



# Ingres II for Unix: Integrating with a NetApp Filer

Bob Jancer and Gurinder P. Brar | Network Appliance | TR 3096

## 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. [Purpose and Scope](#)
2. [Assumptions](#)
3. [Infrastructure](#)
4. [Configuration Overview](#)
5. [Deciding Where to Install II\\_SYSTEM - Filer or Server?](#)
6. [Running Ingres II with Databases on a Filer](#)
7. [Migrating an Existing Ingres II Installation onto the Filer](#)
8. [Migrating Existing Ingres II Databases onto the Filer](#)
9. [Ingres II Database Tuning Tips](#)
10. [Enabling Database Support on the Filer](#)
11. [Caveats](#)

[TR3096]

---

## 1. Purpose and Scope

This document describes the steps necessary to integrate Ingres II for UNIX with a NetApp filer. Specifically, we cover the following issues:

- Preparing both the UNIX server machine and the NetApp filer for the Ingres II installation
- Creating of an Ingres Administrator/User account (ingres)
- Installing of Ingres II with II\_SYSTEM on a NetApp filer or on local disk on the UNIX server machine
- Setting of default database path (II\_DATABASE) to a location on the filer
- Setting of default paths for Ingres II transaction log (II\_LOG\_FILE), checkpoint (II\_CHECKPOINT), journal (II\_JOURNAL), dump (II\_DUMP), and work (II\_WORK) to locations on the filer
- Migrating an existing Ingres II installation from local disk onto the filer
- Migrating existing Ingres II databases and related files from local disk onto the filer

## 2. Assumptions

We assume that you are familiar with Ingres II, as well as the operation of NetApp filers, and possess basic UNIX administration skills. All examples in this technical report are for Ingres II 2.0 running on IBM AIX 4.3. The examples contained in this paper may require modifications to run under your version of UNIX. This paper also assumes that you have available the *Ingres II Getting Started Guide* for your particular operating system version and generally follow its steps for installation and configuration. In particular, issues involving user IDs, group IDs, and the setting of database configuration parameters need to be resolved following these manuals. Where these documents and this technical report contradict, you should assume that the Ingres manual is correct. Please inform Network Appliance of any such contradictions so that this paper can be corrected. This paper describes the implementation of Ingres II into an already functioning AIX and Filer environment.

The sample scripts in this technical report assume the following:

- Name of the filer is "netapp".
- Ingres administrator account on UNIX server machine is "ingres".
- Ingres user group on UNIX server machine is "ingres".

- Volume on filer to store Ingres installation and database-related files is: "/vol/ingres\_vol".
- II\_SYSTEM
  - Local disk installation: path on UNIX server machine is "/export/home/ingres20".
  - Directory in volume on filer: "netapp:/vol/ingres\_vol/iisystem".
  - Mount point on UNIX server machine is "/mnt/ingres\_vol/iisystem".
- II\_DATABASE (Default location for database files)
  - Directory under II\_SYSTEM directory on filer: "netapp:/vol/ingres\_vol/iisystem/ingdata".
- II\_LOG\_FILE (Location for the transaction log file)
  - Directory under II\_SYSTEM directory on filer: "netapp:/vol/ingres\_vol/iisystem/inglog".
- II\_CHECKPOINT (Default location for full backup files)
  - Directory under II\_SYSTEM directory on filer: "netapp:/vol/ingres\_vol/iisystem/ingckp".
- II\_JOURNAL (Default location for incremental backup files)
  - Directory under II\_SYSTEM directory on filer: "netapp:/vol/ingres\_vol/iisystem/ingjnl".
- II\_DUMP (Default location used to store transactions during online backups)
  - Directory under II\_SYSTEM directory on filer: "netapp:/vol/ingres\_vol/iisystem/ingdump".
- II\_WORK (Default location for temporary files)
  - Directory under II\_SYSTEM directory on filer: "netapp:/vol/ingres\_vol/iisystem/ingwork".

Obviously, you will need to make the appropriate changes to the above settings to suit your environment.

### 3. Infrastructure

The following components and infrastructure are needed to run Ingres II for UNIX in conjunction with a filer:

- Ingres II UNIX server machine running AIX 4.3 (or other supported UNIX operating system)
- NetApp F700 or later series filer
- Network
- "ingres" administrator account on the Ingres II UNIX server machine and the NetApp filer
- "root" account on the Ingres II UNIX server machine
- Filer NFS mount point(s)
- One or more symbolic links on the Ingres II UNIX server machine referencing one or more directories or files on the filer (optional, but highly desirable)

#### 3.1. Ingres II UNIX Server Machine

You need Ingres II running on a UNIX server machine or workstation. We used Ingres II Enterprise Edition 2.0 and AIX 4.3 as the working platform to create this document. In your installation, be sure that your equipment satisfies the system requirements for running Ingres II. For more information on this issue, check the *Ingres II Getting Started Guide* for your target operating system. Other versions of UNIX are also supported but each will have their own idiosyncrasies and may require some modifications to the procedures outlined in this document.

If any problems arise, please contact your Network Appliance representative for assistance.

### 3.2. NetApp Filer

Any NetApp F700 (or later) series filer running Data ONTAP version 5.3 or later will work. The NFS license on the filer must be activated, and the NFS protocol must be set up and running.

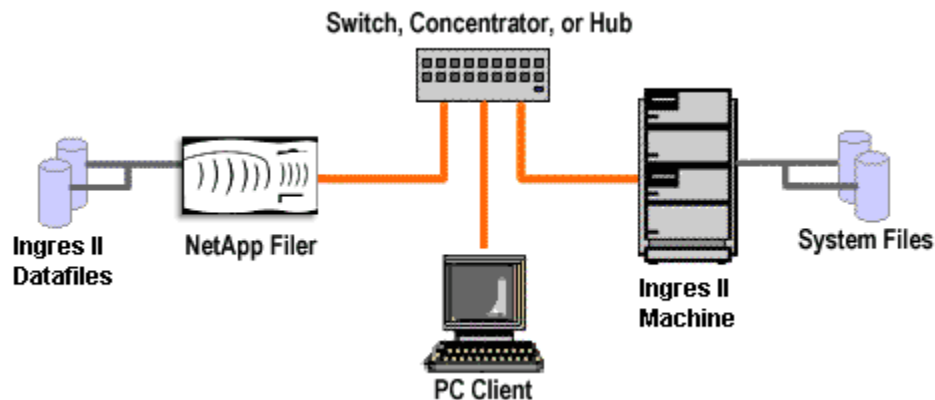
### 3.3. Network

You will need at least one network connection between the Ingres II UNIX server machine and the NetApp filer to enable database storage on the filer.

Network performance notes:

- It is **strongly** recommended that you use gigabit Ethernet interfaces to support your Ingres II database environment on the filer for optimal performance.
- Enable flow control and full-duplex operation on the Ethernet interfaces on BOTH the UNIX server machine and the NetApp filer.

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 Ingres II UNIX server machine and the filer. This can be done using a cross-over cable on most networking topologies (fast Ethernet and gigabit Ethernet being two examples we have tested successfully).

Dedicated network connections between the Ingres II UNIX server machine and the filer are recommended for the following reasons:

- Any issues of contention or latency are eliminated if the Ingres II UNIX server machine and the filer are the only two nodes in the network.
- Security is ensured by creating a private network connection. There is no issue of protecting the Ingres II data files from tampering, as would be the case on a shared network.

Each dedicated network connection between the filer and the Ingres II UNIX server machine can be thought of as functionally equivalent to an SCSI or FC-AL connection. (High-end networking protocols like gigabit Ethernet provide similar bandwidth as well.)

### 3.4. Ingres II User Account

If you have not done so already, start the installation process by following the Ingres II installation instructions for creating the "ingres" administrator user and group. For example, create a group "ingres" and then a user "ingres" that is included in the group "ingres". This account should be added on both the Ingres II UNIX server machine and the filer.

User "ingres" on the Ingres II UNIX server machine will own the installed software and can perform system management tasks such as starting and stopping Ingres II. Use the UNIX system command appropriate for your operating system to add this user. This will affect the Ingres II UNIX server machine only.

You will also 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 from the Ingres II UNIX server machine into the "/etc/passwd" file on the filer. (Alternatively NIS can be used on the filer. For more information on this issue, see the *Data ONTAP System Administrator's Guide*)

### 3.5. Filer NFS Mount Points

You need to create volumes/directories on the filer for the Ingres II installation as well as the appropriate Ingres II data-related files you wish to store there (e.g. database, checkpoint, journal, dump, work, log) with security set to include:

```
ingres rwx
```

In addition, you need to create at least one mount point on the UNIX server machine pointing to the highest-level volume/directory used for your Ingres II installation and database-related files. Depending on how you set up your system, this could be a single mount point if you install all Ingres system and database-related files under the one directory represented by the Ingres II environment variable, `II_SYSTEM`, or it may be multiple mount points if you separate the Ingres II system files from the database-related files into diverse volumes on the filer.

Some nondefault NFS mount options should be set when configuring the method with which the Ingres II UNIX server machine should 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.

Option	Description
hard	This option says that the mount point should never time out and that the Ingres II UNIX server machine should not run without it. This will cause the Ingres II UNIX server machine to hang if the filer is not responding to NFS for any reason. If the Ingres II UNIX server machine is booting and the filer is not found, then it will not complete the boot and Ingres II will not start. If it is already up and running and the filer disappears, all I/O to and from the filer will be suspended until the filer is available again.
intr	This option allows operator generated keyboard interrupts to kill a process that is hung while waiting for a response from a hard-mounted file system (in this case the filer). This is sometimes useful.

vers	This option is supported in recent releases of UNIX. It specifies which NFS version 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 should be used. For NFS over local area networks, UDP offers less overhead (and therefore better performance) than TCP. However, if your network connection path between the NetApp filer and the Ingres II host is prone to lose packets, drop frames, or introduce checksum errors, then TCP can improve performance compared to UDP. We therefore recommend that you run using UDP on a dedicated network connection with a crossover cable between the Ingres II server and the filer. For more information on this issue, see <a href="#">3.3. Network</a> . If you use UDP, be sure to enable UDP checksums on the Ingres II UNIX server machine.
suid	This option tells the server that it should honor the set-uid bit on files mounted at this mount point. If you have any of the Ingres II executables located on the filer, then this setting is important. If you are putting only the database files on the filer, then this option can be omitted. If you use this option, you must also export the file system with the "-anon=0" option. For example, the "/etc/exports" file on the filer should read something like: <pre> /vol/vol0 -anon=0,root=somepc /vol/ingdata -anon=0 /vol/ingckp -anon=0 </pre>

### 3.6. Setting Ingres II Volume Security Styles on the Filer

You must ensure that the volume(s) to be used for Ingres II storage on the filer has either "UNIX" or "mixed" security style by using the qtree security command on the appropriate volume(s) as indicated below:

```
qtree security /vol/ingres_vol/ [ UNIX | mixed ]
```

### 3.7. Link Issues

It is often desirable to use symbolic links to map the Ingres II II\_SYSTEM directory structure to the filer. This is true for three reasons:

1. If you are migrating from local disk to a filer, the use of symbolic links prevents you from having to remove and reinstall your Ingres II installation.
2. If you later wish to reorganize the location of your Ingres II installation and/or related files, you can do so easily by simply taking the Ingres II server down, moving the files, and then editing the link. The symbolic link thus provides an extra layer of abstraction to your file system. Without the symbolic link, you would need to remove and reinstall your Ingres II installation.
3. If you wish to later optimize your Ingres II installation by adding another dedicated network link between the filer and the Ingres II server, this can be easily accomplished using a symbolic link. Again, the same thing is possible without the link, but it requires more steps.

## 4. Configuration Overview

Traditional installation of Ingres II on a UNIX platform called for the storage of all Ingres II components—executables, database files, transaction log files, etc., on storage local to the UNIX server machine with the various Ingres database-related default directories located on separate physical devices to maximize availability and file recovery.

Network Appliance's approach uses network-attached storage, which allows you to store both the Ingres II installation/ product files as well as database and database-related files in a single volume consisting of multiple disks on a filer to accomplish a simpler, faster, easier to administer, and more reliable solution.

The functionality and flexibility of the standard Ingres II installation procedure makes it simple to set up the default Ingres installation (II\_SYSTEM) and database directory paths (as well as checkpoint, journal, work, dump, and temporary directories) to be directed to the filer rather than to the UNIX server machine local disk. As described in section 5 below, you may choose to install Ingres II with II\_SYSTEM located either on the filer or on local disk on the UNIX server machine. The advantages of installing II\_SYSTEM on the filer are described below in section 5 as well.

### 4.1. Volume Configuration

The key to ease of use and manageability of an Ingres II installation and database files stored on a filer is the fact that the entire Ingres II installation can be stored on one volume, requiring minimal attention by the DBAs and system administrators to ensure high performance. However, some physical design considerations need to be made upfront that will ensure these benefits.

#### Databases

To enable the future use of NetApp Snapshots as part of the Ingres II checkpoint backup procedure ("ckpdb") for your databases, ideally each database should be stored on a separate volume of the filer that is not the root volume. However, for all practical purposes in configurations consisting of more than one database this would not be practical due to the number of additional disks that would be required.

#### Transaction Log File

When using a NetApp filer, it is not necessary to store the primary transaction log file in a different volume separate from the database(s). You can utilize the inherent capabilities of the filer and the Data ONTAP operating system and store the primary log file in the same volume on the filer as both the Ingres II installation and other database-related files. In addition, there are also multiple options available during installation related to the setting up of transaction log files. You may select from these options if desired:

- Option #1: You can mirror the transaction log file on the filer by setting up a two disk volume specifically for the log file.
- Option #2: You can utilize the backup transaction log capability of Ingres II to set up the primary log file on the filer with the backup transaction log on either a separate filer or local disk.

#### Backup Transaction Log File (II\_DUAL\_LOG)

With dual logging enabled, the Ingres II installation has an alternate log file. With dual logging a media failure on one of the logs will not result in the loss of data or the interruption of service. If

one of the log disks fails, the logging system automatically switches over to access the other log without interrupting the application. Thanks to RAID4 protection on the NetApp filer, dual logging is not necessary when you use the filer to store your primary log file since the filer will be able to handle the failure of a physical storage disk within the volume without interrupting the application. However, this is not meant to imply that you cannot choose to use the dual-logging capability if you wish to do so. Please keep in mind, however, that there may be additional performance overhead introduced when using this feature.

### Checkpoint, Journal, and Dump Files

The Ingres variables II\_CHECKPOINT, II\_JOURNAL, and II\_DUMP represent the default locations used for full and incremental online Ingres II database backup files. Traditional installation of Ingres II on a UNIX platform suggests that these storage locations should be assigned to a physical device separate from the physical device where the Ingres II database files represented by II\_DATABASE are stored. This is not relevant, however, when using a NetApp filer for Ingres II database storage. The filer provides reliability and availability of the Ingres II files. Therefore, all of these directories can be stored in the same volume on the filer as the II\_DATABASE directory.

### Ingres Work (Temporary) Files

The Ingres variable II\_WORK stores the default location used for all temporary files created during external sorts and other DBMS server operations that require large amounts of temporary file space. Traditional installation of Ingres II on a UNIX platform suggests that you designate a separate storage device as the default location for these files, but as with the other Ingres II database-related files stored on a NetApp filer, this is not necessary. You can reliably store your temporary Ingres II files represented by the II\_WORK directory in the same volume on the NetApp filer as the other Ingres II database-related directories described above. In any event, your choice of the default work location does not affect your ability to recover lost data in the event of a storage device failure.

### Root Volume

The root volume on the filer should be its own volume of just two disks. In our opinion, the extra reliability and flexibility gained by having essentially a mirrored root volume outweighs the cost of consuming two extra drives.

The benefits conferred by this configuration are:

- Slightly higher reliability of the root volume due to mirroring as opposed to RAID4 and a reduced amount of file system activity (less likely for an inadvertent change to cause downtime). Should a database volume fail, having a separate, still functioning root volume will save valuable time in the recovery process.
- Using a small two-disk root volume and locating all user data on other volumes makes physically moving foreign volumes between two filers easier.

### Suggested Volume Configuration

Volume	Description	Raid Group Size Limit	Number of Disks (recommended)
/vol/vol0	Root volume	14	2
/vol/ingres_vol	Ingres II system	14	7 or more (preferably 14)



	volume		
--	--------	--	--

## 5. Deciding Where to Install II\_SYSTEM - Filer or Server?

This section describes the process of installing Ingres II and setting up the required Ingres II parameters.

Before you begin the installation, you must decide whether to have the II\_SYSTEM directory located on the Filer or the UNIX server machine. There are advantages to having II\_SYSTEM located on the filer; they are described in section [5.1. Advantages - II\\_SYSTEM on a filer](#) below.

It is important to note that the installation procedure described below remains the same whether you choose to locate II\_SYSTEM on the filer or on the UNIX server machine. The actual distinction is made when you originally set the II\_SYSTEM variable prior to the installation as a mount point to a volume on the filer or to a directory path on the local UNIX server machine disk.

**Note:** The items described in section [3. Infrastructure](#) are required in order for this to work. See the *Ingres II Installation Guide* for more information.

### 5.1. Advantages - II\_SYSTEM on a Filer

- With II\_SYSTEM on a filer, you can easily configure a manual failover server. This will be the subject of a future technical report.
- You could set up II\_SYSTEM on the filer in its own volume of just two disks. This would, in effect, provide a mirrored II\_SYSTEM volume with slightly higher reliability.

### 5.2. Issues - II\_SYSTEM on a Filer

- You have an additional point of failure in your Ingres II installation. If you plan to store Ingres II database files on the filer anyway, this is probably not a serious problem. Loss of the filer would temporarily bring down the database anyway.
- Since some Ingres II programs may need to run with specific permissions, the mount point exported by the filer to the Ingres II UNIX server machine must enable this provision.

### 5.3. Installing Ingres II (New Installation)

Follow the procedures below to perform a new installation of Ingres II system files onto the filer:

1. Before beginning with the Ingres II installation procedure, create the required volume(s) on the filer along with the desired mount point(s) on the UNIX server machine. This is essential since the installation procedure will require existing paths for the Ingres II default storage locations in order to complete successfully.
2. Create the parent directory for the Ingres II installation in a directory in the volume on the filer using the appropriate operating system command. This will become the II\_SYSTEM environment variable for this Ingres II installation and will be used by the installation procedure to set up Ingres II in the proper locations:

```
mkdir /mnt/ingres_vol/iisystem
```

3. Set the II\_SYSTEM environment variable to the directory created above using the appropriate command relevant to your operating system.

4. Create a directory named "ingres" under the II\_SYSTEM parent directory created in the step above. This will be the location where the Ingres II product files get installed:

```
cd $II_SYSTEM
mkdir ingres
```

5. Change your working directory to the newly created ingres installation directory:

```
cd $II_SYSTEM/ingres
```

6. Load the Ingres II "install" directory from the distribution media as per the Ingres II standard installation documentation.

```
tar xvf [path to distribution media Ingres II tar file] install
```

7. Install Ingres II using the standard Ingres II installation utility to install the software into II\_SYSTEM:

```
install/ingbuild
```

8. To configure the installation to support database storage on the filer, you will need to supply appropriate responses to the series of interactive prompts displayed as part of the standard installation procedure.

The section below describes how to tailor your replies to set up the database, transaction log, checkpoint, journal, dump, and work files on the filer regardless of the location of II\_SYSTEM.

## 5.4. Setting Up Default Database and Related File Storage on the Filer

**\*\*\* IMPORTANT NOTE:** You cannot change the values you assign to the default Ingres storage location variables below once the installation/setup procedure has been completed, so choose them wisely. However you will always have the capability to create additional storage locations and assign individual databases to them after this initial installation. For information on creating additional storage locations, refer to the *Ingres II Database Administrator's Guide*.

\*\*\* User responses to the installation procedure prompts are noted in **bold** below.

Please enter the default location for the Ingres database files:

II\_DATABASE: ***/mnt/ingres\_vol/iisystem/ingdata***

**NOTE:** This will be the default path for the location of the Ingres II master database (iibddb) as well as all newly created databases.

Please enter the default location for the Ingres checkpoint files:

II\_CHECKPOINT: ***/mnt/ingres\_vol/iisystem/ingckp***

**NOTE:** The Ingres variables II\_CHECKPOINT, II\_JOURNAL, and II\_DUMP are used to store the default locations used for full and incremental online backup files. Once assigned, the values of II\_CHECKPOINT, II\_JOURNAL, and II\_DUMP can only be changed by removing and re-installing Ingres. Refer to the *Ingres Database Administrator's Guide* for information on moving the default

backup locations. Refer to the chapter on "Backup and Recovery" in the Ingres Database Administrator's Guide for instructions on backup and recovery of Ingres databases.

Please enter the default location for the Ingres journal files:

II\_JOURNAL: */mnt/ingres\_vol/iisystem/ingjnl*

Please enter the default location for the Ingres dump files:

II\_DUMP: */mnt/ingres\_vol/iisystem/ingdmp*

Please enter the default location for the Ingres work files:

II\_WORK: */mnt/ingres\_vol/iisystem/ingwrk*

**NOTE:** Once assigned, the value of II\_WORK can only be changed by removing and reinstalling Ingres. Refer to the *Ingres Database Administrator's Guide* for instructions on moving the default work location in this manner.

## 6. Running Ingres II with Databases on a Filer

There are many advantages to storing Ingres II database and related files on a filer. The section below covers these advantages.

### 6.1. Advantages

Running an Ingres II instance with data files on a filer has the following advantages:

- Backup performance can be significantly improved.
- Administration and tuning requirements are lower when using a NetApp filer compared to a local disk configuration. For example, increasing the size of the file system using most local disk setups is a difficult and complicated process, usually involving a reboot of the host machine. With a NetApp filer, it is a simple operation that takes only a few minutes and requires no downtime to either the NetApp filer or the Ingres II server machine. Also, many of the tasks associated with balancing the load between table space containers can be eliminated because of the high performance of the filer.
- Reliability of the Ingres II database can be improved. For example, loss of a local disk will typically be more disruptive than loss of a disk on a NetApp filer.
- Off-loading the I/O to a network interface will free up some of the Ingres II server's CPU capacity.
- Write performance is typically increased by 10%–30% over well-tuned local disk configurations. (**Note:** This is very environment specific.) Read performance using a NetApp filer should be very close to, or slightly better than, similar local disk configurations.

## 7. Migrating an Existing Ingres II Installation onto the Filer

**\*\*\* IMPORTANT NOTE:** Always ensure that you have fully recoverable backup copy of your II\_SYSTEM directory tree as well all databases and database-related files before proceeding with the migration procedures described below.

### 7.1. Migrating an Existing II\_SYSTEM Directory onto the Filer

Migration of the II\_SYSTEM directory tree can be accomplished using standard operating system procedures. This method involves the actual moving of the files within this directory tree to a volume on the filer and the setting of the II\_SYSTEM Ingres variable as a link to a mount point to the appropriate volume on the filer rather than as an actual local UNIX server machine disk directory. By performing the migration in this manner, II\_SYSTEM will continue to look the same to all Ingres II utilities and programs as it did prior to the migration.

1. As user "ingres", determine the directory path represented by the II\_SYSTEM system variable:

```
echo $II_SYSTEM
```

You will see the actual directory path listed on the line below the command. Use this directory path in the following procedure involving the II\_SYSTEM directory. For our example below, assume that II\_SYSTEM corresponds to "/export/home/ingres20".

2. Shut down the Ingres II installation as the administrative user "ingres":

```
ingstop
```

3. Create a volume on the filer and mount point on the UNIX server machine for the new II\_SYSTEM location. Make a change to the Ingres II UNIX server machine so that the appropriate volume on the filer will always be mounted to the appropriate mount point on the UNIX server machine whenever the server machine boots.
4. Mount the file system on the server using the appropriate mount command for your operating system. In our example using AIX 4.3, this would be:

```
mount -o  
hard,intr,vers=3,proto=udp,suid netapp:/vol/ingres_vol/iisystem /mnt/ingres_vol/iisyste  
m
```

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

5. Move the II\_SYSTEM directory tree over to the filer:

```
cd /export/home/ingres20 (which is the physical location of II_SYSTEM in our  
example but will most likely differ in your installation)  
mv * /mnt/ingres_vol/iisystem
```

6. Verify that all of the directories/files have actually been moved to their new location on the filer
7. Remove the original II\_SYSTEM directory from the UNIX server machine local disk:

```
rmdir /export/home/ingres20
```

8. Create a link so that Ingres II will be able to find the new II\_SYSTEM path when it starts back up.

```
ln -s /mnt/ingres_vol/iisystem /export/home/ingres20
```

9. Start the Ingres II installation:

ingstart

The Ingres II installation should now be ready for normal operation in its new location.

## 8. Migrating Existing Ingres II Databases onto the Filer

**\*\*\* IMPORTANT NOTE:** Always ensure that you have fully recoverable backup copies of your database and related files before proceeding with any of the procedures below involved in migrating your database and related files.

When deciding on the procedures to be used to migrate Ingres II database and related files to the filer, you must keep in mind that the values of the default storage locations as defined during the original Ingres II installation cannot be changed without removing and reinstalling Ingres II. These non-changeable values include II\_DATABASE, II\_CHECKPOINT, II\_JOURNAL, and II\_DUMP. The II\_WORK default storage location is the exception to this rule. It can be changed if you follow the procedure described in the *Ingres Database Administrator's Guide*.

There are two methods available for migrating Ingres II database files to the filer:

1. The first is through the use of standard Ingres II utilities. This procedure is described in section [8.1. Migrating Databases using Ingres II Utilities](#) below.
2. The second is through the modification of any or all of the above default Ingres II storage locations using operating system procedures. This method involves the actual moving of the files within these directories to volumes/directories on the filer and the setting of the original Ingres default storage location values as links to the mount points to the volumes on the filer rather than actual local server machine disk directory paths. By performing the migration in this manner, these locations would continue to look the same to all Ingres II utilities and programs as they did prior to the migration. This procedure is defined in section [8.3. Migrating Databases Using Operating System Procedures](#) below.

There are also two methods available to migrate the files contained in the II\_DATABASE, II\_CHECKPOINT, II\_JOURNAL, II\_DUMP, II\_WORK, and II\_LOG\_FILE directories to the filer and have these Ingres II environment variables represent locations on the filer rather than local disks:

The first would be to perform the procedures described in section [8.3. Migrating Databases Using Operating System Procedures](#) below.

The second would be to remove and reinstall Ingres II and reset the default storage locations pertaining to these variables to point to locations on the filer.

The following sections describe both methods above in detail to accomplish the migration of existing Ingres II databases from local disk onto the filer. The example below assumes that there are multiple "locations" associated with the Ingres II databases. **Note:** The items described in section [3. Infrastructure](#) are required in order for this to work.

**Note regarding the migration of database and related files utilizing alternate locations to the filer:** Regardless of which method you use below to accomplish the migration of your Ingres II database files, it is *imperative* that you ensure that **ALL** of the locations (both primary and alternate) for a specific "type" of Ingres II file (meaning all database locations and/or all dump file locations, etc.) exist on either local server disk **OR** on the filer. DO NOT attempt to split the locations between the local server and filer.

## 8.1. What about Alternate Locations?

The setup and placement of alternate locations on separate physical storage devices for database, checkpoint, etc. files may play a role in a traditional Ingres II installation on UNIX for reasons related to performance, individual local disk storage limitations, and so on. However, these concerns are not relevant when using a NetApp filer for storage of the Ingres II database and related files.

There is one advantage to using alternate locations for database storage, however, that does remain relevant. It is related to the backing up of your databases using the standard Ingres II "ckpdb" database backup utility. If a database has multiple locations (primary and alternate) associated with it, this utility will allow you to specify that the locations be backed up in parallel, thereby shortening the time it would take to complete the backup.

Therefore you can choose to retain the alternate locations during the migration of the Ingres II installation to the filer, setting up the alternate locations as symbolic links to directories in the same volume as the primary volume/directory used for the storage of the Ingres II databases on the filer.

There are standard Ingres II utilities available that will allow you to migrate an existing database comprised of multiple locations on local server disk to a database consisting of a single location in one volume on a filer. These are described below.

\*\*\* **Note:** If you choose to remove the alternate locations from the Ingres II installation during the migration to the filer, you can use the standard Ingres II "unloaddb" utility to do so. Please refer to the *Ingres II Database Administrator's Guide* for information about this utility.

## 8.2. Migrating Databases Using Ingres II Utilities

To accomplish the migration of Ingres II database files from local server machine disk to the filer, you can use either of two standard Ingres II utilities—"unloaddb" or "copydb". An explanation of these utilities is beyond the scope of this technical document. For further information regarding the use of these utilities, please refer to the *Ingres II Database Administrator's Guide*.

## 8.3. Migrating Databases Using Operating System Procedures

1. First, you need to determine all the locations (both primary and alternate) available to the Ingres II databases in this installation. You can use the standard Ingres II ISQL utility to obtain this information from the Ingres II master database (iidbdb) by following the procedures outlined below.

```
isql iidbdb
```

From within the ISQL window:

```
select location_name, location_area from iilocation_info where data_usage = "Y";
```

**Note:** Use the appropriate "Go" command based on your terminal type to process the above command.

You should now see displayed on the screen a list of all the possible database data locations setup within this Ingres II installation. It should look similar to:

location_name	location_area
ii_database	II_DATABASE
location2	/export/home2/ingres20/ingloc2

2. **Note:** There will always be at least one database location. This will be the default database location and it will be listed as II\_DATABASE. This is an Ingres environment variable that represents the default directory path for Ingres II database files and is used exclusively by Ingres II programs and utilities. In order to perform the remaining steps, below it will be necessary for you to know the actual directory path that this variable represents. To determine this, run the Ingres II "ingprenv *[variable]*" utility as user "ingres":

```
ingprenv II_DATABASE
```

You will see the actual directory path listed on the line below the command. Use this directory path in any of the following procedures that would involve the II\_DATABASE directory. For our example below, assume that II\_DATABASE corresponds to "/export/home/ingres20/ingdata".

3. Shut down the Ingres II installation as the administrative user "ingres":

```
ingstop
```

4. Create directories in the Ingres II volume on the filer for all of the database locations indicated in the ISQL output above. Make changes to the Ingres II UNIX Server machine so that the appropriate volumes on the filer will always be mounted to the appropriate mount points on the UNIX server machine whenever the server machine boots.

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

5. Move the database files over to the filer, one location at a time:

```
cd /export/home/ingres20/ingdata (which is the physical location of II_DATABASE
in our
```

```
example but will most likely differ in your installation)
mv * /mnt/ingres_vol/iisystem/ingdata
```

```
cd /export/home2/ingres20/ingloc2 (which is the physical location of "location2" in
our
```

```
example but will most likely differ in your installation)
mv * /mnt/ingres_vol/iisystem/ingloc2
```

6. Verify that all of the files/directories have actually been moved to their new location on the filer.
7. Remove the original database directories for all of the database locations from the UNIX server machine local disk:

```
rmdir /export/home/ingres20/ingdata
rmdir /export/home2/ingres20/ingloc2
```

8. Create links so that Ingres II will be able to find all of these directories when it starts back up.

```
In -s /mnt/ingres_vol/iisystem/ingdata /export/home/ingres20/ingdata
In -s /mnt/ingres_vol/iisystem/ingloc2 /export/home2/ingres20/ingloc2
```

9. Start the Ingres II installation:

```
ingstart
```

The database should now be ready for normal operation in its new location.

## 8.4. Migrating Checkpoint, Journal, Dump, and Work Files Using Operating System Procedures

To migrate the database checkpoint, journal, dump, and/or work files from local disk to the filer, you would use the same procedure as described in section 8.3 above, making slight modifications to the commands described in steps 1 and 2 to make them appropriate for the type of Ingres II files to be migrated. These modifications are outlined below:

II_CHECKPOINT	(Step 1) select location_name, location_area from iilocation_info where ckpt_usage = "Y" (Step 2) ingprenv II_CHECKPOINT
II_JOURNAL	(Step 1) select location_name, location_area from iilocation_info where jrnl_usage = "Y" (Step 2) ingprenv II_JOURNAL
II_DUMP	(Step 1) select location_name, location_area from iilocation_info where dump_usage = "Y" (Step 2) ingprenv II_DUMP
II_WORK	(Step 1) select location_name, location_area from iilocation_info where work_usage = "Y" (Step 2) ingprenv II_WORK

Once you have completed the migration procedure, take whatever steps you feel are necessary to ensure that all files are accessible to Ingres II exactly as they were prior to the migration before proceeding to use the system.

## 9. Ingres II Database Tuning Tips

### Database Blocksize

To take full advantage of the filer's WAFL file system, it is **strongly** recommended that you set your Ingres II database blocksize to either 4k or a multiple of 4k.

## 10. Enabling Database Support on the Filer



It is a good idea to enable the [Data ONTAP](#) feature (NetApp's operating system software) which 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 ingres_vol 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 Ingres II 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 contents of NVRAM are 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, including the system running the Ingres II server. These stale filehandle errors will cause the Ingres II instance to hang or crash, and the Ingres II DBA will know that it is necessary to ensure 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 file names. Since this occurs before the filer is providing network service, these files will no longer have the same file name as previously. Thus, the applications using them cannot automatically restart.

The file in the example below corresponds to the Ingres master database (iibdadb), without which Ingres will not start. This master database is always located in the II\_DATABASE directory, wherever that directory may exist (local server machine disk or filer). In the example below, the assumption is that the volume named "ingdata" on the filer corresponds to the mount point for II\_DATABASE on the UNIX server machine in the Ingres II installation.

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:

```
\\netapp\ingres_vol\iisystem\ingdata\ingres\data\default\iibdadb\aaaaaaaa.cnf  
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 line:  
/vol/ingres_vol/iisystem/ingdata/ingres/data/default/iibdadb/aaaaaaaa.cnf  
Upon an NVRAM failure being detected, the file will be renamed to:  
\\netapp\ingres_vol\iisystem\ingdata\ingres\data\default\iibdadb\aaaaaaaa.cnf.nvfail  
This will prevent the Ingres II server from opening the file, which is required for Ingres to complete  
its normal startup process, and thus make the DBA fully aware of the NVRAM failure.
```

## 11. Caveats

Use of network attached storage for storing Ingres II database files is only approved by Computer Associates on Network Appliance filers. Do not use this configuration on any other type of network attached storage device. Further, the configuration presented in this paper has been tested by Network Appliance using only in a limited set of hardware and software options.

Therefore, your experience may differ from that presented here. If you have any problems with the techniques shown in this technical report, please contact [the author](#).



**Network Appliance, Inc.**  
495 East Java Drive  
Sunnyvale, CA 94089  
[www.netapp.com](http://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.