Installation Guide” for your platform.

2.3 NETAPP STORAGE SYSTEM REQUIREMENTS

- A NetApp storage system running Data ONTAP[®] version 7.x or later
- The NFS license on the storage system must be activated
- The NFS protocol must be set up and running

2.4 NETWORK REQUIREMENTS

A network connection between the Sybase Server machine and the storage system is required. In this setup, a single 100BaseT connection is used. Using multiple connections or a faster network technology improves performance. Sybase supports TLI/TCP and TLI/SPX protocols to enhance data integrity. Figure 1 illustrates the network configuration used to test this solution.




Figure 1) Network setup.

2.5 SYBASE IQ/SERVER USER ACCOUNT

On the server, create the user account "Sybase" and provide read/write and execute permission on the directory where NFS is mounted and that holds Sybase IQ installation binaries and default databases.

Sybase IQ by default creates a DBA user account that has all the permissions to perform DBA activities (backup and restore, database administration, performance, and tuning).

2.6 CREATING AN NFS VOLUME

1. Create a new volume on the NetApp storage to install Sybase binaries.

For example, to create the volume `sybhome` in the `aggr1` aggregate with 50GB size:

```
storage> vol create sybaseiq aggr1 50g
```

If you prefer to logically partition the volumes into multiple qtrees for various reasons, including better manageability. Add qtrees and set the security style for the volume as 'unix'.

2. Export the `sybaseiq` volume with `rw` and `root` permission to nodes only using `exportfs`:

```
storage> exportfs -p rw,root=<SybaseIQ Server IP>,anon=0 /vol/sybaseiq
```

3. Export the new volume:

```
storage> exportfs -a
```

2.7 CREATING THE MOUNT POINT

1. Create a directory on the Sybase IQ Server machine with rwx permission to “sybase” user.

```
# mkdir -p /sybaseiq/home
# chown sybase /sybaseiq/home
```

Where `sybase` is the user account created as described in [section 2.5](#).

2. Mount the NFS volume on the Sybase IQ Server machine as described below.
3. Set a few nondefault NFS mount options while mounting the NFS volume on the Sybase IQ Server machine. Based on requirements, the system administrator should define these NFS mounts with the necessary options. For example:

```
-o hard, vers=3, proto=tcp
```

Table 1) NFS mount options.

NFS Mount Options	Description
<code>hard</code>	This option determines that the mount point should never time out and that the Sybase IQ Server machine should not run without it. This causes the Sybase IQ Server machine to hang if the storage system is not responding to NFS for some reason. If the Sybase IQ Server machine is starting and if the storage system is not found, it does not complete the boot and Sybase does not start. If it is already up and running, all I/O to and from the storage system is suspended until the storage system is available again.
<code>vers</code>	This option is supported in recent releases of UNIX. It specifies that the protocol to be used is NFS version 3. This is to take care of compatibility issues. If the system administrator wants to try version 2, no harm occurs and the performance of Sybase is not affected. In our setup we used the default settings (version 3).
<code>proto</code>	Set the value as TCP for this option, because UDP is not a reliable transport for database operations. Note that Sybase supports only TLI/TCP and TLI/SPX protocols, not UDP protocols.

4. Create an entry in the `/etc/vfstab` to make the mount point persistent upon rebooting the server:

```
Storage:/vol/sybaseiq - DatabasePath nfs - yes hard,vers=3,proto=tcp
```

Where `DatabasePath` is the path to the directory to mount the storage system (for example, `/sybaseiq/home`).

Other UNIX operating systems have similar mechanisms for mounting file systems when the system is first started. Mount the file system on the server by using the appropriate mount command:

```
# mount /sybaseiq/home
```

Now the NFS mount point is available to install Sybase IQ software.

Follow the steps in section 2.6 and 2.7 to create additional volumes to use as dbspace, in this setup I have created two FlexVols and mounted on `/iq` and `/db`.

2.8 ADDITIONAL INFORMATION

Before you install the Sybase product, read the Sybase IQ release notes and check for any known issues and their workarounds.

3 SYBASE IQ SERVER INSTALLATION

This section describes the process of installing Sybase IQ over NFS on NetApp storage. This section does not describe the detailed steps involved in installing Sybase IQ. For the detailed, step-by-step procedure for installing Sybase IQ Server, refer to <http://infocenter.sybase.com/help/index.jsp>.

Note: The prerequisites described in [section 2.1](#) are required for this setup to work.

3.1 INSTALLATION PROCEDURE

1. Insert the Sybase IQ Product CD in the appropriate drive or download and extract the Sybase IQ install image from the Sybase Product Download Center (SPDC).
2. Log in to the Sybase IQ Server machine as 'sybase' user.
3. Run the "sybinstall" utility from the directory where the CD is mounted.
4. Select a directory for the installation (in our configuration, change the directory to `/sybaseiq/home`).
5. Follow the steps described in the Sybase IQ Installation Guide.

3.2 TESTING THE SYBASE IQ DATABASE

Once Sybase IQ is installed with all default settings, the server and database are built. Execute the shell script `SYBASE.sh` script created during the installation procedure. This script is used to set the environment variables needed to run the Sybase Server. After executing it, copy the contents of the `SYBASE.sh` to `.profile` under the home directory of 'sybase' user, which was created to install Sybase IQ and explained in [section 2.5](#).

To test the database, you can either test using CLI or using GUI Sybase Central from the Sybase IQ Server. We will invoke Sybase Central to test the default demo database `asiqdemo.db`, which is created as part of the installation.

```
#./scjview
```

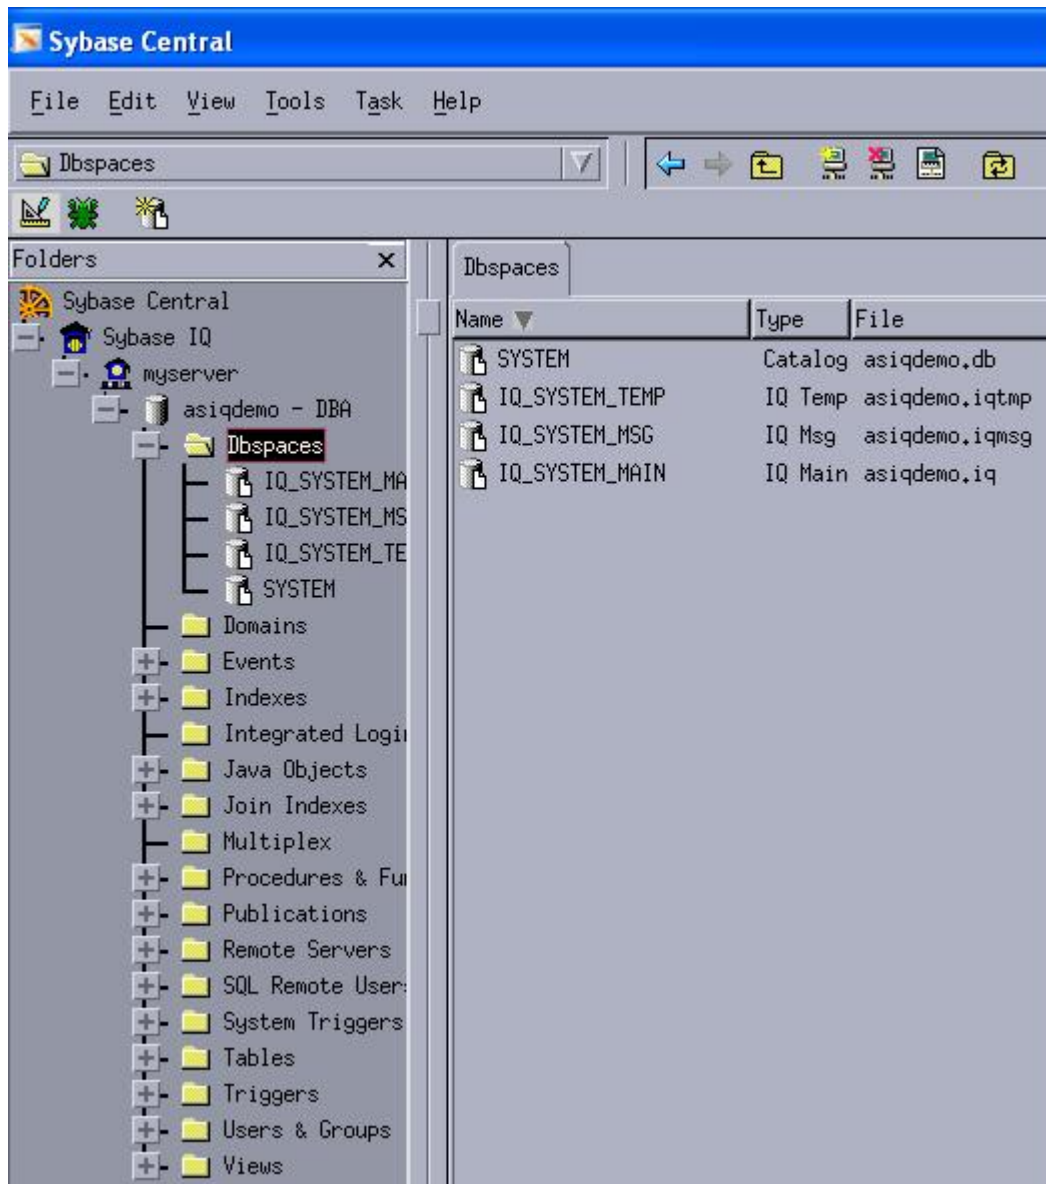


Figure 2) Sybase Central.

If the output of this command (a set of database dictionary objects) is displayed on the screen, the database is working correctly. You can now populate the database as you normally would.

This procedure has installed the Sybase IQ Server, System Main IQ, Catalog db, System TMP, and System MSG files, which can reside on the storage system. Sybase IQ recommends placing the IQ store, Catalog db store, and IQ tmp stores on different FlexVol® volumes on different aggregates on NetApp storage. This helps to improve performance and recovery if the database gets corrupted.

3.3 CREATING A DATABASE

The following command creates a database device and databases on the storage system using the command line `./dbisql`. You can run this command using the open interactive SQL GUI or using CLI. Make sure myserver is started and then connect to the phantom database Utility_db for creating user databases.

```
#./start_asiq -n myserver
```

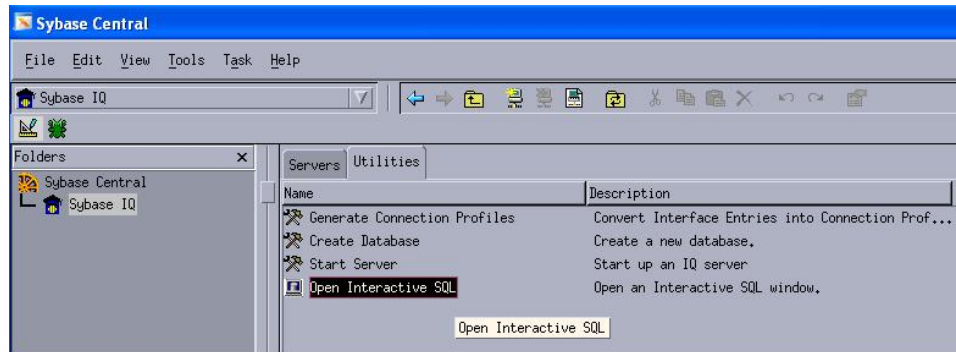


Figure 3) Creating a database (1).

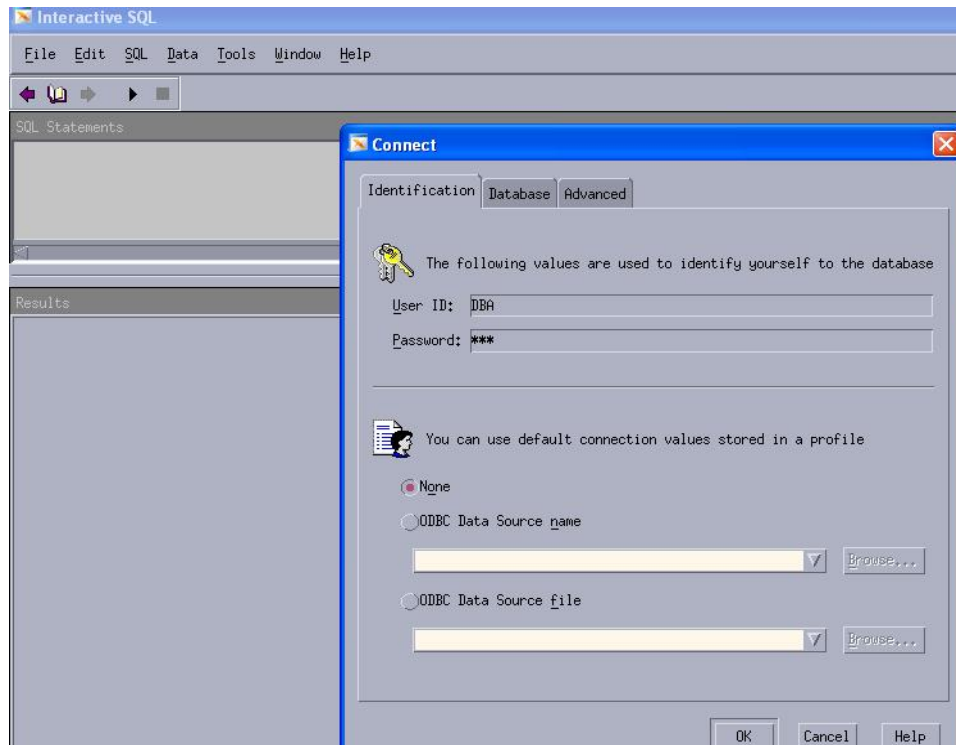


Figure 4) Creating a database (2).

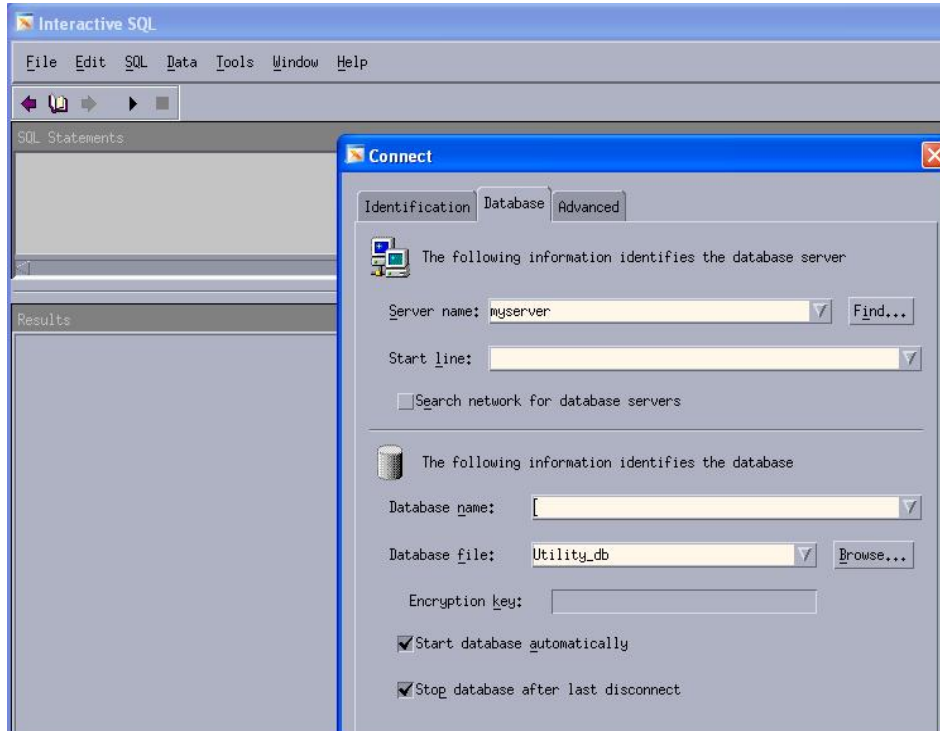


Figure 5) Creating a database (3).

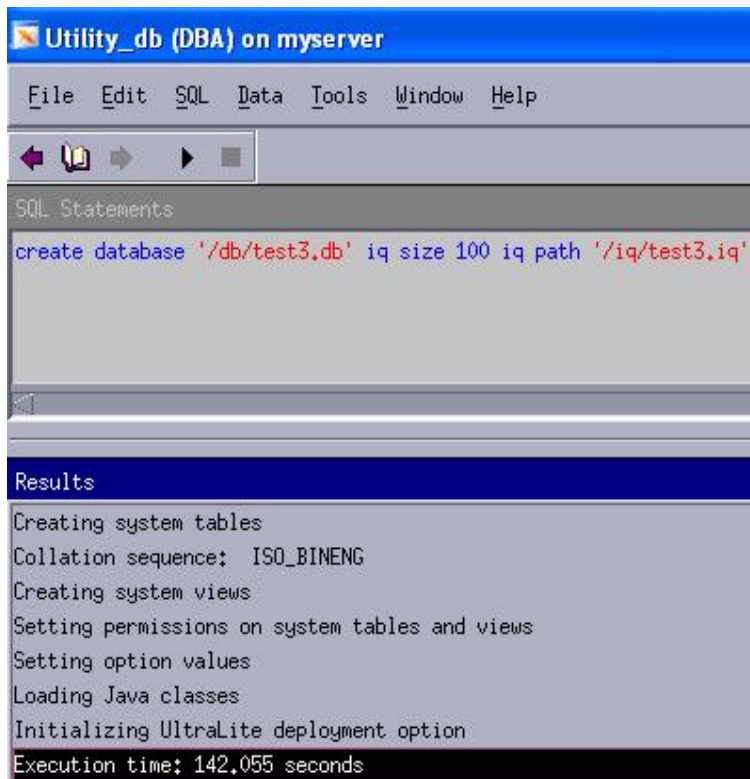


Figure 6) Creating a database (4).

3.4 VERIFYING THE DATABASE

Once the Sybase IQ database is created you can verify it by connecting and creating a table.

Figure 7 shows a sample output.

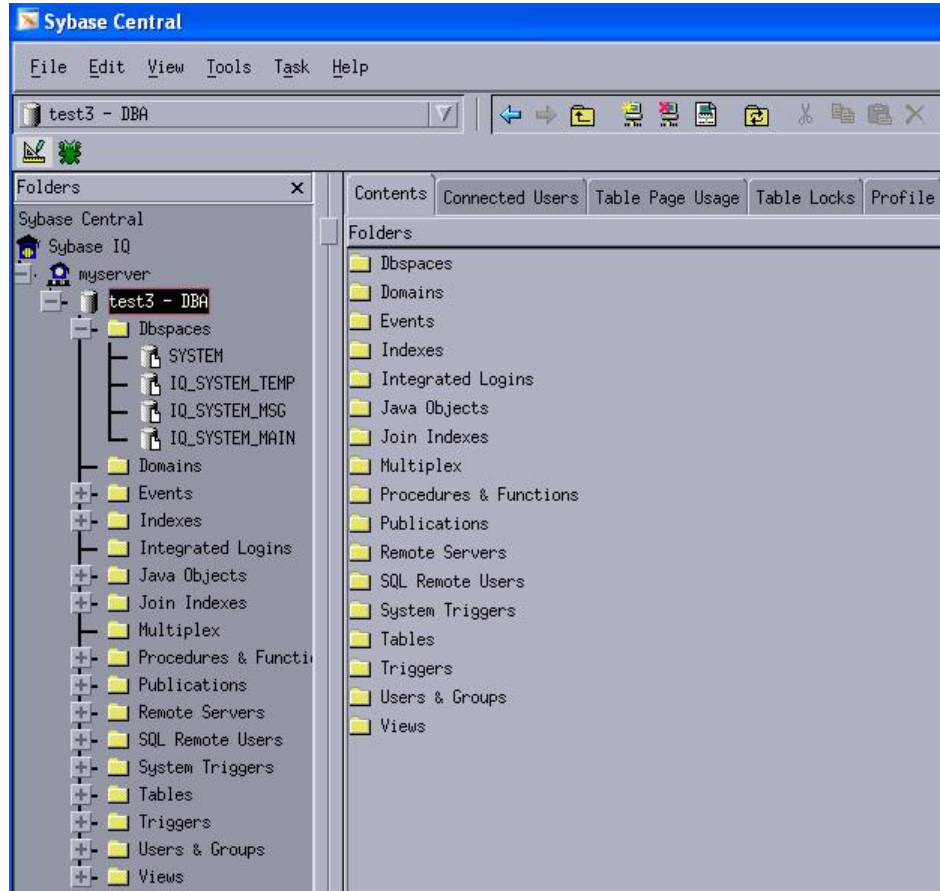


Figure 7) Verifying the database (1).

4 DATABASE AND DISK RESIZING BY USING FLEXIBLE VOLUMES

This section describes how to use the FlexVol feature available in Data ONTAP 7.0 and later to resize flexible volumes on NetApp storage without any application downtime. Flexible volumes are logical data containers that can be sized, resized, and managed. A flexible volume can share its containing aggregate with other flexible volumes to allow better storage space provisioning. It can also run faster because it uses all spindles in the aggregate available to them at all times. The recommended best practice for better performance and proper disk space utilization is to create one large aggregate with more disks and then allocate space for multiple flexible volumes on top of it.

Table 2 lists the logical view of storage and hosts on the NFS and Sybase device entries described in [section 4.1](#) and [section 4.3](#).

Table 2) NFS and Sybase device entries for aggregate "Aggr1."

Flexible Volume Name	Volume Size	NFS Mount Points on Storage	Hosts /etc/vfstab entries	Sybase Devices
Test1	50GB	exportfs -p rw,/vol/ERP1	10.x.x.x:/vol/ERP1 - /ERP1 nfs - yes rw,bg,hard,nointr,rsize=32768,wsiz=32768,proto=tcp,vers=3,suid	ERP1 device

Test2	50GB	exportfs -p rw,/vol/ERP2	10.x.x.x:/vol/ERP2 - /ERP2 nfs – yes rw,bg,hard,nointr,rsize=32768,wsz=32768,proto=tcp,v ers=3,suid	ERP2device
-------	------	-----------------------------	---	------------

4.1 CREATING A FLEXIBLE VOLUME WITH A SPACE GUARANTEE

Space guarantees on a FlexVol volume enable writes to a specified FlexVol volume or writes to files with space reservations enabled to not fail because of lack of available space in the containing aggregate. The space guarantee option reserves space well in advance although the database is not consuming the entire space on the volume. This is really important in an environment in which a critical database needs to be online around the clock and storage space is an important component of keeping on pace with business data. There are situations in which space needs to be reclaimed from reservation because of the space guarantee and needs to allocate for critical database usage. This can be achieved in conjunction with the FlexVol feature, with which volume space can be dynamically decreased or increased while the database and storage are online. The following example describes a scenario in which the FlexVol feature is extensively utilized.

Example:

Let us assume that FlexVol Test1 and FlexVol Test2 with a 50GB size are created with a guarantee set to the volume, which reserves 100GB of storage space in advance from aggregate Aggr1. Sybase IQ dbspace Sales1 is created on FlexVol Test1. FlexVol Test2 is provisioned to a file server. Sybase IQ dbspace Test1 needs more space to grow and hence the underlying storage Test1 volume needs to be grown. However, aggregate Aggr1 does not have adequate free space. The only solution is to decrease the FlexVol Test2 volume size, as it is not consuming all the space and its growth rate is likely not to be high enough to consume the entire space in the near future. Once FlexVol Test2 releases space to aggregate Aggr1, FlexVol Test1 volume can be grown while the Sybase IQ database remains online.

- Storage command to create aggregate Aggr1:

```
storage>aggr create aggr1 -t raid_dp -d 0d.20 0d.21 0d.22 0d.23 0d.24 0d.25  
0d.26 0d.27 0d.28 0d.29 0d.30
```

As Sybase IQ is a highly parallel database, the performance can be increased by providing more spindles.
- Storage command to create FlexVol Test1 and FlexVol Test2:

```
storage>vol create ERP1 -s volume aggr1 50g  
storage>vol create ERP2 -s volume aggr1 50g
```
- NFS configuration command:

```
storage> exportfs -p rw /vol/ERP1 or wrfile /etc/exportfs  
storage>exportfs -p rw /vol/ERP2 or wrfile /etc/exportfs  
storage>exportfs -a
```
- Sybase IQ commands to create dbspace:




Figure 8) Creating dbspace.

- Storage commands to resize volumes:

```
storage>vol size Test2 -10G
```

```
storage>vol size Test1 +10G
```
- Sybase commands to resize a device:

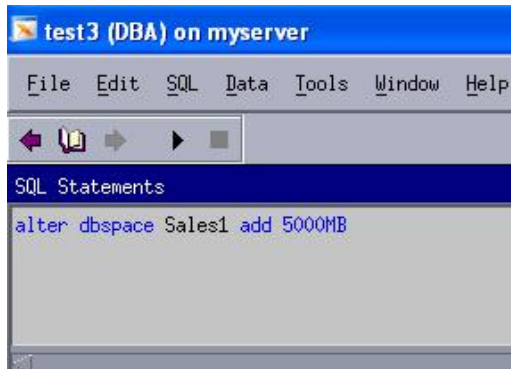


Figure 9) Resizing a device.

4.2 CREATING A FLEXIBLE VOLUME WITH THIN PROVISIONING

You can set the volume space higher than the available physical space on an aggregate by changing the space guarantee option to none on the volume. This allows you to overprovision more storage than available, a method known as thin provisioning. However, be sure to monitor the available space on the aggregate, which must be increased accordingly, or you will run out of space and end up losing data. The three types of space guarantee options are:

- **Volume.** This setting means that the amount of space required by the flexible volume is always available from its aggregate. This is the default setting for flexible volumes.
- **File.** This setting means that space is always available for the write operation on the NFS volume.
- **None.** This setting does not allow reserving space and overrides any existing space reservation settings on the volume. The write operations to the volume might fail if the containing aggregate does not have enough available space.



Example:

Let us assume the thin-provisioned flex volume Test1 is created with a 100GB size and a space guarantee of none, so that it does not reserve any space on the aggregate aggr1 with 60GB size. By using this thin-provisioning feature, you can overprovision the storage size against your available usable aggregate size. Sybase Sales1 IQ dbspace with 45GB is created on a Test1 flex volume. It is observed that the Sales1 IQ database needs 10GB space and hence aggr1 needs provisioned with more disks. Then the Test1 IQ main database needs to be grown.



- Storage command to create aggregate aggr1:

```
storage>aggr create aggr1 -t raid_dp -d 0d.20 0d.21 0d.22 0d.23 0d.24
```
- Storage command to create FlexVol Sales1 with thin provisioning enabled:

```
storage>vol create Test1 -s none aggr1 100g
```
- Sybase commands to create IQ dbspace Sales1:



- Storage command to grow aggregate:
storage>aggr add aggr1 -g raid_dp 0d.25 0d.26
- Sybase commands to resize IQ dbspace Sales1:



For best practices when thin provisioning volume space, see <http://now.netapp.com/NOW/knowledge/docs/ontap/rel713/html/ontap/bsag/c4crf3.htm>.

4.3 DISADVANTAGES OF USING TRADITIONAL VOLUMES

Suppose that Sybase databases are created on traditional volumes named `trad1` and `trad2`, each of which is 10GB in size and uses RAID 4. After a period of time, the `trad1` volume consumes only 1GB and the `trad2` volume consumes 9GB. Traditional volumes can be scaled by adding more disks, but they cannot be decreased in size. In this scenario, because the `trad2` volume has to be increased, the only solution is to add an extra disk. The disadvantage of a traditional volume is that the unused space under the `trad1` volume cannot be reused because it cannot be compressed to free up storage space that can then be added to scale `trad1`. Also, when it comes to performance, you are restricted with limited disks when IOPS numbers have to be calculated in sizing for database performance.

5 DATA PROTECTION USING SNAPSHOT TECHNOLOGY

A Snapshot copy is a frozen read-only image of a volume that captures the state of the file system at a point in time. Data ONTAP provides a configurable Snapshot scheduler that defines the number of Snapshot copies to be created and retained, and specifies the time and days to take the Snapshot copies.

The Snapshot scheduler is unaware of applications writing data to volumes when it performs the Snapshot operation, and therefore NetApp recommends driving the Snapshot technology from the host where the Sybase IQ database is installed. This is important because when Data ONTAP creates a Snapshot copy on the volume where Sybase IQ is actively writing or updating the database from memory, you may end up having inconsistent data in the Snapshot copy. When `snap restore` is run on the volume, you may end up with a corrupt database.

5.1 CREATING A SNAPSHOT COPY

This section describes the best practices for taking Snapshot copies in a Sybase IQ environment.

Note: You can run storage commands by using RSH from the host. To enable RSH on storage, see <https://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb7872>.

1. Turn off the Snapshot scheduler on the volume where the Sybase IQ database is installed. Running the Snapshot scheduler when the RSH is active results in the creation of incorrect Snapshot copies and consumption of space.

To estimate the maximum number of Snapshot copies that can be retained on a volume, see <http://now.netapp.com/NOW/knowledge/docs/ontap/rel1003/html/ontap/admin/guide-volume-tasks-volume-snapshot.html>. The maximum number of Snapshot copies to be retained also depends on the Data ONTAP version.

To turn off the Snapshot scheduler, using the CLI or FilerView[®] enter:

```
storage>snap sched Test1 0 0 0 0
```

2. Use the Sybase Central tool and Open Interactive SQL mode to perform a test Sybase IQ database backup on the server.

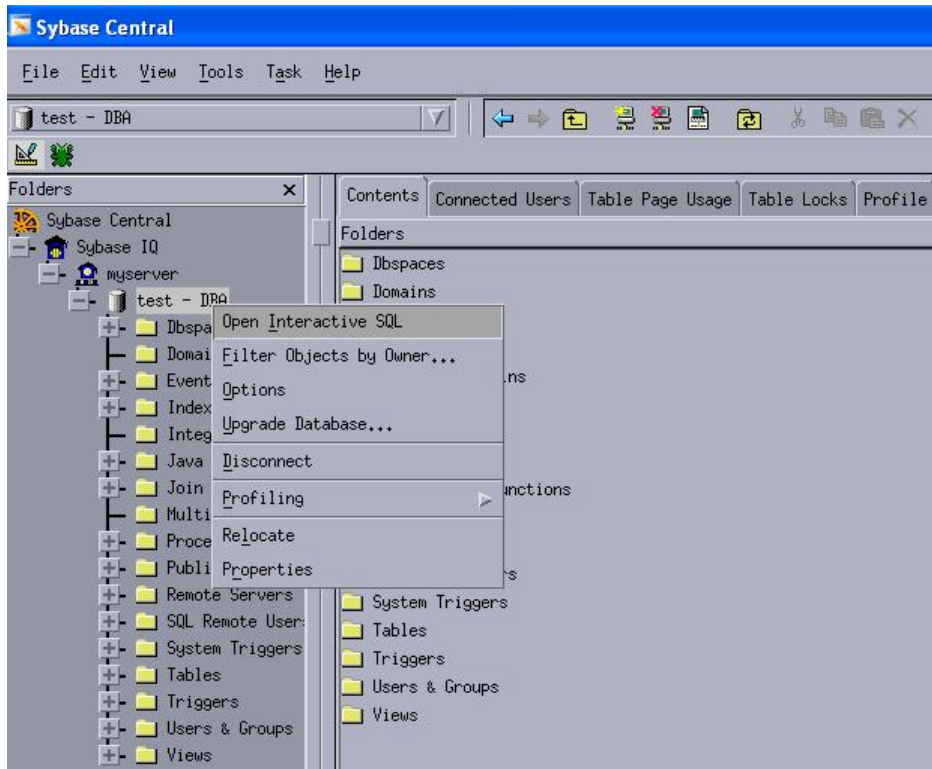


Figure 10) Open Interactive SQL.

3. Perform a Sybase IQ backup as virtual decoupled and take a full backup to storage volume and then take a Snapshot copy of the volume.




Figure 11) Backing up the database.

4. After you perform the virtual decoupled backup using RSH or FilerView, take a Snapshot copy of the Test1 volume where the database is created.

```
rsh 10.x.x.x snap create -V Test1 testsnap.1
```
5. After you take a consistent Snapshot copy on the storage, perform an incremental virtual decoupled backup on the Test1 IQ database.




Figure 12) Backing up the database.

5.2 PERFORMING SNAPRESTORE

If there is database corruption, you can restore the database back to a point in time by choosing the appropriate Snapshot copy. The following example uses a Testsnap .1 Snapshot copy to perform a restore. Before performing `snap restore`, stop the Sybase IQ database and unmount the NFS.

1. Make sure you stop the Sybase IQ server and then unmount the NFS mount point on the server.
`./stop_asiq`
2. Perform SnapRestore using the consistent Snapshot copy taken and explained in backup section 5.1.
`rsh snap restore Test1 -s Testsnap.1`
3. After the Test1 volume is restored with the point-in-time Testsnap.1 Snapshot copy, mount the Test1 volume on the server.
4. Start the Sybase IQ Utility_db phantom database on the server using Sybase Central Interactive SQL Utilities to perform the Sybase IQ database restore.
`./start_asiq -n myserver`

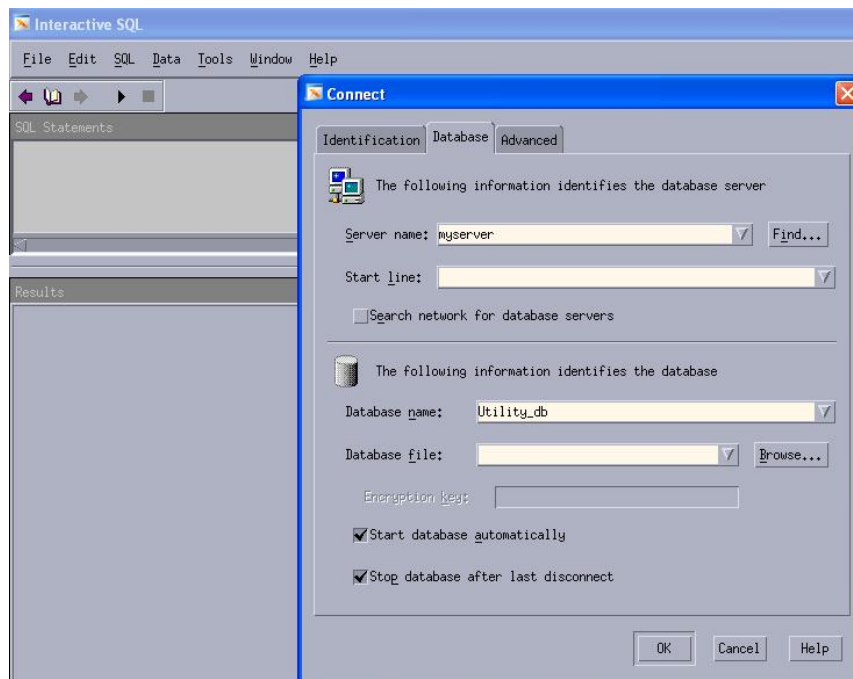


Figure 13) Performing database restore.

5. Perform the full virtual decoupled full backup restore using the following commands. Before performing this, make sure your server has access to the shared storage where the Sybase IQ database backup files are located.



Figure 14) Performing database restore.

6. Perform a virtual decoupled incremental restore using the following commands:




Figure 15) Performing database restore.

7. Perform data validation and integrity checks once the IQ database is restored.

6 DEVELOPMENT AND TEST ENVIRONMENT

NetApp FlexClone technology instantly replicates data volumes and data sets as transparent virtual copies which increase productivity and saving storage space without compromising performance.

FlexClone saves space with minimal overhead, so you can create all the clones you need to improve productivity.

Each cloned volume is a transparent virtual copy that you can use for your essential business operations:

- Testing and bug fixing
- Platform and upgrade tests
- Multiple simulations against large data sets

6.1 CREATING A VOLUME CLONE

This section describes how to leverage FlexClone to rapidly deploy the test and development environment, debug any issue, and roll out an application or operating system patch in a test environment before it is applied to production.

1. To create a FlexClone volume for a test and development environment, use the `Testsnap.1` Snapshot copy created on the parent FlexVol `Test1` volume for performing a virtual decoupled backup, explained in [section 5.1](#). Then mount a FlexClone volume to a test host:

```
Storage>vol clone create Test1clone -b Test1 Testsnap.1
```

6.2 MOUNTING A VOLUME CLONE ON THE TEST SERVER

Create an NFS mount on the `Test1clone` volume on the storage as follows:

1. On the storage system, enter:

```
Storage>exportfs -p rw /vol/Test1clone  
Storage>exportfs -a
```
2. On the host, enter:

```
rsh 10.x.x.x exportfs -p rw /vol/Test1clone  
rsh 10.x.x.x exportfs -a
```
3. To mount the Sybase IQ database `Test1` on the test host, mount the `Test1clone` volume and also provide access to where virtual decoupled full and incremental backup dumps reside.

```
TestHost#mount 10.x.x.x:/vol/Test1clone - /Test
```

Make an entry in `/etc/vfstab`:

```
10.x.x.x:/vol/Test1clone - /Test nfs - yes rw,bg,hard,nointr,  
rsize=32768,wsize=32768,proto=tcp,vers=3,suid
```
4. After mounting the `Test1clone` volume on the test server, perform a virtual decoupled full and incremental restore on the Test IQ database following the steps explained in [section 5.2](#).
5. After you restore the virtual decoupled backup, start the IQ database using Sybase Central or the command line as explained in [section 3.5](#).

ACKNOWLEDGEMENT

The author would like to thank the following individuals for their contributions to this technical report:

NetApp

Michelle Nguyen, Paul Miach, and Sridhara Gangoor

Sybase

Joydeep Das and Andrew Neugebauer

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

© Copyright 2010 NetApp. 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, Data ONTAP, FilerView, FlexClone, FlexVol, and Snapshot are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. UNIX is a registered trademark of The Open Group. Linux is a registered trademark of Linus Torvalds. SQL Server is a registered trademark of Microsoft Corporation. All other brands or products are trademarks of registered trademarks of their respective holders and should be treated as such. TR-3851