



Informix IDS 9.2: Integration with a NetApp® Filer Using NFS

Bruce Clarke, Ed Sassi, and Maury Tiller | Network Appliance | TR 3075

TECHNICAL REPORT

Network Appliance, a pioneer and industry leader in data storage technology, helps organizations understand and meet complex technical challenges with advanced storage solutions and global data management strategies.

Table of Contents

1. [Introduction](#)
2. [Audience](#)
 - 2.1. [Terminology](#)
3. [Infrastructure](#)
 - 3.1. [IDS Server Machine](#)
 - 3.2. [Filer](#)
 - 3.3. [Network](#)
 - 3.4. [IDS User Account](#)
 - 3.5. [Filer NFS Mount Point](#)
 - 3.6. [Using Symbolic Links](#)
4. [Installing INFORMIXDIR onto a Filer](#)
 - 4.1. [Advantages](#)
 - 4.2. [Disadvantages](#)
5. [Quick Start Guide](#)
 - 5.1. [Establish a Persistent Mount Point](#)
 - 5.2. [Migrating the Informix Software](#)
 - 5.3. [Migrating the Informix Instance](#)
 - 5.4. [Modify the ONCONFIG File](#)
 - 5.5. [Start the Database](#)
6. [Installing INFORMIXDIR onto a Filer from Scratch](#)
 - 6.1. [Supporting Multiple Informix Versions on a Filer](#)
7. [Enabling Database Support on the Filer](#)
8. [Caveats](#)
9. [Revision History](#)

[TR3075]

1. Introduction

This document describes the steps necessary to integrate Informix® Dynamic Server™ (IDS) Version 9.2 with a Network Appliance™ filer. This paper also applies to the IDS 7.3X engine. Specifically, we cover the following issues:

- Creating an INFORMIXDIR directory on a filer.
- Placing the Informix software on a filer using NFS.
- Migrating Informix data from local disk onto a filer.
- Sharing one set of Informix executables with multiple database engines

2. Audience

This paper is intended for DBAs, system administrators, system engineers, or those with a good understanding of Informix, including the terminology used throughout this paper. All examples in this report are for Informix Dynamic Server 9.2 running on Sun™ Solaris™ version 7. The examples listed may, however, require modifications to run under your version of UNIX®.

It is advisable to have the *IDS Installation and Administrator's Guide* available for your version of Informix and generally follow its steps for installation and configuration. Where the Informix documents and this paper conflict, you should assume the Informix material is correct. Please inform Network Appliance of any such conflicts.

Of course, you can safely ignore the line in the *IDS Installation Guide* that reads: "Your Informix directory can be any directory that is local and is not mounted from another system on the network."

When using NetApp filers, it is perfectly safe to use a directory mounted onto your server to store your Informix files.

2.1. Terminology

The purpose of this paper is to describe the concepts and give examples of the steps required to get Informix database products working with a Network Appliance storage system. Informix refers to the database storage space as "chunks." The ultimate goal of this paper is to get the Informix "chunks" located on the Network Appliance storage system. Alternatively, the other Informix files may be stored there as well, including the execution binaries, and configuration and documentation files.

This paper makes extensive use of examples. Rather than repetitively saying that a command should be run as the UNIX root user and another command should be run as a regular user, this paper uses the generally accepted notation that commands requiring root are shown with the command-line prompt "#". Commands that can be run as a user other than a root user are shown with the command-line prompt "%".

3. Infrastructure

The following items are needed to run UNIX versions of IDS 9.2 in conjunction with a filer:

- IDS 9.2 server machine (running a version of UNIX supported by Informix)
- NetApp filer with Data ONTAP™ software version 5.3 or later
- Network connections
- Informix user account in the Informix group
- Root account on the IDS server machine
- Filer NFS mount point(s)
- One or more symbolic links on the IDS server machine referencing one or more directories or files on the filer (optional, but highly desirable)

3.1. IDS Server Machine

A UNIX server is required for running Informix. This paper was developed using Informix Dynamic Server 9.2 running under Sun Solaris 7. In your installation be sure that your system satisfies the requirements for running IDS. For more information on this issue, check the *IDS Installation Manual* for your target platform.

Other versions of UNIX should be compatible, including AIX, HP-UX, Linux®, SCO, IRIX®, etc. Each of these environments has its own idiosyncrasies and will require some slight modification to the procedures outlined in this document. If any problems arise, don't hesitate to contact your Network Appliance technical representative for assistance.

3.2. Filer

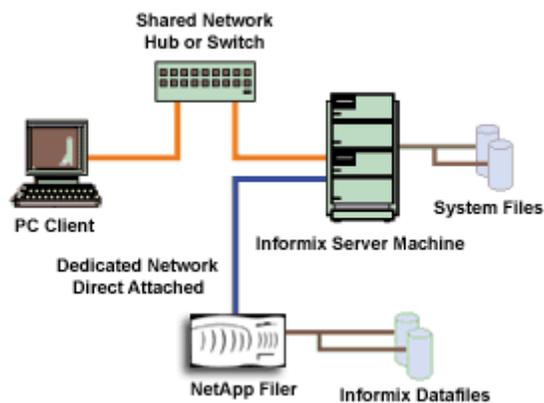
Any NetApp filer running Data ONTAP version 5.3 or better will work. The NFS license on the filer must be activated and the NFS protocol must be setup and running. The filer should also have

the "nvfail" option enabled on it. For further information, refer to Section 7, Enabling Database Support on the Filer.

3.3. Network

You need a network connection between the IDS server machine and the filer. We have used Fast Ethernet and Gigabit Ethernet, both of which work fine. A faster network connection (or possibly multiple network connections) will improve performance and offer resiliency in case of network problems.

The following shows the network configuration we used to test this solution:



As shown in the network diagram, we recommend that you dedicate at least one private network connection between the IDS server and the filer. This can be done using a crossover cable with most networking technologies (Fast Ethernet and Gigabit Ethernet being two examples we have tested successfully). Of course, network switches can also be effectively used and are required when trunking links and clustering filers together.

Dedicated network connections between the IDS server machine and the filer are recommended for the following reasons:

- Any issues of contention or latency are eliminated if the IDS server machine and the filer are the only two nodes in the subnetwork.
- Creating a private network connection ensures security. There is no issue of protecting the IDS chunk files from "snooping," which may happen on a shared network.

The above procedures generally describe creating a storage area network. This will provide all the normal SAN characteristics of reliability, scalability, and ease of management. Each dedicated network connection between the filer and the IDS server machine can be thought of as functionally equivalent to a SCSI or FC-AL connection that can alternatively be used to build a storage area network. High-end networking protocols such as Gigabit Ethernet provide similar bandwidth as well. This model maps well to traditional database tuning techniques for ensuring that no single I/O path bottlenecks performance. This topic is discussed later.

Of course, additional connections to the filer can be configured to provide nondatabase I/O. All filers can support multiple network connections, so the picture above would normally include additional connections to the shared network to provide access to the filer.

3.4. IDS User Account

If you have not done so already, start the installation process by following the Informix installation instructions for creating the Informix user and group. For example, create a group **informix** and then a user **informix** that is included in the group **informix**. You need to add a user account under which the Informix processes will run. This account should be added on both the IDS server machine and the filer. In the Sun Solaris 7 operating system, you can use the "useradd" command to accomplish this task. However, this will affect the IDS server machine only. You need to add an identical user on the NetApp filer. To add the user to the filer, you can copy the appropriate line of the "/etc/passwd" file on the IDS server machine into the "/vol/vol0/etc/passwd" file on the filer. (Alternatively NIS can be used on the filer. For more information on this issue, see *Data ONTAP System Administrator's Guide*)

The **informix** user account should have its login scripts initialize the environment variables required for proper Informix operation. These variables include (but may not be limited to):

INFORMIXDIR - the directory where Informix is installed
PATH - includes \${INFORMIXDIR}/bin in the path
ONCONFIG - the name of the current "onconfig" file
INFORMIXSERVER - the name of the IDS instance you want to control

Examples later in this paper assume the correct setting of these variables prior to using the Informix product specific commands.

3.5. Filer NFS Mount Point

You need to create a directory on the filer with security set to:

```
rw-rwx--- informix informix
```

In other words, create a directory owned by the user **informix** and the group **informix** and only accessible to that user and members of the **informix** group. This directory can be used for the **INFORMIXDIR** directory described in the *IDS Installation Guide*. Please note that this is not absolutely necessary but rather a suggestion. It is certainly possible to put the **INFORMIXDIR** directory on the server's local disk and then use symbolic links to point parts of the installation to the filer, but this is often more complicated and therefore less reliable. However, this approach works well for migrating a NetApp filer into an existing Informix environment. This will be discussed later in this document. Some nondefault NFS mount options should be set when configuring the method that the IDS server machine uses to connect to the filer. The system administrator should define these NFS mounts with the following options:

```
-o hard,intr,vers=3,proto=udp,suid
```

The following table explains these options in detail.

hard	This option says that the mount point should never time out and that the IDS server machine should not run without it. This will cause the IDS server machine to hang if the filer is not responding for any reason. If the IDS server machine is booting and the filer is not found, then it will not complete the boot and IDS will not start. If IDS is already up and running and the filer fails, all I/O to and from the filer will suspend until the filer is available again.
intr	This option indicates that the operator wishes to be able to escape out of operations that might hang on this mount point. This is sometimes useful but also can cause problems. Network Appliance recommends starting with this option. If mysterious intermittent disconnections occur, try running without this option.

vers	This option is supported in recent releases of UNIX. This option specifies the NFS version that should be used. Some versions of UNIX have been reported to have serious performance problems when running with NFS Version 3. Others perform better using NFS Version 3 instead of Version 2. The system administrator should try "vers" with both NFS versions, and should run with the NFS version that provides the best performance.
proto	Along with the "vers" parameter, "proto" gives the system administrator the option of choosing whether UDP or TCP protocol should be used. For NFS over local area networks, UDP offers slightly less overhead (and therefore better performance) than TCP. However, if your network connection path between the NetApp filer and the IDS server is prone to lose packets, drop frames, or introduce checksum errors, then TCP can improve performance compared to UDP. (On the other hand, you should not run IDS with its back-end storage being accessed over such an unreliable network.) We therefore recommend that you run using UDP on a dedicated network connection with a crossover cable or other dedicated or semidicated networking hardware between the IDS server and the filer. For more information on this issue, see Section 3.3, Network, above. If you use UDP, be sure UDP checksums are enabled on the IDS server machine. Most UNIX systems have this option set by default.
suid	If the Informix binaries are put on the filer, some of them have the Set-UID bit set configured. If this is the case, you must enable this operation using this parameter.

3.6. Using Symbolic Links

As mentioned earlier, it is desirable to use symbolic links to map the IDS server directory structure to the filer. This is true for several reasons.

1. If you are migrating from local disk to a filer, the use of symbolic links prevents you from having to make any modifications to the database setup.
2. If you later wish to reorganize the location of your IDS chunk files, you can do so easily by simply taking the IDS server down, moving the files, and then editing the link. The symbolic link thus provides an extra layer of abstraction to your file layout. Without the symbolic link, you need to edit the database configuration file, which involves more steps.
3. If you wish to optimize your Informix installation later by adding another dedicated network link between the filer and the Informix server, this can be easily utilized using a symbolic link. Again, the same thing is possible without the link, but it requires more steps.

Network Appliance strongly encourages the use of symbolic links to maximize flexibility and control over the Informix environment.

4. Installing INFORMIXDIR on a Filer

INFORMIXDIR is the environment variable specifying the location in the file system where Informix stores many of its executables and configuration files. On a normal UNIX installation, this is located at something like:

```
INFORMIXDIR=/export/home/ids9.2
```

It is sometimes advantageous for **INFORMIXDIR** to be located on a filer. This configuration creates issues that must be managed as well. Section 4.1 below covers these advantages and issues.

There are pros and cons to installing the entire Informix database server on the filer. These are outlined below.

4.1. Advantages

Placing INFORMIXDIR on a filer has the following advantages:

- With **INFORMIXDIR** on a filer, you can easily configure a manual failover server. If INFORMIXDIR and the Informix chunk files are located on the filer and a failure occurs in the primary UNIX server, a second UNIX box can be quickly brought online providing Informix database service. This is done by configuring the second server with the same host name and IP addresses, connecting it to the same network resources as the failed server, and pointing it at the same mount points as the primary server uses. This entire operation can be accomplished in 10 to 20 minutes. At this point the Informix server can be brought back online and production can continue while the primary server is repaired.
- **NOTE:** If desired, application cluster technologies from vendors such as Veritas, Legato, Integratus, IBM, and HP can be used to automate this procedure completely.
- With **INFORMIXDIR** on a filer, you can run multiple instances of IDS on separate database servers that use the same set of executables. This saves the space otherwise required to store multiple copies of the IDS executables on separate database servers. It also ensures that the same version of Informix is running on several different servers, since when they are all using the same copy, that simplifies configuration management.
- With **INFORMIXDIR** on a filer, you can easily set up multiple versions of IDS on the same server, and then switch between them. This is ideal in a lab environment, where multiple versions of IDS must be tested. This is covered in detail in Section 6.1 of this technical report.

4.2. Disadvantages

Placing INFORMIXDIR on a filer creates the following issues:

- In the event of a filer outage, the MSGPATH file and online.log file that Informix normally keeps in the **INFORMIXDIR** directory will not be accessible. Thus, there is a chance that information required to diagnose the problem might not be available. For this reason, you should relocate this file onto the local hard drive of the IDS server machine or onto a separate NetApp filer. Do this either by manipulating the symbolic links used to point the Informix server at the filer mount point or by modifying the **MSGPATH** parameter in the **onconfig** file to point to a location on the server.

5. Quick Start Guide

If you already are familiar with both Informix installation procedures for UNIX and basic operations of a Network Appliance filer, the following section quickly outlines the steps for migrating the Informix data onto a filer. Read on in this paper for the reasons and theory behind some of the recommended steps.

Assumptions

These steps assume that you wish to migrate existing Informix database storage to a NetApp filer. If you have not already installed Informix, follow the Informix installation instructions for your platform for getting Informix installed *using locally stored files* (not RAW disks), and then follow these steps for migrating the database. Furthermore, we also assume that the entire **INFORMIXDIR** directory (binaries and data files alike) are going to be moved. Further, we assume the Informix userid required for running the Informix database is "informix" and the

Informix group is also "informix". Finally, we assume that `${INFORMIXDIR}` is currently at `/export/home/Informix`.

5.1. Establish a Persistent Mount Point

Follow the steps outlined later for getting a persistent mount point to the filer that Informix will use for its data I/O. This usually involves editing the `/etc/vfstab` or equivalent file so that this mount point is established automatically anytime the UNIX system is booted. Be sure to specify the options listed in Section 3.5.

5.2. Migrating the Informix Software

After shutting down the Informix database ("onmode -yuk", for instance), copy the Informix files to the filer:

```
< as informix >
% cd $INFORMIXDIR
% onmode -yuk
% su root
# cd ..
# tar cf - Informix |(cd /ifx/test/home;tar xvFp -)
```

Note: the 'p' parameter on the tar extract ensures that those programs in the Informix directory with set-UID information associated with them will be correctly moved and set. It is always a good idea to double-check and be sure that these permissions did make it over. If not, delete the files that were copied, correct the problem preventing the set-UID information from being extracted, and try again.

Reset the pointers

Now create a symbolic link pointing to the info on the filer.

```
< continuing from the "tar" step above >
# mv Informix Informix.save
# ln -s /ifx/test/home/Informix Informix
```

5.3. Migrating the Informix Instance

Simplest Method to Migrate an Informix Database onto a NetApp filer

This method assumes you are currently using Informix on locally attached drives with symbolic links pointing to the disks themselves. By this point, the DBA should have a volume created on the filer to store the database. This volume should be correctly mounted to the database server. The following steps outline how to migrate the Informix data over to a Network Appliance filer.

1. Collect the general information about the current database configuration. Save this information to a file.
2. Bring down Informix.
3. Create new files that correspond to the original chunk pathnames.
4. "dd" the contents of each chunk to its new location.
5. Reset the symbolic links to point to the new location on the filer.
6. Bring up Informix.

Each of the steps listed above will be described below in more detail.

Step 1: Collect the output of an `onstat -d`, `onstat -c`, `onstat -b`, and `onstat -a`. Save information about where each chunk's symbolic link is pointing. In our example, the "chunks" portion of an `onstat -d` would look like the following:

```
chunks
address  chk/dbs  offset  size    free    bpages  flags  pathname
a402210  1 1 0      50000  27149          PO-   /database/IFX_CHUNKS/rootdbs.ln
a402448  2 2 0      100000 49947          PO-   /database/IFX_CHUNKS/custdbs.1.ln
a402528  3 3 0      100000 49947          PO-   /database/IFX_CHUNKS/orderdbs.1.ln
a402608  4 4 0      100000 49947          PO-   /database/IFX_CHUNKS/manudbs.1.ln
a4026e8  5 5 0      100000 49947          PO-   /database/IFX_CHUNKS/stockdbs.1.ln
a4027c8  6 6 0      100000 49947          PO-   /database/IFX_CHUNKS/itemsdbs.1.ln
a4028a8  7 7 0      100000 49947          PO-   /database/IFX_CHUNKS/calldbs.1.ln
a402988  8 8 0      100000 49947          PO-   /database/IFX_CHUNKS/ccallsdbs.1.ln
a402a68  9 9 0      100000 49947          PO-   /database/IFX_CHUNKS/catdbs.1.ln
a402b48  10 2 0     125000 124997         PO-   /database/IFX_CHUNKS/custdbs.2.ln
      8 active, 2047 maximum
```

The pathname to each chunk is really a symbolic link to the raw local raw device. An "ls -l" output shows the following:

```
# ls -l /database/IFX_CHUNKS
lrwxrwxrwx 1 root root          30 Jan  2 13:02 calldbs.1.ln ->/vol/calldbs.1
lrwxrwxrwx 1 root root          29 Jan  2 13:02 catdbs.1.ln ->/vol/catdbs.1
lrwxrwxrwx 1 root root          31 Jan  2 13:02 ccallsdbs.1.ln ->/vol/ccallsdbs.1
lrwxrwxrwx 1 root root          30 Jan  2 13:02 custdbs.1.ln ->/vol/custdbs.1
lrwxrwxrwx 1 root root          30 Jan  2 13:02 custdbs.2.ln ->/vol/custdbs.2
lrwxrwxrwx 1 root root          31 Jan  2 13:02 itemsdbs.1.ln ->/vol/itemsdbs.1
lrwxrwxrwx 1 root root          30 Jan  2 13:02 manudbs.1.ln ->/vol/manudbs.1
lrwxrwxrwx 1 root root          31 Jan  2 13:02 orderdbs.1.ln ->/vol/orderdbs.1
lrwxrwxrwx 1 root root          28 Jan  2 12:59 rootdbs.ln ->/vol/rootdbs
lrwxrwxrwx 1 root root          31 Jan  2 13:02 stockdbs.1.ln ->/vol/stockdbs.1
#
```

Step 2: Bring down Informix gracefully ("`onmode -sy`").

Step 3: Create new files on the filer that correspond with the original chunks. For each original chunk, create a new file. In our example, the new location for the chunks are in `/filer`, which is NFS mounted from the Network Appliance filer. Make sure the file permissions are set accordingly.

```
# cd /ifx/test/data
# touch calldbs.1
# touch catdbs.1
# touch ccallsdbs.1
# touch custdbs.1
# touch custdbs.2
# touch itemsdbs.1
# touch manudbs.1
# touch orderdbs.1
# touch rootdbs
# touch stockdbs.1
#
# chmod 660*
# chown Informix *
# chgrp Informix *
```

Step 4: Move the Informix data to its new location.

This should be done utilizing the `dd` command. The `dd` command can read raw devices as well as regular UNIX files. You will use the following arguments:

```
"dd if=(Input File)of=(Output File)skip=(Input File Offset)seek=(Output File Offset) count=(size of chunk in pages)bs=(Informix pagesize)"
```

The skip-and-seek parameter should always be the same. The value for the (Informix pagesize) can be found in the output of "onstat -b" at the bottom. It is the value of "buffer size".

The following is an example after running onstat -b:

```
# onstat -b | tail
```

```
Informix Dynamic Server Version 9.20 UC2 -- On-Line --Up 00:00:37 -- 11808KB
```

```
Buffers
address userthread flgs pagenum memaddr nslots pgflgs xflgs owner waitlist
0 modified, 0 resident, 200 total, 256 hash buckets, 2048 buffer size
#
the buffer size is 2048
```

In order to move the chunk /database/IFX_CHUNKS/rootdbs.In to /filer/rootdbs, the dd command should be:

```
"dd if= /database/IFX_CHUNKS/rootdbs.In of= /filer/rootdbs skip=0 seek=0 count=50000 bs=2048"
```

Repeat the above command for each chunk by modifying the parameters accordingly.

Step 5: Recreate the symbolic links.

This is simply the process of pointing the symbolic links to the corresponding data files that are now on the filer instead of the datafiles on local attached disk. The links must first be removed before pointing them to the data on the filer.

In our example, it would be done with the following commands:

```
# ln -s /filer/calldbs.1 /database/IFMX_CHUNKS/calldbs.1.In
# ln -s /filer/catdbs.1 /database/IFMX_CHUNKS/catdbs.1.In
# ln -s /filer/ccallsdbs.1 /database/IFMX_CHUNKS/ccallsdbs.1.In
# ln -s /filer/custdbs.1 /database/IFMX_CHUNKS/custdbs.1.In
# ln -s /filer/custdbs.2 /database/IFMX_CHUNKS/custdbs.2.In
# ln -s /filer/itemsdbs.1 /database/IFMX_CHUNKS/itemsdbs.1.In
# ln -s /filer/manudbs.1 /database/IFMX_CHUNKS/manudbs.1.In
# ln -s /filer/orderdbs.1 /database/IFMX_CHUNKS/orderdbs.1.In
# ln -s /filer/rootdbs /database/IFMX_CHUNKS/rootdbs.In
# ln -s /filer/stockdbs.1 /database/IFMX_CHUNKS/stockdbs.1.In
```

Step 6: At this point, all of the links should now be pointing to data that resides on the filer. To ensure that all the steps were done correctly, it is advised to run "oncheck -cr" before bringing Informix online. If oncheck reports errors, it is probably due to the links not being setup properly.

Once Informix is running on the filer to your satisfaction, you can get rid of the local disk array.

5.4. Modify the ONCONFIG File

One change is recommended to the **ONCONFIG** file. Modify the **MSGPATH** parameter to point to a path on a local drive or on another filer. The reasons for this are outlined later on. While this

step is not required, it is recommended. While the **ONCONFIG** file is being modified, check other settings to see if any other settings need modification.

If you are using or intend to use the ON-Bar (On-line Backup and Restore) facility of Informix, you may also need to modify the **BAR_ACT_LOG** parameter to be similar to the **MSGPATH** parameter. Refer to your Informix documentation for details.

5.5. Start the Database

It is recommended that you completely log out of the UNIX system now and log back in as the Informix user. Then start the Informix database server.

```
% oninit
```

The Informix server should start normally.

6. Installing INFORMIXDIR onto a Filer from Scratch

The procedure for installing INFORMIXDIR from scratch onto a filer requires two simple steps. First, be sure that a reasonable path on the IDS server machine is chosen that resolves to a location on the filer. Second, make sure that this path will automatically mount correctly whenever the IDS server machine is booted. To ensure this, the system administrator should follow these steps:

1. Choose a path where Informix will reside. For example, `"/export/home/informix"`. This becomes the setting for **\$(INFORMIXDIR)**.
2. Choose where INFORMIXDIR will reside on the filer. For example, `"/vol/vol1/home/informix"` on `"toaster"`. On the IDS server machine, this is `"toaster:/vol/vol1/home/informix"`. Make changes to the IDS server machine so that this point will always be mounted whenever the server boots. For more information on this issue, see 3.5, Filer NFS Mount Point.
3. The filer volume must be mounted prior to installing the Informix software. Mount the file system on the server using the mount command. For Solaris (and most other UNIX systems) this would be:

```
# mount /export/home/informix
```

If everything has been defined correctly, the space on the filer will now be available with the correct settings.

4. On the IDS server machine extract the Informix installation files and run the Informix installer per the installation instructions. Use as the install directory `"/export/home/informix"`. Since this is the location of the filer's file system, the Informix files will be extracted to and installed on the filer.
5. Finally, you need to move the message files off the filer and onto the IDS server machine's local disks or a different NetApp filer. One simple way to do this would be to modify the **MSGPATH** parameter in the `onconfig` file to point to a location on the server.

At this point, the Informix server software should be installed correctly with all Informix chunk files and executables located on the filer and correctly mounted on the IDS server machine.

6.1. Supporting Multiple Informix Versions on a Filer

It is possible to store multiple versions of Informix on a NetApp filer, and then bring up Informix engines that use these different versions of Informix. These different versions are all stored in separate directory structures and accessed by appropriately setting the `$(INFORMIXDIR)` variable.

7. Enabling Database Support on the Filer

It is a good idea to enable the feature of Data ONTAP (NetApp operating system software) that supports special error processing when connected to a database server. To do this, enter the following command from the filer's console or a telnet session:

```
vol options vol1 nvfail on
```

This will cause the filer to issue appropriate error messages in the `/etc/messages` file in the case of system failure that might affect the Informix database. The administrator will learn of these errors either by examining the message logs or by the autosupport e-mail notification feature of the filer.

In particular, this option enables some additional status checking when the filer goes through its initialization sequence at boot time to verify that the NVRAM is in a valid state. This should be the case for both a clean (normal) shutdown or a dirty (crash, power failure, etc.) outage. Only a failure of the NVRAM card itself should cause it to become invalid. If the content of NVRAM is found to be invalid, an error message will be put on the system console and into the filer log file, and all attempts by existing NFS clients to access the filer will fail with "stale filehandle" errors. This will affect all NFS clients accessing this volume, including the system running Informix server. These stale filehandle errors will cause the Informix server to hang or crash and the INFORMIX DBA will know that it is necessary to check that the state of the database is correct and valid.

Furthermore, additional protection is provided by an optional feature that renames certain files that the system administrator or DBA may wish to ensure are not accessible to the network until after they have been carefully examined. The file `/etc/nvfail_rename` controls this option. If it exists, the files in it are renamed by having the string `".nvfail"` appended to their original filenames. Since this occurs before the filer provides network service, these files will no longer have the same file name. Thus, the applications using them cannot automatically restart (including an Informix server that accesses these files).

The format of the `/etc/nvfail_rename` file is simply the name of the file, one per line, as viewed from the filer. So causing the file:

```
acme:/vol/vol1/home/Informix/etc/onconfig
```

to be renamed when the NVRAM failure is detected in version 4.3 of Data ONTAP or later would involve creating on the filer the file `/etc/nvfail_rename` with the one line:

```
/vol/vol1/home/Informix/etc/onconfig
```

Upon an NVRAM failure being detected, the file will be renamed to:

```
acme:/vol/vol1/home/Informix/etc/onconfig.nvfail
```

This will prevent the Informix server from opening the file, and thus make the DBA fully aware of the NVRAM failure.

8. Caveats

Informix supports the use of network-attached storage only in the context of Network Appliance filers. Further, NetApp has not tested this configuration with any version of UNIX other than Sun

Solaris, and has certainly not tested it with all of the combinations of hardware and software options available on Solaris. Informix has tested the use of NetApp filers with HP-UX, and various customers have tested them with other flavors of UNIX. There may be significant differences in your configuration that will alter the procedures necessary to accomplish the objectives outlined in this paper. If you find that any of these procedures do not work in your environment, please contact Network Appliance immediately.

9. Revision History

Date	Name	Description
11/1999	Bruce Clarke	Creation
4/10/2001	Ed Sassi	Update
4/16/2001	Maury Tiller & Bruce Clarke	Update



Network Appliance, Inc.
495 East Java Drive
Sunnyvale, CA 94089
www.netapp.com

© 2005 Network Appliance, Inc. All rights reserved. Specifications subject to change without notice. NetApp, NetCache, and the Network Appliance logo are registered trademarks and Network Appliance, DataFabric, and The evolution of storage are trademarks of Network Appliance, Inc., in the U.S. and other countries. Oracle is a registered trademark of Oracle Corporation. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.