

# **DAS to SAN Migration**

**A Technical White Paper**

**September 2002**



**Scott Hanson**

**ArdentA Limited**

**<http://www.ardenta.com>**

---

## TABLE OF CONTENTS

<b>Introduction</b>	<b>2</b>
<b>Why Migrate?</b>	<b>2</b>
<b>Typical DAS Configuration</b>	<b>3</b>
<b>Migration Phase</b>	<b>4</b>
<b>Migration Process</b>	<b>5</b>
<b>New SAN Configuration</b>	<b>7</b>
<b>About the Author</b>	<b>8</b>
<b>About ArdentA</b>	<b>8</b>
<b>Glossary</b>	<b>9</b>
<b>Trademarks</b>	<b>10</b>

## Introduction

Many organisations are recognising that the Total Cost of Ownership (TCO) of Direct Attached Storage (DAS) is outweighing the traditionally prohibitive cost of implementing a Storage Area Network (SAN).

According to "The Storage Report - Customer Perspectives & Industry Evolution" by Merrill Lynch & Co., a leading financial management company, and McKinsey & Company, a leading international management consulting firm, the TCO of two terabytes of DAS can cost up to two and a half times the cost of the same amount of networked storage - SAN or NAS.

The task of migrating a live system from one disk array to another can however appear daunting. The process can be fraught with pitfalls and can potentially lead to a significant downtime for the live server.

The migration methods in this white paper have been proven by Ardenta consultants in migrating many servers with minimal downtime. All the downtime required can be scheduled at times convenient to the user.

## Why Migrate?

There are many reasons to migrate to a SAN environment, of which the following are just a few:

### Improved disk utilisation.

Companies using DAS may be forced to purchase new hardware to provide storage capacity even though free capacity is available elsewhere but cannot be used.

In a SAN environment, storage can be "virtualised", so that is made to look and behave like a complete disk volume to an operating system, even if it is only a disk partition or stripe set. This means that the customer application drives the storage allocation, not the other way around.

### Ease of growing storage pools

Additional disk can be added to a SAN or NAS solution without the need for downtime on the server. Traditionally, this is not possible using a DAS solution.

## Ease of management

SAN management solutions such as Hewlett-Packard's SAN Manager product have removed the problem of visualising how much disk is available and where it physically is located. This allows the storage administrator to execute capacity planning and other maintenance exercises more easily.

## Centralised backup and recovery

A main advantage of using a central storage pool is the ease of backup and recovery. Utilising the latest NDMP<sup>1</sup> technology, serverless backup and recovery techniques can be employed, thus allowing backup with no interruption to live service. In a DAS environment, backup places excessive demands upon a live server. These demands can be mitigated using a SAN solution.

## High availability

Modern disk arrays allow for multiple mirrors of data transparent to the host server. This can mean outage-free disk replacement and sophisticated disaster recovery solutions that can be beneficial to a customer's business.

## Improved Bandwidth

Fast/Wide SCSI has a sustained I/O rate of 14 MBps on a High Speed Connect bus. By comparison, 1 GB Fibre Channel has a sustained I/O rate of 92 MBps. Therefore the migration from DAS to SAN can yield much higher bandwidth. This translates to significant improvements in application response time.

## Dual Hosting

To facilitate migration, and if hardware permits, an effective method involves connecting the servers to both the existing DAS environment and the new SAN. This allows for fast and direct copying between the DAS and SAN environment and little, if any, downtime for cutover. There will always be some downtime to connect the new SAN and to disconnect the old DAS solution, but this can be scheduled for a time when it will have least impact.

---

<sup>1</sup> Network Data Management Protocol

## Typical DAS Configuration

The following diagram shows a typical DAS configuration with one storage array connected to multiple UNIX hosts.

The migration from DAS to SAN often proves problematic. In particular, migrating a live service from one environment to another can typically result in many hours, if not days, of downtime.

This can be overcome by executing the migration carefully, utilising software mirroring where appropriate.

### Take a Backup

Since a migration involves significant data movement, a full backup of all databases and file systems should be taken before initiating the migration activities.

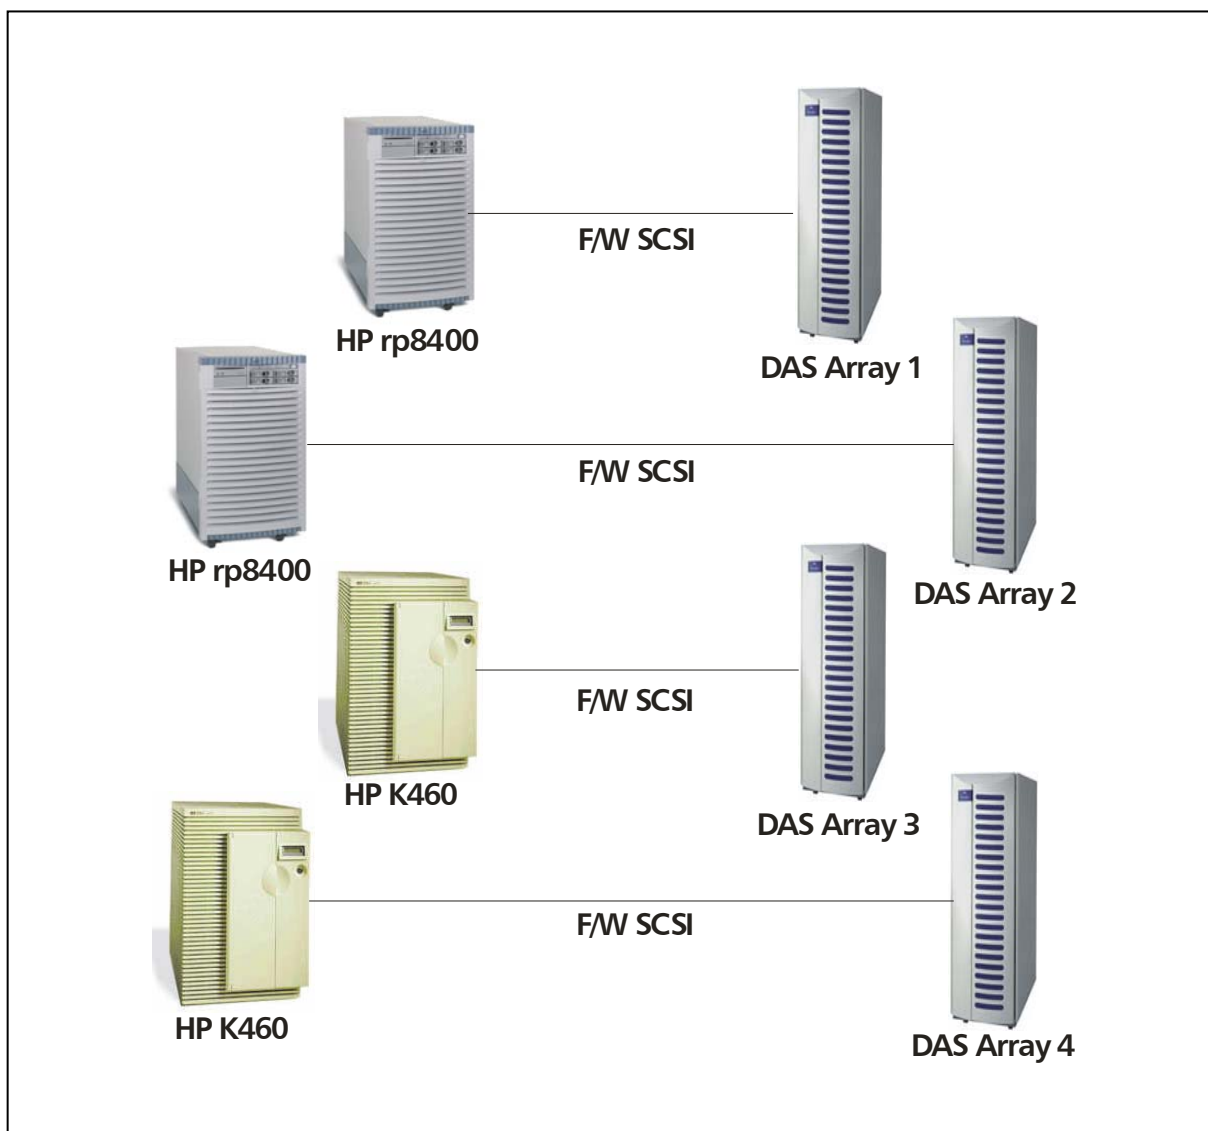


Figure 1: DAS Configuration

## Migration Phase

The new SAN array should be pre-configured using an agreed configuration in which the physical spindles are split into the same size hyper-volumes<sup>2</sup> as those in the DAS array.

The Fibre Channel interface cards in the SAN Array should be connected to the Brocade SAN switches.

Each host should be installed with a new Fibre Channel Host Bus Adapter (HBA). As part of the installation, the server should be tuned to allow the HBAs to function.

Once the hosts are restarted, the configuration of the SAN, also known as "zoning" should be completed. The "zoning" allows the relevant hyper-volumes to be presented to the hosts.




Figure 2 : Migration Configuration

<sup>2</sup> A hyper-volume is a logical section of disk or disks, which are presented to the host as a single SCSI disk.

## Migration Process

In order to migrate the data, software mirrors of the data should be established between the disk on the DAS array and the disk on the SAN array.

The advantage of this approach is that the data from the existing DAS array is kept intact until all mirrors are established on to the new SAN array. Once this has been done, each mirror copy from the DAS array can then be removed safely.

In the following two scenarios, this was done using HP Logical Volume Manager (LVM) MirrorDisk and IBM Informix Dynamic Server (IDS) Mirroring. However, almost all major operating systems and relational database management systems have equivalent functionality, so it is easy to see how the principle could be extended to most configurations.

In both cases, a section of disk (volume group or chunk) in the DAS array was mirrored to an identically sized piece of disk in the SAN array. This mirror was left in an established state until the cutover time.

### Technique for Filesystem Migration using HP MirrorDisk/UX.

Check each volume group for the maximum number of physical volumes that can be allocated to the volume group.

```
#vgdisplay vg01

--- Volume groups ---
VG Name                /dev/vg01
VG Write Access        read/write
VG Status               available
Max LV                 255
Cur LV                8
Open LV                8
Max PV                 16
Cur PV                4
Act PV                4
Max PE per PV         2129
VGDA                   8
```

```
PE Size (Mbytes)      4
Total PE              8516
Alloc PE              7753
Free PE               763
Total PVG             0
Total Spare PVs      0
Total Spare PVs in use 0
```

As can be seen, the maximum number of physical volumes in the volume group is 16. This may not be enough if there are already more than eight physical volumes in the volume group. If this has to be increased, it will be necessary to export and import the volume group with a larger number of Max PVs.

To begin the migration, initialise the new disks as Logical Volume Manager Physical Volumes.

```
#pvcreate /dev/dsk/c14t7d5
#pvcreate /dev/dsk/c15t7d6
```

Add new disk into existing volume group as a physical volume group.

```
#vgextend -g PVG1 /dev/vg01
/dev/dsk/c14t7d5 /dev/dsk/c15t7d6
```

For each logical volume in the volume group, set a PVG-strict allocation policy. This means that mirrors of a logical extent cannot share the same physical volume group. This prevents the newly created mirrors from accidentally being created on the original disk.

```
#lvchange -s g /dev/vg01/lvol1
```

Extend the logical volume on to the new array by creating a new mirror in the new physical volume group.

```
#lvextend -m 1 /dev/vg01/lvol1 PVG1
```

When the time comes to cut over to the new array, for each logical volume simply remove the mirror copy on the DAS array.

```
#lvreduce -m 0 /dev/vg01/lvol1
/dev/dsk/c9t13d5 /dev/dsk/c9t13d6
```

Once all the mirrors have been removed, remove the physical volumes on the DAS array from the volume group.

```
#vgreduce vg01 /dev/dsk/c9t13d5
/dev/dsk/c9t13d6
```

The volume group is now fully operational on the SAN Array.

## Technique for Database Migration using IBM Informix Mirroring

Before IBM Informix Dynamic Server Mirroring can be used, an identical disk logical volume structure must be configured on the SAN array for each logical volume that is used by the IBM Informix instance on the DAS array.

Once all the logical volumes have been created, ensuring that the permissions are correct on the files in /dev, set up links to the new logical volumes.

```
$ cd /opt/Informix/dbspaces1
$ umask 117
$ ln -s /dev/vginformix/rrootdbs
  rootdbs_m
$ ln -s /dev/vginformix/rphysdbs
  physdbs_m
...
```

The next step is to mirror the database instance, chunk by chunk. It is advisable to create a shell script to do this. In addition the MIRROR flag in the onconfig configuration file should be set to "1",

```
$ cd $INFORMIXDIR/bin
$ ./onspaces -m rootdbs -p
/opt/Informix/dbspaces1/rootdbs -p \
/opt/Informix/dbspaces1/rootdbs_m 0
$ ./onspaces -m physdbs -p
/opt/Informix/dbspaces1/physdbs -p \
/opt/Informix/dbspaces1/physdbs_m 0
...
```

Verify the mirrors are active.

```
$ onstat -d

INFORMIX-Universal Server Version
9.14.UC4 -- On-Line -- Up 00:00:49 --
29384 Kbytes

Dbspaces

address number flags fchunk nchunks
flags owner name
```

```
b586108 1 a 1 1 MX informix rootdbs
1 active, 2047 maximum

Chunks

address chk/dbs offset size free bpages
flags pathname

b586180 1 1 0 25000 12400 PO-
/opt/informix/dbspaces1/rootdbs

b586258 1 1 0 25000 0 MX-
/opt/informix/dbspaces1/rootdbs_m

1 active, 2047 maximum
```

All chunks should have an "M" flag next to them.

When the time comes to cut over, simply shut down the engine and swap the links over.

```
$ cd /opt/Informix/dbspaces1
$ umask 117
$ mv rootdbs_m temp
$ mv rootdbs rootdbs_m
$ mv temp rootdbs
```

This will have the effect of swapping the primary link for the mirror. This should be done for all links.

Restart the instance and then remove the mirror on the DAS array.

```
$ cd $INFORMIXDIR/bin
$ ./onspaces -r rootdbs
$ ./onspaces -r physdbs
...
```

Now the instance is fully operational on the new SAN Array.

## Take another backup

Before removing the old DAS array, a full backup of all databases and filesystems should be taken.

## New SAN Configuration

Now that all the data from the DAS array have been migrated to the SAN Array, the DAS Array can be removed. All hosts should be shut down, all SCSI connections disconnected and terminated correctly.

A schematic of the new SAN can be seen below.

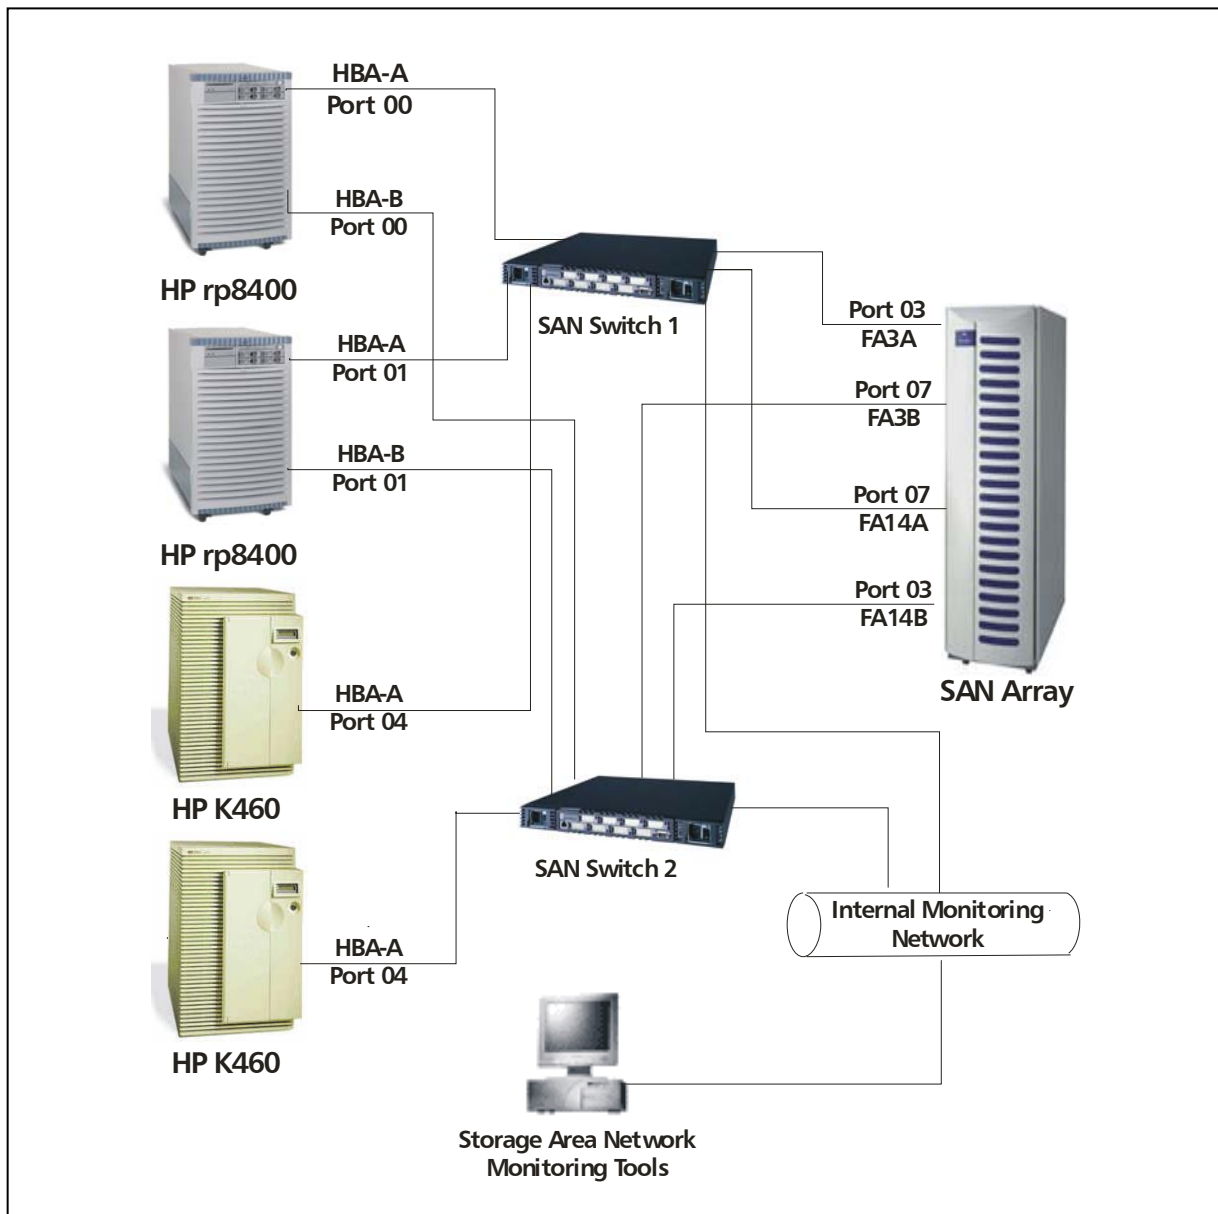


Figure 3 : SAN Configuration

---

## About the Author

After achieving a BSc (Hons) in Computer Science from the University of Newcastle, Scott pursued a career as a technical IT consultant. Scott has had training from companies like Hewlett-Packard, Informix, and EMC and is an Oracle and Informix database administrator (DBA).

Scott has worked for a number of organisations including DHL, British Telecom, Parallel IT Consulting, MacDermid Europe, and BSO/Origin BV.

Scott planned and executed the migration to a new European data centre for DHL Worldwide Express from servers based in individual countries. This included implementing storage solutions, building a complete migration environment and designing a backup and recovery solution.

## About ArdentA

ArdentA is an independent IT consultancy specialising in managed technical services, remote technical support and project management to organisations running database management systems on a range of computing platforms.

Based in Sunbury on Thames, the company's mission is to help enterprises to minimise their total cost of ownership by facilitating optimal use of their technologies.

## Glossary

### Direct access storage device (DASD)

A mass storage medium, on which a computer stores data, that permits direct retrieval of information as opposed to sequential retrieval.

### Fibre Channel (FC)

A high-speed networking standard for disks whose underlying medium is fibre optics.

### Host Bus Adapter (HBA)

An HBA is an I/O adapter that sits between the host computer's bus and the Fibre Channel loop and manages the transfer of information between the two channels. In order to minimize the impact on host processor performance, the host bus adapter performs many low-level interface functions automatically or with minimal processor involvement.

### Informix Dynamic Server (IDS)

Informix Dynamic Server is a multithreaded relational database server that manages data that is stored in rows and columns. It employs a single processor or symmetric multiprocessor (SMP) systems and dynamic scalable architecture (DSA) to deliver database scalability, manageability and performance.

### Logical Volume Manager (LVM)

Allows logical volumes to span multiple physical volumes. Data on logical volumes appears to be contiguous to the user, but can be discontinuous on the physical volume. This allows file systems, paging space, and other logical volumes to be resized or relocated, span multiple physical volumes, and have their contents replicated for greater flexibility and availability.

### Network Data Management Protocol (NDMP)

An open protocol for network-based backup of data.

### Storage Area Network (SAN)

A network where a small number of computers share a large amount of data, usually within a single server room, where performance is critical. Examples of SAN protocols include Fiber Channel and the Small Computer Systems Interface (SCSI).

### Small Computer Systems Interface (SCSI)

A parallel interface standard used by Apple Macintosh computers, PCs, and many Unix systems for attaching peripheral devices to computers. SCSI interfaces provide for faster data transmission rates (up to 80 MBps) than standard serial and parallel ports. In addition, you can attach many devices to a single SCSI port, so that SCSI is really an I/O bus rather than simply an interface.

### Switch

A Fabric device providing bandwidth and high-speed routing of data via link-level addressing.

### Total cost of ownership (TCO)

A model developed by Gartner Group to analyse the direct and indirect costs of owning and using hardware and software. Managers of enterprise systems use various versions of TCO to lower costs while increasing the benefits of information technology deployments.

### UNIX operating system

An operating system developed by Bell Laboratories that features multiprogramming in a multiuser environment. UNIX was originally developed for use on minicomputers but has been adapted for mainframes and microcomputers.

---

## Trademarks

Brocade® is a registered trademark of Brocade Communications Systems, Inc.

Hewlett-Packard®, HP®, and HP-UX® are registered trademarks of Hewlett-Packard Company.

IBM UNIX is a registered trademark in the United States, other countries, or both and is licensed exclusively through X/Open Company Limited.

Informix® is a trademark of the IBM Corporation in the United States and/or other countries:

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other trademarks and trade names mentioned herein are the property of their respective owners.

## Disclaimer

ArdentA Limited believes that the information in this publication is accurate as of its publication date. The information is subject to change without notice.

ArdentA Limited makes no representations or warranties of any kind with respect to the information in this publication and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.