# Migrating an SAP environment from legacy storage to IBM XIV Storage System


This paper can be found in the IBM Techdocs library.

*Hans-Günther Hörhammer (hoerhammer@de.ibm.com)*

*Axel Westphal (awestphal@de.ibm.com)*

*IBM Systems and Technology Group ISV Enablement*
*IBM System Storage - Advanced Technical Support (ATS)*
*European Storage Competence Center (ESCC), Mainz, Germany*

*November 2010*

© Copyright IBM Corporation, 2010

## Table of contents




© Copyright IBM Corporation, 2010

## Abstract

*There are many options available for migrating data from one storage system to another. This white paper describes two methods of data migration from legacy storage to IBM XIV Storage System in an SAP environment. The first method describes a procedure for data migration in conjunction with a volume layout change such as logical unit number (LUN) consolidation. A reason for LUN consolidation can be that it might help to gain more advantage of the XIV Storage System architecture. The second method describes the embedded XIV migration feature that enables the easy movement of data from an existing storage system to the XIV Storage System.*


## Introduction

Whenever storage hardware has to be replaced, the migration of the data from the old storage system to the new storage system has to be planned and performed. There are various tools available for data migration. The question, which tool or method is the appropriate one depends on various aspects. One of those aspects is the differences between the existing and the new volume layout. The traditional approach for a volume layout on a storage system is the use of multiple small volumes distributed all over the arrays, ranks, pools, and adapters in the storage backend. The intention is to use every physical disk drive and distribute the workload equally.

This leads to the result that on various storage systems the volume layout consists of multiple small volumes. As the IBM® XIV® Storage System has a different internal architecture; the volume layout is also different. Every volume you create on the XIV system will be automatically distributed over all physical disks. Therefore, a small amount of big volumes is the desired volume layout. Changing the volume layout from multiple small volumes to a few big volumes requires a volume consolidation. This white paper describes how to consolidate the volumes of the legacy storage system in order to meet the new volume layout on the XIV system.

In case a change of the volume layout is not necessary, the XIV system provides a data migration feature that allows data migration in a one-to-one volume relationship. This feature enables the production environment to continue functioning during the data transfer with only one brief period of downtime for your business applications. The XIV system Data Migration service does not need a complex configuration or additional software. A small downtime of the server is required to redirect the host and the SAP system from legacy storage to XIV Storage System. As soon as the SAP volumes are redirected to the XIV system, the SAP system can be started, while the migration is running in the background. The XIV system migration feature requires maintaining the original data layout and volume sizes. Both described migration procedures in this whitepaper also work for non-SAP applications on the host.


© Copyright IBM Corporation, 2010

# IBM XIV Storage System Product overview

IBM XIV Storage System is a next-generation, high-end, open disk-storage system. As part of the IBM broad spectrum of system-storage and storage area network (SAN) offerings, IBM XIV is an innovative grid-based storage system, using off-the-shelf hardware components. The XIV system architecture is designed to provide the highest levels of performance, reliability, and functions — combined with unprecedented ease of management and exceptionally low total cost of ownership (TCO).

The IBM XIV Storage System provides the following features:

- A revolutionary, high-end, disk storage system designed to eliminate the complexity of administration and management of tiered storage
- Fully-virtualized grid architecture, supporting mission-critical workloads and innovatively designed to simplify storage, virtually in every way
- High reliability and data availability through active-active N+1 redundancy of all key components, partition mirroring, unique self-healing, rapid rebuild times, and nondisruptive upgrades
- Unmatched total cost of ownership (TCO) and environment-friendly attributes, with breakthrough ease of management, full system virtualization, and dramatic efficiencies in capacity, energy, and space
- Enterprise-level consistent performance through massive parallelism, disk utilization, and unique caching mechanisms.
- Open and integrated system, offering strong integration with the IBM portfolio and diverse interoperability with open-system servers and technologies
- Transparent, no-impact scalability, in on-demand increments, without the need for manual data migration or performance tuning, providing automatic online capacity growth
- Rich, enterprise-class function set bundled with the system software including virtually unlimited, near-instant differential snapshots, thin provisioning, data migration, asynchronous and synchronous mirroring, as well as an intuitive management console
- Data migration can be used to migrate volumes from legacy storage systems to the XIV Storage System with practically no downtime


© Copyright IBM Corporation, 2010

You can find more details about the XIV system features and benefits at the following websites:

http://www.redbooks.ibm.com/abstracts/sg247659.html?Open
http://www-03.ibm.com/systems/storage/disk/xiv/index.html

# Hardware and software setup (in the Mainz lab)

The following sections describe the hardware and the software setup requirements for the test environment.

## Diagram of the test environment

The hardware and software setup used for the migration tests explained in this white paper is as shown in Figure 1.

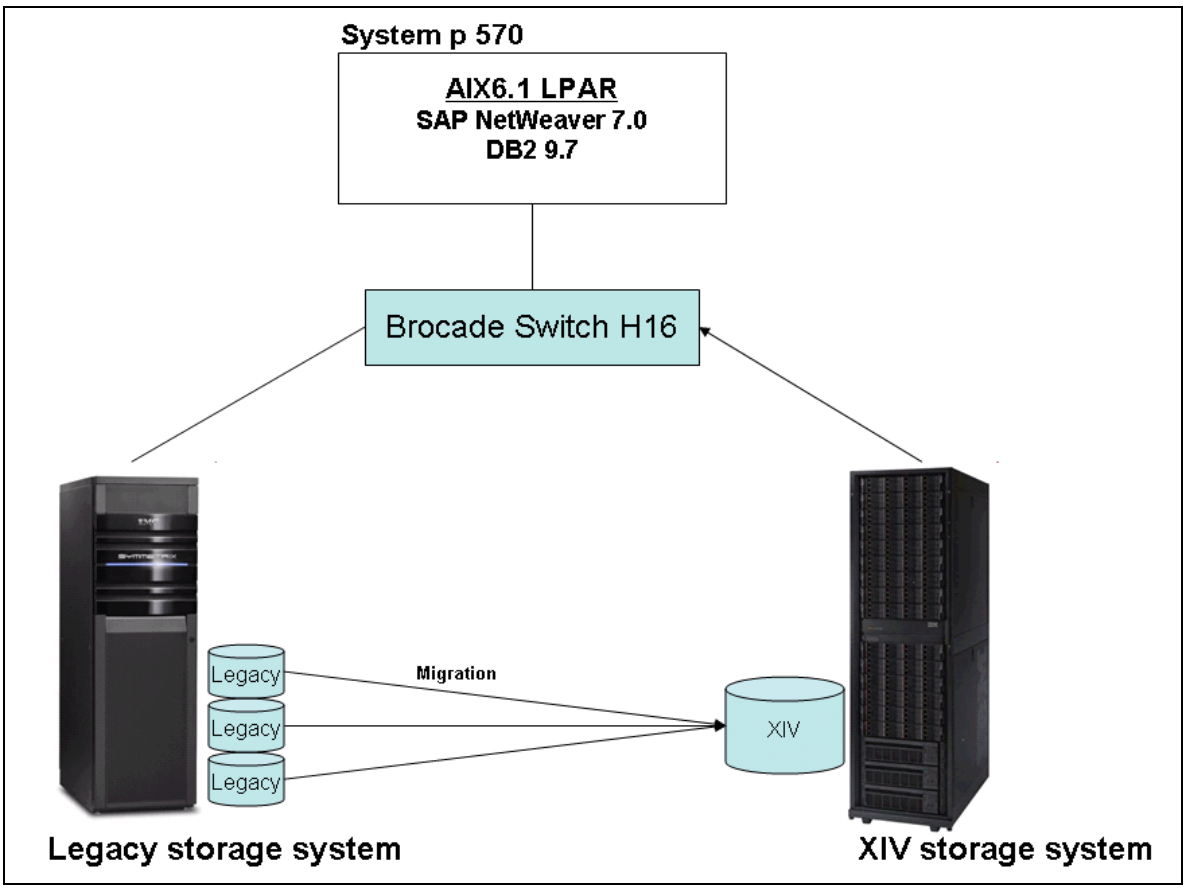


*Figure 1: Hardware setup for the XIV system migration scenario*

## SAP central instance server

The information about the hardware and software versions that were used on the SAP server is shown in Table 1. One logical partition (LPAR) on an IBM Power® 570 system with one processor was configured as the SAP application server. This LPAR was used as the production server. In the Logical Volume


© Copyright IBM Corporation, 2010

Manager (LVM) migration scenario, this LPAR was also used for migrating data from the legacy storage to the XIV Storage System.

Table 1 summarizes the hardware and software configuration for the IBM AIX® server in the test environment.

| **Feature** | **AIX 6.1 / SAP central instance** |
|---|---|
| Server model | IBM System Power 570 |
| Processors | One |
| RAM | 8 GB |
| Boot disk | 50 GB |
| Operating system version | AIX 6.1<br>6100-04-03-1009 |
| SAP application | NetWeaver 7.0 and SAP Retail System ERP 6.0 |
| IBM DB2® version | DB2 9.7 |
| Multipath driver | Powerpath 5.3 for AIX |

*Table 1: SAP central instance server*

Note: AIX does not support more than one multi-pathing driver stack per Host Bus Adapter. To use MPIO and powerpath simultaneously there is the requirement to use separate Host Bus Adapters for each storage system.

## Storage hardware

Two storage systems were used to test the migration scenarios. For the source storage system, a legacy storage system was used. The XIV Storage System represents the target storage system in the tests.

Table 2 shows the hardware and software levels of the XIV storage system

| **Feature** | **Version** |
|---|---|
| Storage model 1 | IBM XIV Storage System 2810 |
| Storage manager server level | 10.2.2 |
| Storage manager GUI | 2.4.3 |

*Table 2: XIV Storage System hardware and software*


© Copyright IBM Corporation, 2010

## Volume and volume group layout

For the test scenario, the SAP system was installed in three volume groups residing on multiple small volumes from the legacy storage. The names of the volume groups are *sapdatavg*, *saplogvg* and *sapswvg*. Each test started with SAP running on these volume groups. This configuration of the legacy storage volumes represented the multiple small volumes configuration. Depending on the migration scenario, the volumes were copied by the XIV Storage System migration feature or were consolidated using the LVM. In the LVM scenario, the volume layout is changed in order to obtain only a few volumes on the XIV Storage System. As the steps for migration of each volume group is identical, only the migration of the *sapdatavg* is shown as example in the LVM migration scenario.

| SAP LPAR | Volume group | Volume IDs | Volume size |
|---|---|---|---|
| SID: MTX | sapdatavg | hdiskpower0<br>hdiskpower1 | 100 GB |
| | saplogvg | Hdiskpower2<br>Hdiskpower3 | 100 GB |
| | sapswvg | hdiskpower4<br>hdiskpower5 | 100 GB |

*Table 3: SAP volume layout before migration*

# Configuration of AIX in an XIV system environment

This section explains specific considerations for host connectivity and describes the host attachment-related tasks for the AIX operating system platform. At the time of writing the Host Attachment Guide for AIX Version 1.5.2 was used. For the latest information, refer to the hosts attachment kit publications at

http://www-01.ibm.com/support/docview.wss?uid=ssg1S4000802

## Prerequisites

If the current AIX operating system level installed on your system is not a level that is compatible with XIV system, you must upgrade the operating system prior to attaching it to the XIV Storage System. To determine the maintenance package or technology level installed on your system, use the oslevel command as shown in Example 1: Determine current AIX version and maintenance level.

```
root@sap-xiv-1: / > oslevel -s
```


© Copyright IBM Corporation, 2010

```
6100-04-03-1009
```

*Example 1: Determine current AIX version and maintenance level*

In this test scenario, the system is running AIX 6.1.0.0 technology level 4. Use this information in conjunction with the System Storage Interoperation Center (SSIC) to ensure that the attachment will be an IBM-supported configuration.

In the event that AIX maintenance items are needed, consult the IBM Fix Central Web site to download fixes and updates for your system's software, hardware, and operating system at:

http://www-933.ibm.com/support/fixcentral/main/System+p/AIX

Attaching the XIV Storage System to an AIX host using Fibre Channel (FC) involves the following activities from the host side:

- Identify the Fibre Channel Host Bus Adapters (HBAs) and determine their Worldwide Port Name (WWPN) values.
- Install AIX Host Attachment Kit that is specific to XIV system.
- Configure multipathing (optional)

In order to allocate XIV system volumes to an AIX host, the first step is to identify the Fibre Channel adapters on the AIX server. Use the *lsdev* command to list all the FC adapter ports in your system, as shown in Example 2. In the test environment the AIX LPAR is connected to a VIO server over NPIV. Therefore the output shows virtual fibre channel adapters.

```
root@sap-xiv-1: / > lsdev -Cc adapter | grep fcs

fcs0   Available 12-T1 Virtual Fibre Channel Client Adapter

fcs1   Available 13-T1 Virtual Fibre Channel Client Adapter

fcs2   Available 14-T1 Virtual Fibre Channel Client Adapter

fcs3   Available 15-T1 Virtual Fibre Channel Client Adapter
```

*Example 2: Listing Fibre Channel adapters*

To obtain the WWPN of each of the IBM POWER® system FC adapters, you can use the ***lscfg*** command, as shown in Example 3.

```
root@sap-xiv-1: / > lscfg -vl fcs0  fcs0

     U9117.MMA.06B33F2-V40-C12-T1  Virtual Fibre Channel Client Adapter

        Network Address.............C0507600C4220072

        ROS Level and ID............
```


© Copyright IBM Corporation, 2010

*Example 3: WWPN of the Fibre Channel adapter*

At this point, you can define the AIX host system on the XIV Storage System and assign the WWPN to the host. If the FC connection was correctly created, the zoning enabled, and the Fibre Channel adapters are in an available state on the host, these ports will be selectable from the drop-down list in the XIV system graphical user interface. After creating the AIX host, map the XIV system volumes to the host.

### Installing the XIV Host Attachment Kit for AIX

For AIX to correctly recognize the disks mapped from the XIV Storage System as MPIO 2810 XIV Disk, a specific software package known as the XIV Host Attachment Kit for AIX (XIV HAK) is required on the AIX system. This package will also enable multipathing. At the time of this writing the XIV HAK 1.5.2 was used. You can download the fileset from:

http://www-01.ibm.com/support/search.wss?q=ssg1*&tc=STJTAG+HW3E0&rs=1319&dc=D400&dtm

For more details, refer to the XIV Host Attachment Guide for AIX at:

http://www-01.ibm.com/support/docview.wss?uid=ssg1S4000802

## Migrating to XIV system from legacy storage

This section describes two methods for data migration of an SAP system from a legacy storage system to an XIV Storage System. The first section explains how to perform a migration combined with LUN consolidation. This is the preferred method if the original volume layout does not meet the desired volume layout on the XIV Storage System. The second section shows how to use the XIV Storage System feature for migration if LUN consolidation is not necessary.

### Migration with LUN consolidation

As mentioned in the introduction, on many storage systems it is for performance reasons that it is recommended to have a volume layout, which is based on multiple small volumes. The XIV system architecture is designed in a way that only a few volumes are needed. This simplifies the disk management not only on the storage system, but also on the AIX host. The following section shows how to use the LVM to migrate and consolidate the LUNs from the legacy storage system to the XIV Storage System using the AIX LVM *mklvcopy* and *syncvg* commands The AIX LVM is very flexible. Under certain constraints or requirements for data spreading or consolidation across LUNs, it may be easier to use the *mklvcopy* and *syncvg* commands to accomplish the migration from legacy storage to XIV Storage System. If you choose the *mklvcopy* command, it is only important to check that there is enough space in the volume groups and on the XIV system LUNs to create the mirrors.

The following steps describe how to migrate and consolidate the volumes form the legacy storage to the XIV Storage System using the *mklvcopy* and *syncvg* commands. As the steps for migration of the volume groups are identical, only the migration of one of the volume group (*sapdatavg*) is shown as example in the LVM migration scenario:


© Copyright IBM Corporation, 2010

1. Identify the source volumes of the legacy storage system. Use the *lsdev* command and the *bootinfo –s* command (as shown in Example 3 and Example 4) to gather information about the AIX host such as the number of LUNs and sizes.

```
root@sap-xiv-1: / > lsdev -c disk

hdisk0      Available          Virtual SCSI Disk Drive

hdisk1      Available 13-T1-01 MPIO 2810 XIV Disk

hdisk2      Available 12-T1-01 EMC Symmetrix FCP VRAID

hdisk3      Available 12-T1-01 EMC Symmetrix FCP VRAID

hdisk4      Available 12-T1-01 EMC Symmetrix FCP VRAID

hdisk5      Available 12-T1-01 EMC Symmetrix FCP VRAID

hdisk6      Available 12-T1-01 EMC Symmetrix FCP VRAID

hdisk7      Available 12-T1-01 EMC Symmetrix FCP VRAID

hdisk8      Available 12-T1-01 EMC Symmetrix FCP VRAID

hdisk9      Available 14-T1-01 EMC Symmetrix FCP VRAID

hdisk10     Available 14-T1-01 EMC Symmetrix FCP VRAID

hdisk11     Available 14-T1-01 EMC Symmetrix FCP VRAID

hdisk12     Available 14-T1-01 EMC Symmetrix FCP VRAID

hdisk13     Available 14-T1-01 EMC Symmetrix FCP VRAID

hdisk14     Available 14-T1-01 EMC Symmetrix FCP VRAID

hdisk15     Available 14-T1-01 EMC Symmetrix FCP VRAID

hdiskpower0 Available 14-T1-01 PowerPath Device

hdiskpower1 Available 14-T1-01 PowerPath Device

hdiskpower2 Available 14-T1-01 PowerPath Device

hdiskpower3 Available 14-T1-01 PowerPath Device

hdiskpower4 Available 14-T1-01 PowerPath Device

hdiskpower5 Available 14-T1-01 PowerPath Device

hdiskpower6 Available 14-T1-01 PowerPath Device
```

*Example 3: Output of the lsdev command*


© Copyright IBM Corporation, 2010

```
root@sap-xiv-1: / > bootinfo -s hdiskpower0

105000

root@sap-xiv-1: / > bootinfo -s hdiskpower1

105000
```

*Example 4: Output of the bootinfo –s command*

2. Use the *lsvg –p* command (as shown in Example 5:) to collect information about the volume group which will be migrated. In this example, the volume group name is *sapdatavg.*

```
root@sap-xiv-1: / > lsvg -p sapdatavg

sapdatavg:

PV_NAME           PV STATE          TOTAL PPs   FREE PPs    FREE DISTRIBUTION

hdiskpower0       active            410         4           00..00..00..00..04

hdiskpower1       active            410         5           00..00..00..00..05
```

*Example 5: Output of the lsvg –p command*

3. Collect information about the logical volumes in the volume group. Use the *lsvg –l vg_name* command, (as shown in Example 6) where *vg_name* is the name of the volume group.

```
root@sap-xiv-1: / > lsvg -l sapdatavg

sapdatavg:

LV NAME           TYPE         LPs     PPs    PVs    LV STATE      MOUNT POINT

sapdata1lv          jfs2       200     200     1    open/syncd  /db2/MTX/sapdata1

sapdata2lv          jfs2       200     200     1    open/syncd  /db2/MTX/sapdata2

sapdata3lv          jfs2       200     200     1    open/syncd  /db2/MTX/sapdata3

sapdata4lv          jfs2       200     200     1    open/syncd  /db2/MTX/sapdata4

siddb2sidlv         jfs2       10      10      1    open/syncd  /db2/MTX/db2mtx

loglv02             jfs2log    1       1       1    open/syncd    N/A
```

*Example 6 : Output of the lsvg –l vg_name command*

4. Assign the target LUNs on XIV Storage System to the host.


© Copyright IBM Corporation, 2010

5. Discover the XIV system LUNs using the *cfgmgr* command (as shown in 7) on the AIX host.

```
root@sap-xiv-1: / > cfgmgr

root@sap-xiv-1: / >
```

*Example 7: Output of the cfgmgr command*

6. Determine the difference between the source LUNs and XIV system LUNs. Again run the *lsdev -c disk* command (as shown in Example 8: Output of the *lsdev -c disk* command.

```
root@sap-xiv-1: / > lsdev -c disk

hdisk0      Available          Virtual SCSI Disk Drive

hdisk1      Available 13-T1-01 MPIO 2810 XIV Disk

hdisk2      Available 12-T1-01 EMC Symmetrix FCP VRAID

hdisk3      Available 12-T1-01 EMC Symmetrix FCP VRAID

hdisk4      Available 12-T1-01 EMC Symmetrix FCP VRAID

hdisk5      Available 12-T1-01 EMC Symmetrix FCP VRAID

hdisk6      Available 12-T1-01 EMC Symmetrix FCP VRAID

hdisk7      Available 12-T1-01 EMC Symmetrix FCP VRAID

hdisk8      Available 12-T1-01 EMC Symmetrix FCP VRAID

hdisk9      Available 14-T1-01 EMC Symmetrix FCP VRAID

hdisk10     Available 14-T1-01 EMC Symmetrix FCP VRAID

hdisk11     Available 14-T1-01 EMC Symmetrix FCP VRAID

hdisk12     Available 14-T1-01 EMC Symmetrix FCP VRAID

hdisk13     Available 14-T1-01 EMC Symmetrix FCP VRAID

hdisk14     Available 14-T1-01 EMC Symmetrix FCP VRAID

hdisk15     Available 14-T1-01 EMC Symmetrix FCP VRAID

hdisk16      Available 13-T1-01 MPIO 2810 XIV Disk

hdiskpower0 Available 14-T1-01 PowerPath Device

hdiskpower1 Available 14-T1-01 PowerPath Device

hdiskpower2 Available 14-T1-01 PowerPath Device

hdiskpower3 Available 14-T1-01 PowerPath Device
```


© Copyright IBM Corporation, 2010

```
hdiskpower4 Available 14-T1-01 PowerPath Device

hdiskpower5 Available 14-T1-01 PowerPath Device

hdiskpower6 Available 14-T1-01 PowerPath Device
```

*Example 8: Output of the* lsdev -c disk *command*

7. Identify the sizes of the XIV Storage System target LUNs. Use the *bootinfo -s hdisk#* command (as shown in Example 9: Output of the bootinfos –s command), where # is the number of the hdisk.

```
root@sap-xiv-1: / > bootinfo -s hdisk16

212992
```

*Example 9: Output of the bootinfos –s command*

Before the team continued with the migration and LUN consolidation procedure, the team used the SAP load generator command SGEN to introduce some load on the SAP system. The SAP transaction SGEN is used to generate the ABAP loads of a number of programs, function groups, classes, and so on, as well as Business Server Page applications.

This shows that the LUN consolidation procedure will work with load from the server. The *iostat* command (as shown in Example 10: Output of the iostat command) shows that now *sgen* is running without having started the logical volume copy:

```
root@sap-xiv-1: / > iostat -d hdiskpower0 hdiskpower1 hdiskpower2 hdiskpower3
hdiskpower4 hdiskpower5 hdisk16 3 10

hdiskpower0     89.0     7680.0     1136.7      22264       776

hdiskpower1     27.0     1796.0     338.7       4188      1200

hdiskpower2      0.0       4.0       1.0          0        12

hdiskpower3      3.3     549.3      98.0          0      1648

hdiskpower4      0.3     134.7      13.0          0       404

hdiskpower5      0.3      12.0       3.0          4        32

hdisk16          0.0       0.0       0.0          0         0
```

*Example 10: Output of the iostat command*


© Copyright IBM Corporation, 2010

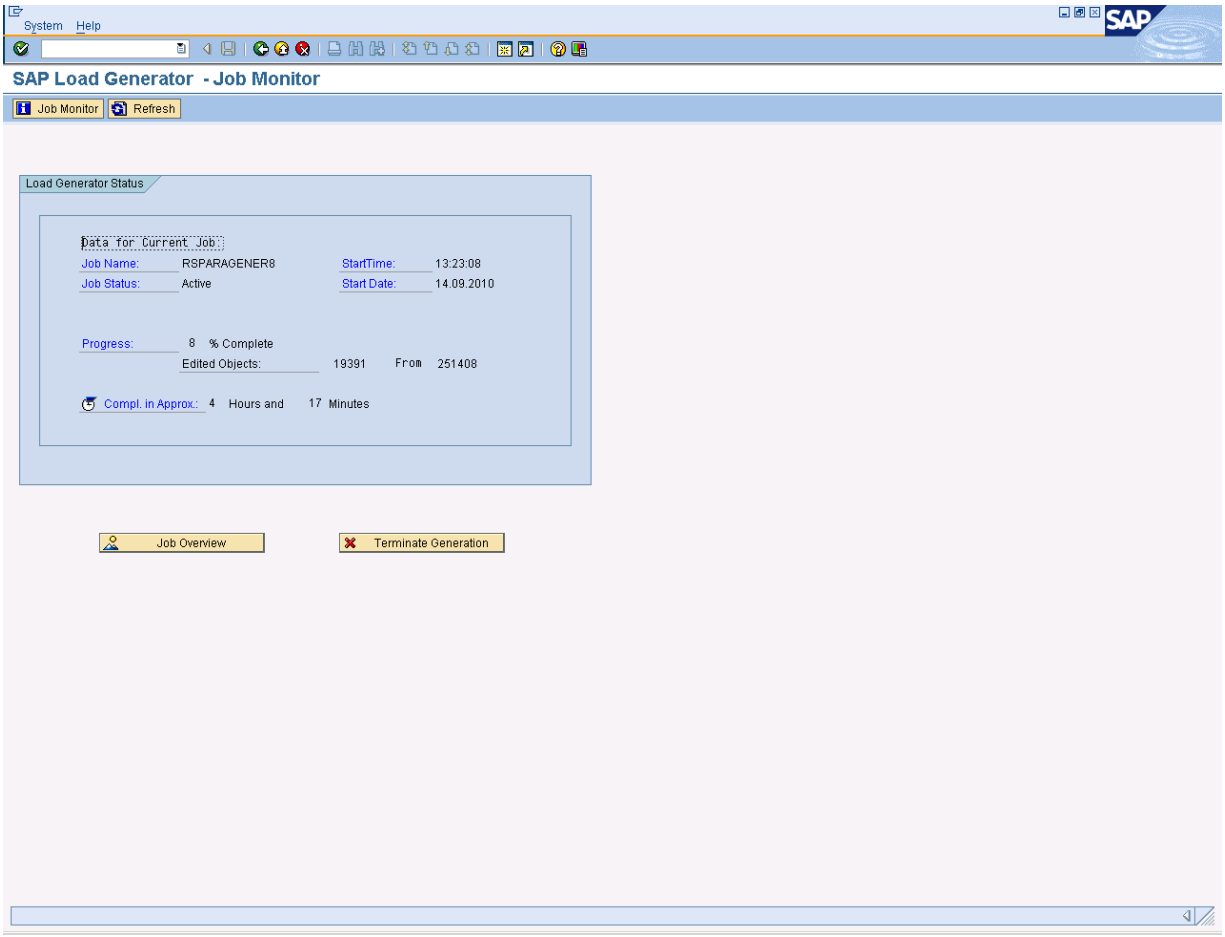


*Figure 2: Status of the SAP load generator job*

8. Move the XIV system LUNs into the volume groups appropriately. Use the *extendvg vg-name hdisk#* command where *vg_name* is the name of the volume group and # is the number of the hdisk. In this example, *hdisk16* has to be moved into *sapdatavg (as shown in* Example 11: extendvg command*)*.

```
root@sap-xiv-1: / > extendvg sapdatavg hdisk16
```

*Example 11: extendvg command*


© Copyright IBM Corporation, 2010

9. Verify that the XIV system LUNs are added to the volume group using the *lsvg -p vg_name* command (as shown in Example 12: Output of the lsvg –p vg_name command), where *vg_name* is the name of the volume group.

```
root@sap-xiv-1: / > lsvg -p sapdatavg

sapdatavg:

PV_NAME           PV STATE     TOTAL PPs   FREE PPs    FREE DISTRIBUTION

hdiskpower0       active         410         4           00..00..00..00..04

hdiskpower1       active         410         5           00..00..00..00..05

hdisk16           active         831         831         167..166..166..166..166
```

*Example 12: Output of the lsvg –p vg_name command*

10. Determine how the logical volumes are spread across the hdisks. Use the *lv_name* command (as shown in Example 13: Output of the lsvg –l lv_name command) where *lv_name* is the name of the logical volume.

```
root@sap-xiv-1: / > lslv -l sapdata1lv

sapdata1lv:/db2/MTX/sapdata1

PV                COPIES        IN BAND       DISTRIBUTION

hdiskpower0       100:000:000   82%           000:082:018:000:000

hdiskpower1       100:000:000   82%           000:082:018:000:000

root@sap-xiv-1: / > lslv -l sapdata2lv

sapdata2lv:/db2/MTX/sapdata2

PV                COPIES        IN BAND       DISTRIBUTION

hdiskpower0       100:000:000   0%            036:000:064:000:000

hdiskpower1       100:000:000   0%            036:000:064:000:000

root@sap-xiv-1: / > lslv -l sapdata3lv

sapdata3lv:/db2/MTX/sapdata3

PV                COPIES        IN BAND       DISTRIBUTION

hdiskpower0       100:000:000   0%            046:000:000:054:000

hdiskpower1       100:000:000   0%            046:000:000:054:000

root@sap-xiv-1: / > lslv -l sapdata4lv

sapdata4lv:/db2/MTX/sapdata4
```


© Copyright IBM Corporation, 2010

```
PV                COPIES          IN BAND       DISTRIBUTION

hdiskpower0       100:000:000     0%            000:000:000:028:072

hdiskpower1       100:000:000     0%            000:000:000:028:072
```

*Example 13: Output of the lsvg –l lv_name command*

The server output shows that the logical volumes reside on *hdiskpower0* and *hdiskpower1*, which are both volumes from the legacy storage.

11. Copy the logical volume data from the source LUNs to the XIV system target LUN. Use the *mklvcopy lv_name 2 hdisk#* command (as shown in Example 14: mklvcopy command), where *lv_name* is the name of the logical volume and # is the number of the *hdisk* on the target XIV system LUN.

```
root@sap-xiv-1: / > mklvcopy sapdata1lv 2 hdisk16

root@sap-xiv-1: / > mklvcopy sapdata2lv 2 hdisk16

root@sap-xiv-1: / > mklvcopy sapdata3lv 2 hdisk16

root@sap-xiv-1: / > mklvcopy sapdata4lv 2 hdisk16

root@sap-xiv-1: / > mklvcopy siddb2sidlv 2 hdisk16

root@sap-xiv-1: / > mklvcopy loglv02 2 hdisk16
```

*Example 14: mklvcopy command*

12. Verify that the logical volume copies are correct using the *lslv -l lv_name* command where *lv_name* is the name of the logical volume.

The iostat command (as shown in Example 15: Output of the iostat command) after mklvcopy has started shows that the data is now copied to hdisk16 which resides in the XIV Storage System.

```
Disks:        % tm_act     Kbps      tps    Kb_read   Kb_wrtn

hdiskpower0     14.3     2281.0     481.3       5988       855

hdiskpower1     13.3     1937.0     353.3       4676      1135

hdiskpower2      1.3     529.3      29.7           0      1588

hdiskpower3      0.0       0.0       0.0           0         0

hdiskpower4      0.0       0.0       0.0           0         0

hdiskpower5      0.0       0.0       0.0           0         0

hdisk16          1.3       5.3      10.7           0        16
```


© Copyright IBM Corporation, 2010

*Example 15: Output of the iostat command*

13. Synchronize the logical volume data from the legacy storage LUNs to the XIV system target LUNs using the command,*syncvg -v vg_name (as shown in* Example 16: syncvg command*)*, where *vg_name* is the name of the volume group. In this test, the name of the volume group is *sapdatavg.*

```
root@sap-xiv-1: / > syncvg -v sapdatavg
```

*Example 16: syncvg command*

14. Verify that the synchronization is not showing stale, it should show as sync'd. Use *lsvg -l vg_name* command (as shown in Example 17: Output of the lsvg vg_name command), and if the lv still shows stale, then you need to resynchronize it before proceeding.

```
root@sap-xiv-1: / > lsvg sapdatavg

VOLUME GROUP:       sapdatavg                VG IDENTIFIER:
00cb33f200004c000000012b0b7705bf

VG STATE:           active                   PP SIZE:        256 megabyte(s)

VG PERMISSION:      read/write               TOTAL PPs:      1651 (422656
megabytes)

MAX LVs:            256                      FREE PPs:       29 (7424 megabytes)

LVs:                6                        USED PPs:       1622 (415232
megabytes)

OPEN LVs:           6                        QUORUM:         2 (Enabled)

TOTAL PVs:          3                        VG DESCRIPTORS: 3

STALE PVs:          0                        STALE PPs:      0

ACTIVE PVs:         3                        AUTO ON:        yes

MAX PPs per VG:     32512

MAX PPs per PV:     1016                     MAX PVs:        32

LTG size (Dynamic): 256 kilobyte(s)          AUTO SYNC:      no

HOT SPARE:          no                       BB POLICY:      relocatable
```

*Example 17: Output of the lsvg vg_name command*

The server output on sap-xiv-1 show zero stale partitions. That means that all copies of the logical volumes are in synchronization.


© Copyright IBM Corporation, 2010

15. Verify the source and target LUNs for each logical volume using the *lslv -l lv_name* command (as shown in Example 18: Output of the lslv –l lv_name commands), where *lv_name* is the name of the logical volume.

```
root@sap-xiv-1: / > lslv -l sapdata1lv

sapdata1lv:/db2/MTX/sapdata1lv

PV                COPIES        IN BAND       DISTRIBUTION

hdiskpower0       100:000:000   82%           000:082:018:000:000

hdisk16           200:000:000   83%           000:166:034:000:000

hdiskpower1       100:000:000   82%           000:082:018:000:000

root@sap-xiv-1: / > lslv -l sapdata2lv

sapdata2lv:/db2/MTX/sapdata2

PV                COPIES        IN BAND       DISTRIBUTION

hdiskpower0       100:000:000   0%            036:000:064:000:000

hdisk16           200:000:000   0%            068:000:132:000:000

hdiskpower1       100:000:000   0%            036:000:064:000:000

root@sap-xiv-1: / > lslv -l sapdata3lv

sapdata3lv:/db2/MTX/sapdata3

PV                COPIES        IN BAND       DISTRIBUTION

hdiskpower0       100:000:000   0%            046:000:000:054:000

hdisk16           200:000:000   0%            099:000:000:101:000

hdiskpower1       100:000:000   0%            046:000:000:054:000

root@sap-xiv-1: / > lslv -l sapdata4lv

sapdata4lv:/db2/MTX/sapdata4

PV                COPIES        IN BAND       DISTRIBUTION

hdiskpower0       100:000:000   0%            000:000:000:028:072

hdisk16           200:000:000   0%            000:000:000:065:135

hdiskpower1       100:000:000   0%            000:000:000:028:072

root@sap-xiv-1: / > lslv -l siddb2sidlv

siddb2sidlv:/db2/MTX/db2mtx
```


© Copyright IBM Corporation, 2010

```
PV                COPIES        IN BAND       DISTRIBUTION

hdiskpower0       005:000:000   0%            000:000:000:000:005

hdisk16           010:000:000   0%            000:000:000:000:010

hdiskpower1       005:000:000   0%            000:000:000:000:005

root@sap-xiv-1: / > lslv -l loglv02

loglv02:N/A

PV                COPIES        IN BAND       DISTRIBUTION

hdiskpower0       001:000:000   0%            000:000:000:000:001

hdisk16           001:000:000   0%            000:000:000:000:001
```

*Example 18: Output of the lslv –l lv_name commands*

16. Remove the source copy of the logical volume from the legacy storage LUNs. The *rmlvcopy lv_name 1 hdisk#* command (as shown in Example 19: rmlvcopy command) removes the logical volume where *lv_name* is the name of the logical volume and # is the number of the hdisk of the source LUN.

```
root@sap-xiv-1: / > rmlvcopy sapdata1lv 1 hdiskpower0 hdiskpower1

root@sap-xiv-1: / > rmlvcopy sapdata2lv 1 hdiskpower0 hdiskpower1

root@sap-xiv-1: / > rmlvcopy sapdata3lv 1 hdiskpower0 hdiskpower1

root@sap-xiv-1: / > rmlvcopy sapdata4lv 1 hdiskpower0 hdiskpower1

root@sap-xiv-1: / > rmlvcopy siddb2sidlv 1 hdiskpower0 hdiskpower1

root@sap-xiv-1: / > rmlvcopy loglv02 1 hdiskpower0
```

*Example 19: rmlvcopy command*

17. Verify that all the source LUNs on the legacy storage are free with no data using the *lsvg -p vg_name* command (as shown in Example 20) where *vg_name* is the name of the volume group.

```
root@sap-xiv-1: / > lsvg -p sapdatavgsapdatavg:

PV_NAME           PV STATE          TOTAL PPs   FREE PPs    FREE DISTRIBUTION

hdiskpower0       active            410         410         82..82..82..82..82

hdiskpower1       active            410         410         82..82..82..82..82

hdisk16           active            831         20          00..00..00..00..20
```

*Example 20: Output of the lsvg –p vg_name command*


© Copyright IBM Corporation, 2010

18. Use the iostat command (as shown in Example 21: Output of the iostat command) and verify that no more I/O is going to the disks from the legacy storage.

```
Disks:        % tm_act     Kbps      tps       Kb_read   Kb_wrtn

hdiskpower0      0.0       0.0       0.0          0         0

hdiskpower1      0.0       0.0       0.0          0         0

hdiskpower2      0.0       0.0       0.0          0         0

hdiskpower3      0.0       0.0       0.0          0         0

hdiskpower4      0.0       0.0       0.0          0         0

hdiskpower5      0.0       0.0       0.0          0         0

hdisk16         81.0     5797.3     933.3      13184      4208
```

*Example 21: Output of the iostat command*

19. Remove the source LUNs from the volume groups and verify that they are removed using the *reducevg vg_name hdisk#* and *lsvg -p vg_name* commands (as shown in Example 22) where vg_name is the name of the volume group and # is the hdisk number.

```
root@sap-xiv-1: / > reducevg sapdatavg hdiskpower0

root@sap-xiv-1: / > reducevg sapdatavg hdiskpower1

root@sap-xiv-1: / > lsvg -p sapdatavg

sapdatavg:

PV_NAME           PV STATE          TOTAL PPs   FREE PPs    FREE DISTRIBUTION

hdisk16           active            831         20          00..00..00..00..20
```

*Example 22: Output of the reducevg vg_name hdisk# and the lsvg –p vg_name commands*

The output in Example 23 shows that *hdiskpower0* and *hdiskpower1* are no longer members of the volume group sapdatavg.

20. Use the lspv command (as shown in Example 23: Output of the lspv command) and verify that the previously migrated hdisks are no longer member of a volume group

```
root@sap-xiv-1: / > lspv

hdisk0          00cb33f21d40001a                    rootvg          active

hdisk1          00cb33f2f43e39ba                    sapISRvg        active

hdisk2          none                                None
```


© Copyright IBM Corporation, 2010

```
hdisk3          none                                None
hdisk4          none                                None
hdisk5          none                                None
hdisk6          none                                None
hdisk7          none                                None
hdisk8          none                                None
hdisk9          none                                None
hdisk10         none                                None
hdisk11         none                                None
hdisk12         none                                None
hdisk13         none                                None
hdisk14         none                                None
hdisk15         none                                None
hdiskpower0     00cb33f2fda738bd                    None
hdiskpower1     00cb33f2fdaaf7e2                    None
hdiskpower2     00cb33f2fdac10e1                    saplogvg        active
hdiskpower3     00cb33f20351a5ac                    saplogvg        active
hdiskpower4     00cb33f21cc8e907                    sapswvg         active
hdiskpower5     00cb33f238d614ce                    sapswvg         active
hdiskpower6     00cb33f238d7548d                    swapvg          active
hdisk16         00cb33f20a8d0fdb                    sapdatavg       active
```

*Example 23: Output of the lspv command*

21. Delete the device definitions from the host. Use the *rmdev -dl hdisk#* command where # is the hdisk number.
22. Verify that the device definitions are removed. Use the *lsdev -Cc disk* command and check that there are no defined disks.
23. Remove the source zone definition from the switch fabric.
24. In the legacy storage, unassign the LUNs from the AIX host.


© Copyright IBM Corporation, 2010

## Using XIV data migration function

There are many options available for migrating data from one storage system to another; the XIV Storage System includes an embedded data migration feature that enables the easy movement of data from an existing storage system to the XIV Storage System. This feature enables the production environment to continue functioning during the data transfer with only one brief period of downtime for your business applications.

The IBM XIV Data Migration solution offers a smooth data transfer, because it:

- Requires only a single short outage to switch LUN ownership. This enables the immediate connection of a host server to the XIV Storage System, providing the user with direct access to all the data before it has been copied to the XIV Storage System.
- Synchronizes data between the two storage systems using transparent copying to the XIV Storage System as a background process with minimal performance impact.
- Supports data migration from practically all storage vendors.
- Must be set up using Fibre Channel.
- Can be used to migrate SAN boot volumes.

The only drawback is that the XIV migration feature does not allow LUN consolidation. There always has to a one to one relationship between source and target volumes.

The basic migration steps performed in the XIV system GUI or command line interface (CLI) are:

1. **Install the new XIV and establish connections**
   The XIV storage system supports different versions of the AIX operating system, either through Fibre Channel or iSCSI connectivity. The setup used in this white paper is shown in Figure 1.

2. **Cable and zone the XIV system to legacy storage system for migration**
   Because the non-XIV system device views the XIV system as a host, the XIV system must be connected to the non-XIV Storage System as a Small Computer System Interface (SCSI) initiator. The initiator ports can then be directly attached to the legacy storage system as shown in Figure 3. The legacy storage also has to be defined as migration target on the XIV system. This includes defining the WWPNs of the legacy storage system ports. In the XIV system GUI, the connection between XIV Storage System and legacy storage can easily be done by drag and drop.


© Copyright IBM Corporation, 2010

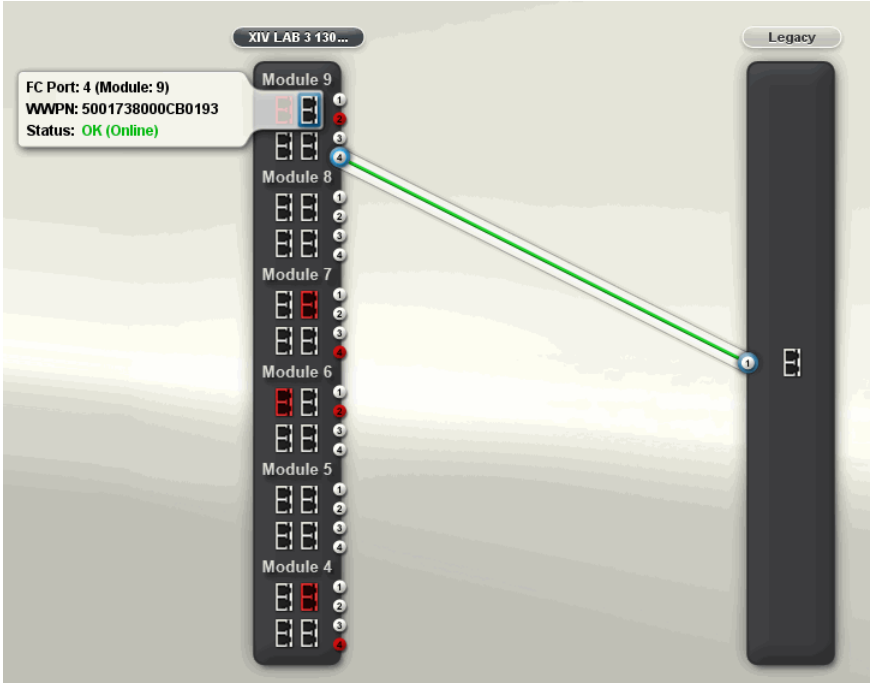


*Figure 3: Migration connectivity to the legacy storage system*

3. **Define XIV Storage System as host on the legacy storage system**
   This process is vendor and device dependant because the legacy storage system management interface has to used. The XIV system has to be defined as Linux® or Microsoft® Windows® host to the legacy storage system.

4. **Define data migration volumes on XIV system**
   The source volume from the legacy storage system and XIV system must be exactly the same size. In this case, one has an option of creating the volumes prior to defining the data migration object or have the XIV system create the XIV Storage System volumes whenever a data migration object is defined

5. **Define the SAP host being migrated to XIV system**
   Before starting data migration, the host has to be defined in the XIV system and has to be mapped to the XIV system target volumes. The XIV system is acting as a proxy between the host and the legacy storage system. The host must no longer access the legacy storage system after the data migration is activated. The host must perform all I/O through the XIV system.


© Copyright IBM Corporation, 2010

6. **Prepare the SAP host for data migration**
   Backup all the volumes being migrated and stop SAP. The host must be zoned only to the XIV system that is working as a proxy between the host and the legacy storage system. Install the XIV system host attachment kit and map the XIV system target volumes to the host. Remove the non-XIV multipathing drivers. Do not allow the host to detect the LUNs from both the XIV Storage System and the legacy storage system.

7. **On XIV system activate data migration and start SAP**
   When the data migration starts, the data is copied sequentially in the background from the legacy storage system to the XIV system. As soon as the migration is active, the SAP system can be restarted on the host. The host reads and writes data to the XIV system without being aware of the migration process. Do not deactivate the migration while the host I/O is running.

8. **Complete data migration**
   After all the volume data has been copied, the migrations process has the status, **synchronized**. The data migration object can be deleted after the migration process has reached the synchronization status without host interruption.

Visit IBM Redbooks to find *IBM XIV Storage System: Copy Services and Migration,* which provides a more detailed description of the XIV system migration feature and the setup steps with GUI and CLI at:

http://www.redbooks.ibm.com/abstracts/sg247759.html?Open


© Copyright IBM Corporation, 2010

# Conclusion

The usage of any new storage system frequently requires the transfer of a large amount of data from the existing or legacy storage system to the new storage system. Most businesses cannot afford their applications such as SAP to be offline while migrating data from a legacy storage system to a new storage system.

This white paper describes the advantages and the disadvantages of the two data migration procedures. The XIV system migration feature is embedded in the XIV Storage System and is easy to use. It does not require deep knowledge of the operating system and logical volume manager on the host. XIV system Data Migration feature enables online migrations without complex configurations, software, services, or capital expenditures that are required by legacy migrations. A small outage window is required to redirect the host from the legacy storage to the XIV System Storage. After redirecting the host to the XIV system array, no further outages are required. The data migration can be set up with the XIV system graphical user interface or the XIV system CLI. So, if the LUN layout from the legacy storage system should be identical on the XIV system, the XIV system migration feature is the right choice for data migration. The monitoring process of the data migration is also easy to follow.

The second procedure in this white paper describes data migration with the AIX logical volume manager. This procedure is purely operating system dependent and requires a good knowledge of the AIX logical volume manager and file system. For LUN consolidation, this is the right choice and the migration can also be done, while the SAP application is up and running. To obtain the best performance with the XIV system, it might be necessary to consolidate small LUNs from the legacy storage system to big XIV system LUNs. Generally with XIV system, there is no need to create a large number of small LUNs. For database applications such as SAP it is only important to have separate LUNs for log data and tablespace data. Using the LVM can reduce the outage required for migrating the data. Having the ability to consolidate LUNs while the system/applications are still running, is a huge benefit in an operational computing environment.


© Copyright IBM Corporation, 2010

## Resources

These Web sites provide useful references to supplement the information contained in this paper:

- IBM XIV Storage System Information Center
  http://publib.boulder.ibm.com/infocenter/ibmxiv/r2/index.jsp
- IBM System p and AIX Information Center
  http://publib.boulder.ibm.com/infocenter/pseries/index.jsp
- IBM Redbooks®
  **ibm.com**/redbooks
- System Storage Interoperation Center (SSIC)
  http://www-03.ibm.com/systems/support/storage/ssic/interoperability.wss
- IBM XIV Host Attachment Kit for AIX
  http://www-01.ibm.com/support/docview.wss?rs=1319&context=STJTAG&context=HW3E0&dc=D400&q1=ssg1*&uid=ssg1S4000802&loc=en_US&cs=utf-8&lang=en


© Copyright IBM Corporation, 2010

# Trademarks and special notices

© Copyright IBM Corporation 2010. All rights Reserved.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

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

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

SET and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local IBM office or IBM authorized reseller for the full text of the specific Statement of Direction.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is

presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.


© Copyright IBM Corporation, 2010