IBM Smart Analytics System

Best Practices

AIX Operating System Level Backup and Recovery for an IBM Smart Analytics System

Garrett Fitzsimons
IBM Smart Analytics Systems Best Practices

Richard Lubell
IBM Smart Analytics Systems Information Development

Ron Parrish
IBM Power Systems Lab Services

Issued 28 September 2010

Introduction

This paper will help you to develop a backup and recovery strategy for the operating system and configuration of an IBM® Smart Analytics System[1] solution based on the AIX® operating system and POWER™ processors.

The paper focuses on recovering the configuration of a replacement device and the configured operating system of a replacement server from the backup files of the following IBM Smart Analytics System components:

- Hardware Management Console or HMC (an IBM System xSeries® 3650 server)
- Server Management, Administration, User, Data, Application and Standby servers (IBM p6-550 servers)
- IBM System Storage ® SAN40B switches
- IBM System Storage DS5300 storage controllers
- Juniper EX4200-48T Ethernet switches
- IBM Remote Support Manager (RSM) for Storage (an IBM System x3650 M2 server)

The objective of the backup and recovery strategy is to resume normal operating behavior as quickly as possible when replacing a server or device.

An IBM Smart Analytics System 7600 R1 test cluster was used to test the configuration, backup and recovery of all recommendations. Sample commands, scripts and summaries of utilities are provided throughout the document. Use these commands as examples of how to implement a backup and recovery strategy on your cluster but be aware that the commands will need to be modified to work in your environment. Test all commands and scripts on a non-production system before implementing on a production system.

The replacement of the storage subsystem components is out of scope for this paper.

To use the information in this paper, you must have a working knowledge of the IBM Smart Analytics System solution.

[1] *The IBM Smart Analytics System solution is an evolution of InfoSphere Balanced Warehouse solution. They are based on the same storage and database design principles. This paper deals with configurations based on the AIX operating system and POWER processors, however, all content applies to IBM Smart Analytics System configurations and InfoSphere Balanced Warehouse configurations. The content also usually applies to custom data warehouses that are based on similar design principles, although some changes might be needed depending on the specific environment.*

Summary of components

Backing up the current configuration of each device and server in the IBM Smart Analytics System enables you to install a replacement server or device efficiently and effectively.

The frequency of configuration changes to each device and server is different for each implementation. This determines the frequency of backups. For example, in an environment where you add and remove users often, it might be necessary to back up the affected nodes more frequently.

Align the frequency of the backup with the frequency of configuration changes for each device or server.

HMC (IBM System x3650 Server)
The Hardware Management Console (HMC) is used to manage and monitor the IBM Power servers using a separate network called the HMC network. Failure of the HMC affects the ability to manage and monitor the IBM Smart Analytics System; however, other servers operate normally. You must recover the HMC if the server fails or the configuration has been compromised and you cannot correct it.

Management server (IBM p6-550 server)
The management server helps automate cluster deployment and runs management software such as DB2 Performance Expert and the IBM DS Storage Manager Software. If the management server fails or the configuration becomes compromised, the other servers still function, and the data warehouse and user requests are unaffected. However, the ability to manage and monitor the system and to perform backups is impaired, so it is crucial to recover the management server as soon as possible. Recover the management server by writing the backup images to DVD then reinstalling from these DVDs.

Administration, user, data, application and standby servers (IBM p6-550 servers)
These retrieve data for the end user and are critical to system availability. Operating in a high availability (HA) environment minimizes downtime; failure of any of these servers increases downtime if no standby server solution is in place. Use the management server and the HMC to reinstall the original or replacement server over the management network from a backup image.

IBM System Storage SAN40B switches
The storage area network (SAN) provides access to the storage system. For redundancy in case of single failure, switches are configured in pairs. Failure of an individual SAN switch is transparent although network speed is affected. Recover a SAN switch by connecting the new SAN switch to the fiber network and restoring the configuration from the backup file.

DS5300 storage controllers
The IBM Smart Analytics System solution contains a DS5300 Storage Controller Unit, composed of dual controllers for redundancy, to manage the SAN. Failure of both controllers requires a manual recovery process. Recover a storage controller through editing the backup file and reapplying the configuration to the replacement controller through the command-line interface.

Juniper EX4200-48T Ethernet switches

The Juniper Ethernet switches provide the platform for the HMC, management and Fast Communications Manager (FCM) networks used in the IBM Smart Analytics System. The Juniper switches include HA features to help ensure maximum uptime. The switches are configured in a set of two or more and each server is connected to two individual Juniper switches to provide redundancy. Failure of an individual switch is apparent only if network speed decreases for a server that is attached to a failed switch. Recover a Juniper switch by using the setup wizard and restoring the configuration from the backup file.

RSM for Storage (IBM System x3650 M2 Server)

RSM for Storage provides a "callback" mechanism through which faults that are related to the storage system are logged directly with IBM Smart Analytics System Support. Failure of the server hosting RSM for Storage does not affect critical operations. Recover the RSM by reinstalling it using the provided disks and recovering the configuration from the backup files.

Overview of a backup and recovery strategy

When creating a backup and recovery strategy, address the following areas:

- What to back up
- How to back up
- When to back up
- Where to store the backup files
- How to manage the storage of files
- Using a storage manager to archive and retrieve
- How to recover a server or device

What to back up
Back up the configuration of each server and device individually. You cannot recover a server or device from a backup associated with a different server or device without further complications; this is not recommended.

How to back up
Script and automate backups to ensure all files and images are created in the appropriate directory with descriptive names. Perform backups over the management network to avoid unnecessarily affecting the query workload.

When to back up
Back up all components automatically on a schedule determined by the frequency of configuration changes to each device and server. To minimize the effect on the system resources, run backups during an off-peak period. Take additional backups when performing upgrades or when otherwise advised.

Where to store the back up files
Store all backup files centrally on a management server file system on the DS5300 storage system. The backups are saved securely on the storage subsystem when a server or device fails.

How to manage the storage of files
To facilitate recovery, ensure that there is enough space on disk to address the next scheduled backup cycle while retaining the last backup copy for each server or device. The amount of space that is required to accommodate a single backup cycle depends on the number of data modules in your system. For details on the available capacity for the disk size that you choose for your system, see Appendix C.

Use TSM to manage the storage and retrieval of backup files for all backup and recovery processes.

Using a storage manager to archive and retrieve
Tivoli® Storage Manager (TSM) is an optional IBM software tool that protects data from hardware failures and errors by storing data on offline and offsite storage. Use TSM to manage the storage of backup images, operating system files, scripts, and other important files associated with the data warehouse. You can also use TSM to schedule the execution of backup scripts. For information about TSM, see the Further reading section.

How to recover a server or device

Reference the xcluster.cfg file or installation record for details of the failed device or server before starting the recovery process. Minimize recovery time by documenting and testing each recovery process. Create detailed instructions for your environment based on the examples in this paper. Keep the last backup of each server or device on the management server and refer to the storage manager for older backups when needed.

Document and test each recovery process to help ensure successful recovery from server or device failure.

Pre-Requisites for IBM Smart Analytics System Backup and Recovery

Before implementing a backup and recovery strategy for the IBM Smart Analytics System ensure that you meet the following pre-requisites:

Backup space
Calculate the amount of space that you need to complete a single backup of each server or device using the information in the following table. The tilde symbol (~) indicates an approximate value.

Component	Backup Size	Comment
HMC	~4GB minimum	Increases where the number of data nodes is greater than two.
Management server	15GB	Requires four to five DVD's when creating for recovery.
Administration, application or user server	5GB per server	
Data server	5GB x 2 servers = 10GB	
Standby server	5GB per Server	
IBM System Storage SAN40B switch	~24KB per component	
DS5300 storage system controllers	~2.5MB per component	
Juniper EX4200-48T Ethernet switch	~1MB for entire Juniper configuration	
RSM for Storage	~1MB	Backup configuration file

Installation record
Obtain this spreadsheet from IBM Smart Analytics Support; it contains a diagram of the cluster as installed by the IBM Lab Services team. Work with IBM Smart Analytics System Support to keep the spreadsheet current as your cluster changes. Use this record to ensure that you reference all servers or devices in the backup strategy.

Cluster configuration file
Acquire the `xcluster.cfg` file which is used during installation that contains all IP addresses of devices that are used in the cluster. Work with IBM Smart Analytics System Support to maintain this file as your cluster changes. The `xcluster.cfg` file is in either the `/ISAS/cfg/xcluster.cfg` or `/BCU_share/installation/xcluster.cfg` directory on the management server.

Save and maintain a copy of the installation record and `xcluster.cfg` configuration file for your IBM Smart Analytics System.

Configuring IBM Smart Analytics System for Backup and Recovery

Perform the steps in this section before you start to back up your system. Log in as root for all operations in this paper unless otherwise directed.

Setting up a central location for backup files
You can simplify the management of backups by centralizing the backup files on a shared file system on the management server and using TSM to manage the archival and retrieval of the files. Maintaining the backup files for the servers or devices in individual file systems creates unnecessary risks and requires more effort.

To centralize access to backup files:

1. Configure at least 250 GB of disk capacity on the storage subsystem for a new file system on the management server. Step-by-step instructions for configuring disk capacity are in Appendix C.

Configure available disk capacity on the storage system to hold backup files centrally on the management server.

2. On the shared file system, implement a directory structure that has separate subdirectories for each server or device to be backed up. Using the following commands, you can implement a directory structure that supports different archiving and retention policies for each server and device backup process:

.

```
mkdir -p /backups/mgmt_hostname
mkdir -p /backups/admin_hostname
mkdir -p /backups/user_hostname
mkdir -p /backups/data1_hostname
mkdir -p /backups/data2_hostname
mkdir -p /backups/standby_hostname
mkdir -p /backups/appserver_hostname
mkdir -p /backups/hmc_cfg
mkdir -p /backups/ds5300_cfg
mkdir -p /backups/juniper_cfg
mkdir -p /backups/san_cfg
```

3. Specify which files and directories to exclude from server backups. You can exclude files and directories that are not needed and might cause the backup to fail by placing entries in the `/etc/exclude.rootvg` file on each server. Issue the following command on each server that you intend to back up. You might want to add directories that are specific to your environment.

```
cat /etc/exclude.rootvg
/backups/
/tmp/
/var/tmp/
/dev/__
/var/perf/pm/daily
/etc.vg/
```

Configuring passwordless SSH on the Juniper and SAN switch devices

Secure Shell (SSH) is a network protocol installed by default on the IBM Smart Analytics System that provides a secure channel between two network devices for data exchange. DB2 software uses SSH to communicate between servers. By default, when connecting through SSH, the SAN and Juniper switches prompt for a password each time they are accessed. To enable the backups to be scripted and automated, configure SSH to bypass the password prompt. For the steps to configure passwordless SSH for the Juniper and SAN switches, see appendix B.

Use the secure copy protocol (SCP) for the secure transfer of backup files from the device to the management server. SCP is based on the SSH protocol and provides encryption and authentication. SSH uses public-key cryptography to handle authentication between the management server and the device.

Configuring SSH access to the RSM for Storage Server

By having two different SSH daemons for remote and local connections, RSM for Storage can help ensure that there is no remote access to the system using the root user ID. Attempts to log into this SSH port as the admin, lservice or rservice user will cause the session to be dropped.

To establish local SSH connections, choose a port number other than default for security reasons. The port number to be used for local SSH connections is specified by editing the file: /etc/rsm/rsmfirewall.
conf.

To create this file, log in as root and rename /etc/rsm/rsm-firewall.conf.sample to /etc/rsm/rsm-firewall.conf and uncomment the following lines:

```
SSHD_ALLOWED 23
SSHD_ROOT_ALLOWED
```

Specify the port number to be used for local SSH access following the keyword SSHD_ALLOWED. The test cluster used port 23. Run `rsm-start` to read the configuration changes and initialize the SSH daemon for local access. Review the last entries in the file RSM security log file to verify that the SSH setup is complete.

IBM Hardware Management Console

The Hardware Management Console (HMC), based on the IBM System xSeries hardware architecture, is a server that you can use to manage an AIX cluster. The HMC comes with the media that you require to perform an installation, eliminating the need to back up the operating system. The HMC also contains configuration information and console data for the managed cluster servers. The configuration data is captured through the command line on the HMC while the console data is captured through a CLI utility.

Use the configuration diagram and the `xcluster.cfg` file to confirm that all components of the IBM Smart Analytics System are visible to the HMC.

Backing up the HMC

To back up, the SSH utility is used create a copy of the configuration of the HMC.

To recover an HMC server, separate backups of the following are required:

- HMC configuration
- HMC console data

Back up HMC Configuration

Create a file containing information on how the HMC is configured and what software versions are installed by capturing the output from the `lshmc` command. An example follows:

```
ssh hscroot@e-hmc01 "lshmc -V ; lshmc -l; lshmc -b; lshmc -v; lshmc -n ; lshmc
-r" > /backup/hmc_cfg/hmc_config.txt
```

The file created is a relatively small text file which will be less than 1MB in size and is used to verify the recovery process.

Back up HMC Console data

Create the console data file containing detailed information on user configuration and servers managed. An example of the `bkconsdata` command follows:

```
ssh hscroot@e-hmc01 "bkconsdata -r nfs -h 172.22.1.10 -l /backups -d hmc_cfg"
```

where:

- e-hmc01 is the name of the HMC server
- 172.22.1.10 is the management server

The backup file is a minimum of 3GB and takes a minimum of 15 minutes to complete. The size and time will increase as the number of servers being managed increases.

Recovering the HMC

To recover, the SSH utility is used to recover the configuration after a standard installation.

The recovery process for the HMC consists of the following steps:

1. Verify HMC version and acquire the appropriate installation media
2. Perform a standard installation of HMC from media
3. Determine the IP address of the failed HMC
4. Restore the HMC configuration from backup
5. Verify the recovery

Verify HMC version and acquire appropriate installation media

Ensure that the correct version of HMC software is installed as specified in the validated support stack. To identify the version that is installed on the HMC that you are recovering, refer to the HMC backup file. The following command issued against the backup configuration file on our test cluster shows that Version 7.1.0 Service Pack 1 is installed:

```
# cat /backup/hmc_cfg/hmc_config.txt
```

The following output is an excerpt from the backup file:

```
"version= Version: 7
Release: 7.1.0
Service Pack: 1
HMC Build level 20100225.1
```

Perform a standard installation of the HMC from media

To perform a standard installation of the HMC, follow the instructions on the recovery disks that are supplied with the HMC.

If the new HMC has a newer version than that previously installed, use the previous HMC install media or download the install software from Fix Central to configure HMC to the version appropriate to your cluster. Consult with IBM Smart Analytics System Support if you are unsure as to what version is appropriate to your cluster.

Determine the IP address of the failed HMC

Connect the HMC to the cluster network using the network settings in the backup file. The backup file on the test cluster uses the IP address highlighted in the following file excerpt. The management network on the test cluster is represented by `ipv4addr_eth3`:

```
hostname=p7hmc01,domain=mul.ie.ibm.com,"ipaddr=192.168.128.1,0.0.0.0,9.162.88.1
15,172.22.1.115","networkmask=255.255.128.0,255.255.255.0,255.255.255.0,255.255
.255.0",gateway=9.162.88.1,"nameserver=9.64.162.21,9.64.163.21","domainsuffix=m
ul.ie.ibm.com,ibm.com",slipipaddr=10.253.0.1,slipnetmask=255.255.0.0,"ipaddrlpa
r=9.162.88.115,172.22.1.115,192.168.128.1","networkmasklpar=255.255.255.0,255.2
55.255.0,255.255.128.0","clients=192.168.252.255,192.168.255.248,192.168.253.25
3,192.168.251.255,192.168.253.252,192.168.254.250,192.168.254.251,192.168.252.2
54,192.168.255.249,192.168.255.250,192.168.255.254,192.168.254.255,192.168.255.
```

```
253,192.168.255.252,192.168.254.254,192.168.253.255,192.168.254.253,192.168.255
.251,192.168.254.252,192.168.253.254","ipv6addrlpar=fe80:0:0:0:214:5eff:fe76:96
ee,fe80:0:0:0:214:5eff:fe76:96ef,fe80:0:0:0:e61f:13ff:fe2d:7378","slpipaddrs=9.
162.88.115,192.168.128.1,172.22.1.115,fe80::e61f:13ff:fe2d:7378/64,fe80::214:5e
ff:fe76:96ee/64,fe80::214:5eff:fe76:96ef/64",ipv4addr_eth0=192.168.128.1,ipv4ne
tmask_eth0=255.255.128.0,ipv4dhcp_eth0=off,ipv6addr_eth0=fe80:0:0:0:e61f:13ff:f
e2d:7378/64,ipv6auto_eth0=off,ipv6privacy_eth0=off,ipv6dhcp_eth0=off,lparcomm_e
th0=off,jumboframe_eth0=off,speed_eth0=auto,duplex_eth0=auto,tso_eth0=off,ipv4a
ddr_eth1=0.0.0.0,ipv4netmask_eth1=255.255.255.0,ipv4dhcp_eth1=off,ipv6addr_eth1
=,ipv6auto_eth1=off,ipv6privacy_eth1=off,ipv6dhcp_eth1=off,lparcomm_eth1=off,ju
mboframe_eth1=off,speed_eth1=auto,duplex_eth1=auto,tso_eth1=off,ipv4addr_eth2=9
.162.88.115,ipv4netmask_eth2=255.255.255.0,ipv4dhcp_eth2=off,ipv6addr_eth2=fe80
:0:0:0:214:5eff:fe76:96ee/64,ipv6auto_eth2=off,ipv6privacy_eth2=off,ipv6dhcp_et
h2=off,lparcomm_eth2=off,jumboframe_eth2=off,speed_eth2=auto,duplex_eth2=auto,t
so_eth2=off,ipv4addr_eth3=172.22.1.115,ipv4netmask_eth3=255.255.255.0,ipv4
dhcp_eth3=off,ipv6addr_eth3=fe80:0:0:0:214:5eff:fe76:96ef/64,ipv6auto_eth3=off,
ipv6privacy_eth3=off,ipv6dhcp_eth3=off,lparcomm_eth3=off,jumboframe_eth3=off,sp
eed_eth3=auto,duplex_eth3=auto,tso_eth3=off
```

Restore the HMC console configuration file

To restore the configuration data:

1. Click the **Restore HMC Data** menu option.
2. Follow the instructions in the restore wizard. Specify information as follows:
   - Select the NFS option.
   - Enter `/backups` as the mount resource.
   - Enter `hmc_cfg` as the source directory where the backup file(s) exists on the /backups directory.
   - When prompted, verify the correct backup file to use.

   The backup file is copied.
3. When prompted, reboot. The configuration, including the original password, is restored.

Verify the recovery

Run the SSH command used to back up the configuration and confirm that the output on the recovered server matches the contents of the backup file.

IBM Smart Analytics System Management Server

The management server monitors the performance of the data warehouse using DB2 Performance Expert and other DB2 software. The management server does not affect the availability of the data warehouse or associated applications, but does impact the ability to manage and monitor the other servers in the cluster, making speedy recovery critical.

Because the NIM server is installed on the management server it will not be available to the recovery process. Instead, create a set of bootable DVD ISO images and recover the management server from DVD. Make a backup of the management server after any changes, such as the loading of new fixpacks or operating system upgrades.

Burn DVDs for recovery only when recovery is necessary.

A configuration diagram of the management server is in the installation record and the network configuration is listed in the `xcluster.cfg`; reference these during recovery.

Backing up the Management Server

An ISO image is created in the `/backups` file system. The `mkdvd` command as follows makes a set of ISO images (four or more DVD ISO images are made).

```
mkdvd –I /backups/mgmtserver –V rootvg –c –e –S
```

where:

- The -c option suppresses the use of compression when creating the backup image. Using this option is recommended to avoid possible failure when uncompressing encrypted files.
- The -e option excludes files as specified in `/etc/exclude.rootvg`.
- The –S (uppercase) option saves the ISO images instead of writing them directly to DVD RAM

The images in `/backups` will have names similar to `cd_image_245824.vol1`. These images need to be archived off the Management server to TSM.

Recovering the Management Server

To recover the management server:

1. Burn files to bootable DVD
2. Set the management server to boot from the SMS menu
3. Boot from the DVD

Burn files to bootable DVD

Write the backup ISO images to DVD using any standard DVD creation tool. Give the images a filename and suffix (such as `mgmt_disk1.iso` on a PC) that allows the SMS menu software to recognize the files as ISO images. Typically, the backup is approximately 15 GB in size and requires four to five DVDs.

Set management server to boot from SMS menu

The management server has the following attributes defined on the HMC by default:

- The system name for the management server on the HMC is `BCUMGR`
- The LPAR name is `sysNode`
- The LPAR profile is `sysDefault`
- The partition ID is 1

Activate the management server through the HMC GUI (web browser) or the HMC CLI to boot to the SMS menu so that the DVD drive can be set as the IPL device.

To boot to the SMS menu:

1. Access the HMC CLI from the HMC or by entering the following command:

   ```
   ssh hscroot@e-hmc01
   ```

2. Power off the management server by entering the following command:

   ```
   chsysstate –r lpar –m BCUMGR –o shutdown --immed –n sysNode
   ```

3. Power on the management server and issue the following command to boot to the SMS menu:

   ```
   chsysstate –r lpar –m BCUMGR –o on –n sysNode –f sysDefault –b sms
   ```

Boot from DVD

1. Place the first DVD in the DVD drive for the Management server (BCUMGR).

2. Log in to the HMC as the `hscroot` user and issue the `vtmenu` command to access the interactive `vtmenu` utility:

   ```
   hscroot@scschma01~> vtmenu
   ```

3. Select the following menu options to complete the instruction to boot from DVD:

   a. When prompted to enter the number of the Managed System, select the number associated with **BCUMGR**.
   b. When prompted to enter the number of the Running Partition, select the number associated with **sysNode**.
   c. From the main menu, select option 5: **Select Boot Options.**
   d. From the **Multiboot** menu, select option 1: **Select Install/Boot Device.**
   e. When prompted to select a device type, select option 3: **CD/DVD**.
   f. When prompted for a media type, select option 5: **SATA**.
   g. When prompted for a media adaptor, select option 1:

“/pci@800000020000200/pci1014,02BD@1/sata. Ensure that the device listed corresponds to your media adaptor. If not, select option 2 for a list.

h. When prompted to select a task, select option 2: **Normal Mode Boot**.

i. When prompted “Are you sure you want to exit System Management Services?” select option 1.

j. From the **Installation and Maintenance** menu, select option 2: **Change/Show Installation Settings and Install**.

k. A summary of the options selected are displayed. Enter 0 and click **Enter to Install**, ensuring that the settings are as indicated before.

4. Insert each of the DVDs to load as prompted by the HMC interface. The server reboots after the process is complete.

5. Review the server storage, links, settings, and configuration to ensure the recovery is correct and complete.

Administration, User, Data, Application and Standby Server

Each server in the IBM Smart Analytics System performs a critical function in the management and delivery of data to the end user. The same backup and recovery strategy is used for the administration, user, data, application, and standby servers.

A configuration diagram is in the installation record and the network configuration is listed in the `xcluster.cfg` file; reference these during recovery.

About Network Install Manager (NIM)
NIM is an AIX operating system feature that you can use to install, update and maintain AIX servers over a network. NIM is used in the installation of the IBM Smart Analytics System and is installed and configured on the management server by default.

The management server is referred to as the NIM master, where the NIM application and repository are installed. The other servers in the cluster are defined on the NIM master as NIM clients. Each file, directory, network, and server is referred to as a resource.

In this paper, the following three resource types are mentioned in references to recovering a server using NIM:

- Mksysb
  The `mksysb` command is the standard AIX command for creating bootable backup images of the `rootvg` volume group that can be used to reinstall a server to its original state.
- lpp-source
  The *lpp-source* represents a directory in which the software installation images are stored, much like that of a CD or DVD.
- Shared Product Object Tree (SPOT)
  SPOTs are used to support NIM operations that require a server to boot over the network. The SPOT resource contains the equivalent of the `/usr` file system on an AIX operating system. During the installation over the network the client mounts this resource over the network to enable the boot process.

Backing up an Administration, Data, Application, User or Standby Server

Each server, other than the management server, is defined on the NIM master as a NIM client. The same command syntax is used to backup each type of server. Backup files are written to a separate dedicated directory for each server in the shared /`backups` file system on the management server.

Use the following tips to build your backup command:

- Use the `mksysb` command to create a backup of a server
- Name the backup file from the mksysb output with a descriptive name
- Automate the backup process, incorporating `mksysb`

Use the mksysb command

Three files will be created when the `mksysb` command is issued:

1. `20100624_bluejay02_mksysb`

   This is the system backup image. The file name is determined by the command line parameter.
2. `image.data`

   This file contains information on how the physical disks and file systems is configured on the root volume group during installation or recovery. Note that the file is created on absolute path '/' and will overwrite any existing files.
3. `bosinst.data`

   This file contains information used by the installation program when installing or restoring a server. Note that the file is created on absolute path '/' and overwrites any existing files. No parameters are required to create this file.

The command issued on the test cluster to back up with administration server `bluejay02` is:

```
mksysb -iep /backups/bluejay02/20100624_bluejay02_mksysb
```

where:

- The '-i' option creates an `image.data` file which ensures the file system information is known
- The 'e' option excludes files as specified in `/etc/exclude.rootvg`
- The 'p' option suppresses the use of compression. Compression is not recommended for encrypted files
- `20100624_bluejay02_mksysb` is the name of the backup image to be created

Name the backup file

Rename both the `image.data` and `bosinst.data` files with to indicate that they are connected with the backup image, and move them to the appropriate backup directory before archiving to TSM. For example:

```
mv /image.data /backups/bluejay02/20100624_bluejay02_image.data
mv /bosinst.data /backups/bluejay02/20100624_bluejay02_bosinst.data
```

Use backup file names that make it easy to identify the origin and content of the files

Automate the backup process

Use the SSH utility to issue the `mksysb` command from the management server though the management network. In the following example, `bluejay02mgt` is the host name that is configured with the management network for the `bluejay02` server:

```
ssh bluejay02mgt “mksysb -iep /backups/bluejay002/20100624_bluejay02_mksysb”
```

Back up files through the management network to avoid conflict with the Fast Communications Manager (FCM) network.

In the following example, the `mksysb` command is issued on the management server to back up three servers through the management network. Each backup file that is produced is renamed to a unique file name, and all files are placed in the appropriate backup directory:

```
for H in bluejay02, bluejay03, bluejay04
do SVR="${H}mgt"
echo "Backing up ${SVR}"
ssh ${SVR} "mksysb -iep /backups/${H}/${H}_`date +%Y%m%d`
mv /image.data /backups/${H}/${H}.image.data.`date +%Y%m%d`
mv /bosinst.data /backups/${H}/${H}bosinst.data.`date +%Y%m%d`  "
done
```

Recovering an Administration, Data, Application, User or Standby Server

The recovery process for a NIM client uses a backup image that contains the specific configuration of the failed server. To perform a network install, the server is shut down and booted over the network. The HMC, through the NIM master on the management server, drives the process. The steps of this process are:

1. Define the NIM resources for the backup files
2. Prepare the NIM Server to install the AIX operating system
3. Collect information that is needed to recover the server
4. Initiate the NIM restore operation
5. Shut down the server
6. Boot the server from the network from the HMC
7. Monitor the boot process

The steps are explained using examples for the test cluster. Root access on both the management server and HMC is required.

Define the NIM resources

You must define a NIM resource on the management server for each of the three files that you will use to reinstall the server. To define NIM resources:

1. Define a NIM resource for the `image.data` file that is associated with the server backup by issuing the following command on the management server:

   ```
   nim -o define -t image_data -a location=/backups/ bluejay02/image.data.20100723
   -a server=master bluejay02_imagedata
   ```

2. Define a resource for the base operating system file that is associated with the server backup by issuing the following command:

   ```
   nim -o define -t bosinst_data -a location=/backups/
   bluejay02/bosinst.data.20100723 -a server=master bluejay02_bosinstdata
   ```

3. Define a resource for the `mksysb` command file that is associated with the server backup by issuing the following command:

   ```
   nim -o define -t mksysb -a location=/backups/bluejay02/bluejay02_20100723 -a
   server=master bluejay02_mksysb
   ```

4. Create the SPOT resource by issuing the following command:

```
nim -o define -t spot -a location=/backups/bluejay02/ -a source=
bluejay02_mksysb -a server=master bluejay02_spot
```

5. Verify that you successfully defined the resources by issuing the `lsnim` commands, as follows:

```
lsnim -t mksysb
lsnim -t image_data
lsnim -t bosinst_data
lsnim -t spot
```

Prepare the management server to install the AIX operating system

You must verify that the server is ready for a NIM operation before initiating the installation. To verify that a server is ready:

1. List status information for each NIM client by issuing the following `lsnim` command on the management server:

```
lsnim -a Mstate -a Cstate -a info
```

   The output on the test cluster is as follows:

```
master:
Cstate = ready for a NIM operation
Mstate = currently running
adminnode:
Cstate = ready for a NIM operation
Mstate = currently running
```

2. If the 'Cstate' value for the server to be recovered is not "ready for a NIM operation" state, reset the server. In the following example, bluejay02 is the server being reset:

```
nim -o reset -F bluejay02
```

Collect information that is needed to recover the server

To collect the information that is needed to recover the server:

1. Identify the LPAR, managed system name and profile of the server you are restoring by issuing the `lssyscfg` command, as shown in the following example:

```
lssyscfg -r sys -F name
```

2. Use the managed systems, the names of the LPAR, and the name of the LPAR profile for the server you are restoring by issuing the `lssyscfg` command on the HMC, using the name you identified in the previous step:

```
lssyscfg -r lpar -m ADMINNODE -F name:curr_profile
```

   The output from the command lists the managed system name and profile of the failed server:

```
adminnode:adminDefault
```

3. Identify the location code, a parameter of the logical partition (LPAR) boot command, for the Host Ethernet Adaptor (HEA) for the server that you are restoring by logging in to the HMC as `hscroot` and issuing the `lshwres` command as follows:

   ```
   lshwres -r hea -m ADMINNODE --rsubtype phys --level sys
   ```

   where:

   - ADMINNODE is the LPAR name on the HMC for the administration server `bluejay02`, referred to as a managed system on the HMC, which you are restoring.

   The output on the test cluster is as follows:

   ```
   adapter_id=23000000,phys_loc=U78A0.001.DNWH8HG-P1,state=functional
   ```

4. Establish the state of the LPAR by issuing the following command:

   ```
   lssyscfg -r lpar -m ADMINNODE -F name:state
   ```

   The output on the test cluster shows that the LPAR is running:

   ```
   adminnode:Running
   ```

5. Get the IP address of the management server that you are recovering by issuing the `grep` command against the `/etc/hosts` file on the management server, where all server IP addresses are located. In the following example, the `grep` command searches for the IP address of the management server on the FCM network:

   ```
   grep bluejay02 /etc/hosts
   ```

Initiate the NIM restore operation

To start the NIM operation to recover the server, issue the `nim` command on the management server.

Issuing the following command on the management server initiates the NIM restore operation on the test cluster:

```
nim -o bos_inst -a source=mksysb -a mksysb=bluejay02_mksysb -a spot=
bluejay02_spot -a bosinst_data= bluejay02_bosinstdata -a image_data=
bluejay02_imagedata -a no_nim_client=yes adminnode
```

where:

- `bluejay02_mksysb` is the `mksysb` resource that was defined for the backup file restored from TSM
- `bluejay02_bosinstdata` is the `bosinstdata` resource that was defined for the backup file restored from TSM
- `bluejay02_imagedata` is the `imagedata` resource that was defined for the backup file restored from TSM
- `bluejay02_spot` is the `spot` resource that was created

- adminnode is the NIM client name to be restored

Important: You must activate the server within a few minutes of initiating the NIM server to load the network image. Activate the server by issuing the lpar_netboot as described later in this section.

Shut down the server prior to network boot and install

Shut down the server that you are recovering before booting it for the network installation. To shut down the server, log in to the HMC as hscroot and issue the chsysstate command as shown in the following example where bluejay02 is the name of the server that is being recovered:

```
chsysstate -r lpar -m bluejay02 -o shutdown --immed -n bluejay02
```

Boot the server from the network from the HMC

Use the HMC to activate the server that is being recovered so that it can load the image from NIM. You can access the HMC through a web browser or the command line interface (CLI.) The examples here use the CLI to load the AIX backup image over the network.

To use the HMC to activate the LPAR, issue the lpar_netboot command as shown. The SSH protocol is used to access the HMC, referenced here by IP address, from the management server. Activating the LPAR causes it to begin loading the AIX image from the NIM server.

```
lpar_netboot -t ent -T off -s auto -d auto -S 172.22.1.10 -G 172.22.1.10 -C
172.22.1.101 -l U78A0.001.DNWH8HG-P1-C6-T1"adminnode" "adminDefault"
"ADMINNODE"
```

where:

- -S (NIM server) specifies the IP address of the management server on the management network
- -G specifies the gateway IP address on the management network (to the management server)
- -C (NIM client) specifies the IP address on the management network of the LPAR to be restored
- -l specifies the HEA location that is connected to the management network. Add the -C6-T1 (Card 6, Terminal 1) to the end of the location code that you obtained in a previous step.
- "adminnode" is the LPAR name
- "adminDefault" is the profile name
- "ADMINNODE" is the managed system name

The lpar_netboot command takes some time to run. On the test cluster restoring the administration server took 20 minutes but times will vary in different environments.

Monitor the boot process

There are two ways to monitor the boot process, either from the management server or from the HMC. To monitor the progress of the **lpar_netboot** command as it recovers the server, issue the lsnim on the management server:

```
lsnim -a Mstate -a Cstate -a info
```

To monitor the recovery process, you must log on to the HMC as the hscroot user and open a session for the LPAR that you are recovering. Issue the `vtmenu` command to access the interactive `vtmenu` utility. The LPAR boots from the network image on the NIM server and guides you through the recovery process using the `vtmenu` utility interface. Enter the following information:

- Enter the number that is associated with the server that you are restoring
- Select 1 for English. Enter 88 for help with other languages
- Select option 2, **Change/Show Installation Settings and Install** to ensure that your settings are correct
- Enter 0 and press Enter to install with current settings

The server reboots after loading is complete. The `vtmenu` utility reports progress via the console. Log in to the HMC and start the `vtmenu` to monitor the progress of the recovery. The output from the `lpar_netboot` command on the test cluster was as follows:

```
# Connecting to appnode_1
# Connected
# Checking for power off.
# Power off complete.
# Power on appnode_1 to Open Firmware.
# Power on complete.
# Getting adapter location codes.
# Network booting install adapter.
# spanning tree.
# bootp sent over network.
# Network boot proceeding, lpar_netboot is exiting.
# Finished.
hscroot@p7hmc01:~>
```

Verify the server recovery

When the recovery process is complete, shut down and restart the recovered server as described in the IBM Smart Analytics System user guide. Verify that the server is fully functional, indicating the recovery was successful.

If high availability (HA) is enabled, the server should failover or failback as necessary. For detailed information about managing High Availability functionality, see the IBM Smart Analytics System user guide. To ensure that the peer domain and servers in a HA cluster are online, use the `lsrpdomain` and `lsrpnode` commands.

IBM System Storage SAN40B Switch

The SAN switches provide connectivity between the servers and the storage subsystem. The SAN switches are configured independently of each other and are in pairs for redundancy. The SAN switches are configured for optimum performance; you should not modify them.

For assistance with default passwords for these components refer to the documentation or contact IBM Smart Analytics System Support.

Refer to Appendix B for details on how to set up SSH for passwordless access to the SAN switch.

Setup passwordless SSH access to enable single command backup tasks for easier automation.

Backing up a SAN40B Switch

To back up, the `configupload` utility is used to create a copy of the zoning configuration rules for each switch.

To back up a SAN switch, you must:

- Identify the IP addresses of the SAN switches
- Back up the configuration of each SAN switch
- Save the license ID to a file

Identify the IP addresses

Determine the SAN switch IP addresses by querying the `xcluster.cfg` file. The `xcluster.cfg` file is in either `/ISAS/cfg/xcluster.cfg` or `/BCU_share/installation/xcluster.cfg` directory on the management server.
The test cluster configuration is queried by issuing the following command on the management server:

```
grep SAN /BCU_share/SAN/xcluster.cfg
```

The output shows the IP addresses of the two SAN switches:

```
SAN_SWITCHES = 2 # Number of SAN switches to be zoned
SAN_SWITCH1 = 172.22.1.113 # SAN Network Switch IP. It should be Mgmt Net IP.
SAN_SWITCH2 = 172.22.1.114 # SAN Network Switch IP. It should be Mgmt Net IP.
SAN_SWITCH_PASS =
SAN_FRAME1_IP = 172.22.1.111 172.22.1.112
SAN_FRAME1_PASSWD =
```

Back up each configuration

To replace a SAN switch you need a backup copy of the SAN switch configuration file which contains the following items:

- Zoning rules that are unique to each SAN switch
- A unique World Wide Name (WWN) required for licensing activation
- The firmware level that must be on the replacement switch

To back up the configuration file for a SAN switch, issue the `configupload` command. A sample command follows:

```
ssh -t admin@172.22.1.113 "configupload -all -p scp
\"172.22.1.10\",\"root\",\"/backups/san_cfg/san_sw_config.txt\""
```

where:

- 172.22.1.10 is the IP address of the management server
- -p scp specifies the scp protocol.
- 172.22.1.113 is the IP address of the SAN switch to backup
- backups/san_cfg/san_sw_config.txt is the target directory and file on the management server

Save the license ID

You must know the WWN, also referred to as the license ID, of the failed SAN switch to transfer the port activation licenses. Instead of viewing this information in the configuration file, you can view it or save it to a file by issuing the `licenseidshow` command from the management server as shown in the following example:

```
ssh admin@172.22.1.200 "licenseidshow" >\
/backups/san_cfg/san_licence_172.22.1.113
```

Sample output from the command follows:

```
10:00:00:05:1e:89:c0:77
```

Recovering a SAN40B Switch

To recover, the `configdownload` utility is used to reconfigure the replacement switch using the backup configuration file. Assistance from IBM Smart Analytics System Support is required to activate the relevant licenses.

This process requires you to contact IBM Smart Analytics System Support for transfer of license from the failed switch to the replacement

To recover a SAN switch, you must:

1. Determine the backup configuration details
2. Set up a serial or point-to-point network to connect to the SAN switch
3. Change the default passwords on the switch
4. Add the SAN switch to the cluster
5. Restore the configuration from a backup
6. Activate the ports on the SAN switch
7. Verify the recovery

Determine the backup configuration details

Retrieve the most recent configuration file backup to view the WWN (boot_licid) and the firmware level of the failed SAN switch.

Set up a serial or point-to-point network

Set up a temporary point-to-point network between a laptop and the SAN switch to gain access to the CLI on the device, which you can then be configured and added to the cluster network. You might need crossover Ethernet cable if the network adapters do not automatically provide the receive and transmit crossover, although most 1GB adapters provide the crossover automatically.

To set up a temporary point-to-point network:

1. Set the laptop network adapter to an address with the same subnet as that of the SAN switch, with a mask of 255.255.255.0. The default address for the SAN switch is 10.77.77.77, so if you keep the default address, you could set the laptop network adapter to an address such as 10.77.77.1.
2. Test the network connection by using the ping utility to connect to the default IP address of the SAN switch. Contact IBM Smart Analytics System Support if the connection fails.
3. To verify that you can access the switch, log in to the SAN switch using the SSH utility as the admin user id.

For further details on the SAN40B switch, see the manual referenced in the Further reading section of this document.

Change default passwords on the switch

Change and confirm the password for admin, root, factory and other accounts on the new switch to match those used for the other switches by performing the following steps:

1. Log into the SAN switch using the temporary point-to-point network.
2. Change the default passwords on initial login when prompted.
3. Verify the new passwords by exiting and logging back in to the switch.

Add the SAN switch to the cluster

To connect the switch to the management network:

1. On the switch administration window, click the **Network** tab.
2. Configure and apply the network settings.
3. Physically connect the switch to the cluster.

Restore the configuration from a backup

To restore the configuration, issue the `configdownload` command on the SAN switch using the SCP protocol to copy the configuration backup. An example follows:

```
ssh -t admin@172.22.1.113 "switchdisable”
ssh -t admin@172.22.1.113 "configdownload -all -p scp
\"172.22.1.10\",\"root\",\"/backups/san_cfg/san_sw_config.txt\""
ssh -t admin@172.22.1.113 "switchenable”
```

Activate ports on the SAN switch

Ports 0 and 24 through 39 are activated through extra licensing. To activate these ports:

1. Consult IBM Smart Analytics System Support to have them transfer the switch licenses from the old SAN switch to the new SAN switch.
2. Obtain the new license keys and instructions by going to: http://www-912.ibm.com/FruLicenseRequestClient/

The following tables show examples of the information on that website:

Switch identification

New Field Replaceable Unit World-Wide Name	Original Switch World-Wide Name
10:00:00:05:1E:BB:76:0B	10:00:00:05:1E:7D:6E:D5

License key information

Activation key	Feature name
`YgWBQHtB943TYYCfPJTtCmTQYmKQJ3CSB7rPH`	Enhanced Group Management
`RRHY3WDGTm9fKHAf4KM9SrSrtYaPNGArBJ9NL`	8 GB License
`yS9ccQyR91TAddf`	Fabric Watch
`R9QSReycczSARTSg`	Full Fabric
`yS9ccQyR9zjAddt`	Ports on demand, 8 incremental
`yS9ccQyR9zTEddh`	Ports on demand, 8 incremental

3. SSH to the SAN switch as the admin user, as shown in the following example:

   ```
   # ssh 172.22.1.113 -l admin
   ```

4. Issue the `licenceAdd` command using the new activation key, as shown in the following example:

   ```
   San02:FID128:admin> licenceAdd "yS9ccQyR9zjAddt"
   ```

5. Issue the `portenable` command for port 0 and ports 24 through 39, as follows:

   ```
   San02:FID128:admin> portenable 0/24-39
   ```

6. Verify the configuration using the `licenceShow` and `switchShow` commands.

7. Using the CLI, verify that the zoning rules were loaded and are active by issuing the `cfgShow` command, as follows:

   ```
   cfgShow
   ```

IBM System Storage DS5300 Storage Controller

The IBM System Storage DS5300 contains two storage controllers that manage the storage subsystem. Failure of a single controller causes load to be diverted to the remaining controller. Replacing a single failed storage controller does not require recovery from a configuration backup: the operational (primary) controller loads firmware and the storage configuration on the replacement controller. To ensure that you maintain the intended flow of air in the system, it is important not to remove a failed component until you are replacing it.

Do not attempt to recover from a failure of both storage controllers without first contacting IBM Smart Analytics System Support.

In the unlikely scenario where both controllers fail and must be replaced, the configuration can be recovered from a backup file. Do not attempt to replace both controllers without first contacting IBM Smart Analytics System Support because loss of data might occur.

A configuration diagram of the storage subsystem is in the installation record and the network configuration is listed in the `xcluster.cfg` file; reference these during recovery.

Backing up a DS5300 storage controller

To back up the storage controllers, take a copy of the configuration, profile and support data files on both controllers. Use the storage manager command line interface utility (SMcli) to generate these files. The SMcli utility is accessed interactively or by issuing the `SMcli` command with the appropriate parameters.

There are three files to backup:

- Storage controller configuration
- Subsystem profile
- Supporting data

Before backing up for the first time, verify the IP addresses of the storage controllers. On the management server, use the `SMcli` utility:

```
SMcli –d
```

Refer to the installation record and `xcluster.cfg` file for details on the storage controllers that are installed in the cluster to verify the IP network addresses.

Storage controller configuration

This file is a script that is executed on the storage system to automatically configure arrays and virtual disks, and also host mapping, partitioning, and configuration parameters values. The following example shows how to issue the `SMcli` command to generate a copy of the configuration file and save it to the appropriate backup directory:

```
SMcli 172.22.1.130 172.22.1.131 –c 'save storageSubsystem configuration
file="/backups/ds5300_cfg/ds5300_1.cfg" allConfig;'
```

Storage subsystem profile

This file is human readable text that describes the storage system contents, attributes, settings, and configuration. The following example shows how to use the **SMcli** command to generate this file:

```
SMcli 172.22.1.130 172.22.1.131 -c 'show storageSubsystem profile'
>/backups/ds5300_cfg/ds5300_1_profile.txt
```

Supporting data

Information, including configuration, profile, error logs, and counters, is collected in a single binary file which can be provided to IBM Smart Analytics System Support if requested. The process takes 5 or 10 minutes to complete and the output file is approximately 1MB. The following example shows how to use the **SMcli** command to generate this file:

```
SMcli 172.22.1.130 172.22.1.131 -c 'save storageSubsystem supportData
file="/backups/ds5300_cfg/ds5300_1_supportdata.bin" ;'
```

Recovering a DS5300 storage controller

To recover, edit the backup configuration file before using the command line interface to restore the configuration to the storage controllers. The configuration file for a DS5300 contains a list of commands that are downloaded and executed on the storage controller. Perform this procedure only on a new storage system or under direction from IBM Smart Analytics System Support because data loss can occur.

As a prerequisite to recovering the configuration, establish a point-to-point network as specified in the installation guide and configure the storage controller to have the same IP address as the failed controller. Refer to the storage controller documentation for details on the default IP address and password for the controller.

Complete the following steps to recover the storage controller configuration:

1. Prepare the backup file for recovery
2. Initiate the recovery process
3. Verify the recovery

Do not remove a failed storage controller until you are ready to replace it as airflow will be adversely affected.

Prepare the backup file for recovery

Check the following areas of the backup script before initiating the recovery process:

1. To avoid clearing the disks and re-creating the logical unit numbers (LUNs) when restoring the storage controller configuration, comment out the following line:

```
//clear storagesubsystem configuration;
```

2. If you are replacing the controller but keeping the existing LUN and RAID arrays, comment out or remove the configuration lines that create the logical drives. These lines begin with `create logicaldrive`, as shown in the following example for the test cluster:

   ```
   //create logicaldrive diskDrives=(11,1 11,2 11,3 11,4 11,5 11,6 11,7 11,8)
   raidLevel=5 userLabel="LUN01_P7ADM01_1952" arrayUserLabel="ARRAY1" owner=A
   segmentSize=256 capacity=2095944040448 Bytes dssPreAllocate=true
   securityType=none;
   ```

3. Move the following lines, which are in the `cacheFlush` section, to the bottom of the file. This ensures that the LUNS are defined before the script is executed and might avoid script errors:

   ```
   show "Setting the Storage Subsystem cache block size to 16.";
   set storagesubsystem cacheBlockSize=16;

   show "Setting the Storage Subsystem to begin cache flush at 40% full.";
   set storagesubsystem cacheFlushStart=40;

   show "Setting the Storage Subsystem to end cache flush at 40% full.";
   set storagesubsystem cacheFlushStop=40;
   ```

4. Depending on your software version, this script might contain the command parameter `diskDrive` rather than `diskDrives,` which is a syntax error. Correct this error by editing the script file, as shown in the last line of the following file excerpt:

   ```
   // Copies the hot spare settings
   // NOTE: These statements are wrapped in on-error continue and on-error stop
   statements to
   // account for minor differences in capacity from the diskDrive of the
   Storage Subsystem on which // the configuration was saved to that of the
   drives on which the configuration will be copied.
   //on error continue;
   show "Creating hot spare at Enclosure 42 Slot 2.";
   set diskDrive[42,2] hotSpare=true;
   ```

Initiate the recovery process

To download the backup configuration script to the storage controller and execute it, issue the following command on the management server from the directory containing the storage controller configuration file. Specify the correct IP address, filename, and password for the storage controller that you are configuring

```
SMcli 172.22.1.151 -f ds5300_1.cfg -p password
```

You can monitor progress by using the DS5300 storage manager user interface on the HMC.

Verify recovery

Verify that the recovery process has completed successfully by performing the shutdown and startup sequence for the cluster, see the IBM Smart Analytics System User Guide for further details. Perform a backup of the new controller after successful replacement.

Juniper EX4200-48T Ethernet Switch

The Juniper Ethernet switches provide the network infrastructure for the IBM Smart Analytics System and are configured in a set of two or more where one switch in each set is the master. The configuration is maintained as a whole for the set of switches. The Juniper Ethernet switches are in two sets: one for the DB2 FCM network, and one for the management network.

For assistance with default passwords, refer to the manufacturer's documentation or contact IBM Smart Analytics System Support. For details on how to setup passwordless SSH access to the Juniper and SAN switches if not already configured, see Appendix B.

The network configuration is listed in the xcluster.cfg file; reference this during recovery.

Reference the xcluster.cfg configuration file to verify the IP addresses of the Juniper Ethernet switches and other network devices.

Backing up a Juniper EX4200-48T Ethernet Switch

To back up, the `scp` utility is used to take copy of the configuration file on each switch.

To back up a Juniper switch, you must:

- Determine the IP address of the Juniper Switch
- Back up the configuration file

Determine the IP address of the Juniper Switch

Query the `xcluster.cfg` file to find the IP address of the Juniper switch. Issue the following command on the management server:

```
grep NET_SWITCH /BCU_share/SAN/xcluster.cfg
```

The output on the test cluster was:

```
NET_SWITCH1 = 172.22.1.101
```

Back up the configuration file

Create a backup copy of the Juniper switch configuration by issuing the following command:

```
scp root@172.22.1.101:/config/juniper.conf.gz /backups/juniper_cfg/
```

where:

- root@172.22.1.101 refers to the IP address of the Juniper switch
- juniper.conf.gz is the configuration file being copied
- /backups/juniper_cfg refers to the backup directory on the management server

Recovering a Juniper EX4200-48T Ethernet Switch

To recover, the CLI is used to load the backup configuration file to the replacement switch.

The replacement Juniper switch is configured using the following steps:

1. Establish a direct connection to the Juniper switch
2. Use the ezsetup wizard to connect the switch to the management network on the cluster.
3. Restore the configuration of the Juniper switch

The following subsections provide more information about the steps. The network configuration is described in the `xcluster.cfg` file; reference this file during recovery.

Establish a direct connection to the Juniper switch

Configure the Juniper switch using the IP address of the failed switch through a direct connection between a laptop and the switch. For step-by-step details on how to establish a direct connection, refer to the manufacturers' documentation.

Connect the switch to the management network on the cluster.

To connect the switch to the management network:

1. Log in to the Juniper switch as the root user, and enter **`ezsetup`** at the command prompt.
2. Use the ezsetup wizard to perform a basic configuration of the Juniper switch, as follows:
   - Set the host name of the Juniper switch to that of the replacement device
   - Set the root password
   - Select **yes** for **Enable SSH service**
   - Select option 2**, Configure out-of-band management** [vme.o]
   - Set the management IP address (172.22.1.10 in the test cluster)
   - Enter the gateway IP address (172.22.1.101 in the test cluster)
   - Select **No** for **Configure SNMP**.
   - Set the time zone as required
3. Commit the changes and exit.

Restore the configuration of the Juniper switch

The ezsetup wizard performs a basic configuration of the Juniper switch to make it available on the management network. To reapply the configuration of the failed Juniper switch you will need to restore and execute the backup configuration file as follows:

- Log in to the Juniper switch as root

```
ssh p7juniper02
```

- Navigate to the `tmp` directory

```
cd /var/tmp
```

- Use secure copy to retrieve the backup configuration file from the management server

```
scp 172.22.1.101:/backups/juniper_cfg/juniper02.conf.gz .
```

- Access the Juniper switch command line interface

  ```
  cli
  ```

- Enter configuration mode on the Juniper switch

  ```
  configure
  ```

- Load the backup configuration file, overriding the existing configuration

  ```
  load override /var/tmp/juniper02.conf.gz
  ```

- Commit the changes

  ```
  commit
  ```

- Exit the CLI

  ```
  exit
  ```

- Verify that the `vlan` and HMC entries exist

  ```
  show configuration values
  ```

IBM Remote Support Manager for Storage

The IBM Remote Support Manager for Storage (RSM for Storage) is a dedicated server that provides problem reporting information and remote access for IBM Smart Analytics System Support for the IBM DS5300 storage system. Remote access is disabled by default and should only be enabled for the duration of specific operations such as backup. Perform the backup process, done manually, when changing the configuration of the RSM for Storage server.

The RSM server has two hard disk drives (HDD) with mirroring, so you can recover from a single HDD failure using standard system x disk replacement procedures without loss of content. A double (complete) failure requires recovery.

To back up, the `rsm-cfg-save` command is used to create a copy of the configuration of the server.

For RSM for Storage information and software downloads, go to http://www-03.ibm.com/systems/storage/disk/rsm/.

Backing up the RSM for Storage

The RSM configuration data is backed up to file through a console window using the following steps:

1. Enable remote access to the RSM for Storage server through the browser interface by navigating to the IP address of the RSM server and clicking on the **Enable Remote Access** button.

2. Log on to the RSM for Storage server as the root, using the SSH utility through the port designated for SSH access, and open a console window.

3. Issue the `rsm` command to save a copy of the configuration into a compressed file:

   ```
   rsm-cfg-save
   ```

4. The saved configuration file is stored here `/packages/rsm-configuration-files.zip`. Using SCP, copy the file to the management server:

   ```
   scp –P 23 172.22.1.140:/packages/rsm-config-files.zip /backups/rsm_cfg/.
   ```

5. When the backup is copied to the management server, disable remote access to the RSM for Storage server through the browser interface by navigating to the IP address of the RSM server and clicking on the **Disable Remove Access** button.

Recovering the RSM for Storage

To recover, use the HDD recovery DVD to re-install the default installation, and then restore the configuration using the backup file.

A diagram depicting how the RSM server is installed is in the installation record; reference this during recovery.

1. To recover the RSM for Storage server:

2. Recover the default installation by using the recovery DVD as described the RSM for Storage documentation.

3. Configure the network settings to match the original configuration.

4. If required, download and re-install the latest version of the RSM software.

5. The original RSM configuration can then be restored from the backed up file. Retrieve the file from TSM if it is not already on `/backups` file system on the management server and copy to the RSM server before re-applying the configuration.

6. Open a browser and enable remote access to the RSM for Storage server through the graphical user interface.

7. Log in to the RSM for Storage as the root user and open a console window.

8. Place the `rsm-configuration-files.zip` in the `/packages` directory.

9. To restore the configuration, enter the following command:

   ```
   rsm-cfg-restore
   ```

10. To stop RSM for Storage, enter the following command:

    ```
    rsm-stop
    ```

11. To start RSM for Storage, enter the following command:

    ```
    rsm-start
    ```

Conclusion

Implementing a backup and recovery strategy for the operating system and configured devices of the IBM Smart Analytics System can significantly reduce downtime in the event of server or device failure. Centralizing the storage of backup files on the management server simplifies the administration of the backup and recovery process. Incorporate the recommendations in this paper into the overall backup strategy for your IBM Smart Analytics System.

Summary of key recommendations

The following list summarizes the key recommendations made in this best practices paper:

- Align the frequency of the backups with the frequency of configuration changes for each device or server.
- Use TSM to manage the storage and retrieval of backup files for all backup and recovery processes.
- Document and test each recovery process to help ensure successful recovery from server or device failure.
- Save and maintain a copy of the installation record and `xcluster.cfg` configuration file for your IBM Smart Analytics System.
- Configure available disk capacity on the storage system to hold backup files centrally on the management server.
- Burn DVDs for recovery only when recovery is necessary.
- Use backup file names that make it easy to identify the origin and content of the files.
- Back up files through the management network to avoid conflict with the Fast Communications Manager (FCM) network.
- Set up passwordless SSH access to enable single command backup tasks for easier automation.
- Do not attempt to recover from a failure of both storage controllers without first contacting IBM Smart Analytics System Support.
- Do not remove a failed storage controller until you are ready to replace it as airflow will be adversely affected.
- Reference the xcluster.cfg configuration file to verify the IP addresses of the Juniper Ethernet switches and other network devices.

Appendix A. Configuration of test system used

The test system that was used for this paper was an IBM Smart Analytics System 7600 R1 configured as follows:

Module	Management server host name	Management network IP address	FCM host	FCM network IP address	Managed system name	NIM client
HMC	e-hmc01	172.22.1.204				
Management	bluejay01mgt	172.22.1.10	bluejay01	172.23.1.10	BCUMGR	
Administration	bluejay02mgt	172.22.1.11	bluejay02	172.23.1.11	ADMINNODE	adminnode
Data 1	bluejay03mgt	172.22.1.21	bluejay03	172.23.1.21	DATANODE_1	datanode_1
Data 2	bluejay04mgt	172.22.1.22	bluejay04	172.23.1.22	DATANODE_2	datanode_2
Standby	bluejay05mgt	172.22.1.23	bluejay05	172.23.1.23	STANDBY_1	stdbynode_1

The devices configured on the cluster and included in this paper were as follows:

Device name	IP address
SAN40B switch 1	172.22.1.113
SAN40B switch 2	172.22.1.114
DS5300 storage controller A	172.22.1.130
DS5300 storage controller B	172.22.1.131
Juniper switch	172.22.1.101

Appendix B. Enable passwordless SSH access for the Juniper and SAN40B switches

Performing single-command backups from the management server requires passwordless SSH access to the Juniper switch and SAN40B switch. This also facilitates scripting and automation of the backup process.

To enable passwordless SSH access:

- Add the management server public key to each Juniper switch
- Add the management server public key to each SAN switch
- Add each SAN switch public key to the management server

Add the management server public key to each Juniper switch

The Juniper switch must have the public key for the management server to be able to issue a backup command to the Juniper switch through SSH.

To add the management server key to each Juniper switch:

1. Retrieve a copy of the management server public key:

   a. Log in to the management server as root
   b. Navigate to the `/.ssh` directory. The public key is the contents of the `id_rsa.pub` file in the `/.ssh` directory

2. Add the public-key to the Juniper switch:

   a. Log in to the Juniper switch as root from the management server: ssh `root@172.22.1.101`
   b. `enter cli mode -> cli`
   c. `enter configure mode -> configure`

3. Add authentication for root on the management server by using the `set` command, specifying the IP address and the public key the public key for the management server in quotation marks.

   In the following example, 172.22.1.10 is the IP address of the management server and the value in of the `id_rsa.pub` file appears in quotation marks:

   ```
   set system root-authentication ssh-rsa “ssh-rsa
   AAAAB3NzaC1yc2EAAAABIwAAAQEAwD+KNt2eLFEUo96HFJ21MaWkDqJ6/wnZtkJkx7FyQSCAU6Ae
   IjcVPKg6SeGOpOwkalHOmeAJg5hseWDy7gg8HbQsSiFavv+sYorQ== root@bluejay01” from
   172.22.1.10
   ```

4. Enter the commit command and then exit the Juniper switch CLI:

   ```
   commit
   exit / exit / exit
   ```

Add the management server public key to each SAN switch

For the management server to issue a backup command to the SAN switch through SSH the SAN switch must have the public key for the management server.

To add the management server public key to each SAN switch:

1. Identify the IP address of each SAN switch from the management server by entering the following command:

   ```
   grep SAN /BCU_share/SAN/xcluster.cfg
   ```

2. Log in to the chosen SAN switch from the management server by entering the following command. The IP address for the test cluster is shown in the following example:

   ```
   ssh admin@172.22.1.200
   ```

3. From the admin prompt, start the `sshutil` utility to add the public key of the management server by entering the following command:

   ```
   sshutil importpubkey
   ```

4. When prompted enter the IP address of the management server:

   ```
   Enter IP address: 172.22.1.10
   ```

5. When prompted enter the directory where the public key exists on the management server.

   ```
   Enter remote directory: /.ssh
   ```

6. When prompted enter the login name to register for SSH access.
7. When prompted enter the password for root on the management server. The public key will now be imported.
8. Exit the `SMcli` utility.

Add each SAN switch to the management server

The `configupload` SAN switch command to upload the configuration of the SAN switch to a file on the management server. To create the backup file on the management server, the command must connect to the management server. To enable the command to connect without prompting you for a password, you must add the SAN switch public key to the management server.

To add the SAN switch public key to the management server, perform the following steps. The IP address for the test cluster is used in the following example:

1. Log in to the SAN switch as user admin:

   ```
   ssh admin@172.22.1.200
   ```

2. From the admin prompt start the sshutil genkey utility:

   ```
   sshutil genkey
   ```

3. When prompted for a passphrase, press Enter:
4. When prompted to repeat the passphrase, press Enter. The process will complete. The `sshutil genkey` utility generates a public and a private key for the SAN switch.

5. Export he public key to the `.ssh` directory on the management server:

   a. Log in to the SAN switch as user admin:

   ```
   ssh admin@172.22.1.200
   ```

   b. From the admin prompt start the `sshutil exportpubkey` utility:

   ```
   sshutil exportpubkey
   ```

   c. When prompted for the IP address, enter the IP address of the management server:

   ```
   Enter IP address: 172.22.1.10
   ```

   d. When prompted, enter the remote directory:

   ```
   Enter remote directory: /.ssh
   ```

   e. When prompted for a login name on the management server enter `root`:

   f. When prompted enter the password for root on the management server. The `/.ssh/out_going.pub` file, which contains the public key for the SAN switch, is created on the management server.

6. Log in to the management server as root.

7. Take a backup copy of keys file before appending the new key to the file:

   ```
   cp /.ssh/authorized_keys /.ssh/authorized_keys.old
   ```

8. Add the SAN switch key to the `authorized_keys` file using the following command:

   ```
   cat /.ssh/out_going.pub >> /.ssh/authorized_keys
   ```

Appendix C. Creating backup file system on the management server

The `/backups` file system on the management server uses spare capacity on the DS5300 storage system. Use the file system to centralize the storage of backup images and other files from every server and device in the IBM Smart Analytics System. Use TSM to archive these files off the management server.

Use the storage manager command-line interface utility `SMcli` to issue commands from the management server to query and configure the storage subsystem. You must be the root user to issue these commands. The `SMcli` and `smcli` utilities are different; be sure to use the `SMcli` command.

To set up the file system, perform the following processes:

- Identify the DS5300 storage system to configure
- Confirm that there is disk space available on the storage subsystem
- Create a logical drive for the available disk space
- Map the new logical drive to a host group
- Configure the logical device on the management server
- Make the disk space available to the management server
- Configure NFS file system on all servers

Identify the DS5300 storage system

The `SMcli` command requires the storage system name or IP address as a parameter by when configuring the file system. To identify the storage manager name and IP addresses (for redundancy, there are two controllers) issue the following command from the management server:

```
SMcli –d
```

On the test system, the command displayed the following name and IP addresses:

```
E-Class-SAN     172.22.1.130    172.22.1.131
```

Confirm that there is disk space available on the storage subsystem

The IBM Smart Analytics System has unused disk capacity on the storage system. To confirm this, issue as root from the management server command prompt, issue the **`show array`** command. In the following example, `172.22.1.130 and` 172.22.1.131 are the IP address of the storage system and `ARRAY3` is the array on which there is spare disk capacity available:

```
SMcli 172.22.1.130 172.22.1.131 –c 'show array ["ARRAY3"];'
```

For the test system, the following output from the command shows 253.593 GB of free capacity and the primary storage controller (owner) for ARRAY3:

```
Performing syntax check...
Syntax check complete.
Executing script...
Name:                         ARRAY3
```

```
Status:                     Optimal
Capacity                    953.593 GB
RAID level:                 5
Media type:                 Hard Disk Drive
Interface type:             Fibre channel
Enclosure loss protection:  No
Security Capable:           No
Secure:                     No
Current owner:              Controller in slot B
Associated logical drives and free capacity
   Logical Drive         Capacity
   LUN03              400.000 GB
   LUN04              150.000 GB
   LUN05              150.000 GB
   Free Capacity:     253.593 GB
```

Create a logical drive for the available disk space

Use the spare disk capacity by creating a logical drive, which is referred to as a logical unit number (LUN). To create a logical drive, from the management server, issue the **`create logicalDrive`** command. The following command creates a logical drive on the test cluster:

```
SMcli 172.22.1.130 172.22.1.131 -c 'create logicalDrive array=ARRAY3
userLabel="BKUP_LUN_250" owner=B segmentSize=256 capacity=268435456000;' -p
<passwd>
```

where:

- `userLabel` specifies is a unique name to be assigned to the logical drive
- `owner` specifies the controller that is deemed the primary or current owner (see the output of the previous section)
- `segmentSize` specifies the segment size; 256 which is consistent with other logical drives
- `capacity` specifies the disk capacity; 250 GB is calculated as 250 x 1024 x 1024 x 1024. If you have larger drives, then your capacity will be greater.
- `passwd` specifies the password for the storage subsystem

Map the new logical drive to a host group

Mapping the new logical drive to a host group makes the drive visible to the management server for configuration.

To map the logical drive:

1. Identify the host group for the management server by using the **`show storageSubsystem`** hostTopology command. The following command identifies the host group for the test cluster:

```
SMcli 172.22.1.130 172.22.1.131 –c 'show storageSubsystem hostTopology;’
```

The following output shows GROUP1 as the host group for the management server which is BLUEJAY01:

```
   TOPOLOGY DEFINITIONS
      STORAGE SUBSYSTEM
         Default type:                    AIX
         Default Group
         Host Group:                      GROUP1
            Host:                         BLUEJAY01
               Host type:                 AIX
               Interface type:            Fibre Channel
                  Host port identifier:  10:00:00:00:c9:89:cd:81
                  Alias:                  BLUEJAY01_fcs1
                  Host port identifier:  10:00:00:00:c9:89:bc:61
                  Alias:                  BLUEJAY01_fcs3
                  Host port identifier:  10:00:00:00:c9:89:cd:80
                  Alias:                  BLUEJAY01_fcs0
                  Host port identifier:  10:00:00:00:c9:89:bc:60
                  Alias:                  BLUEJAY01_fcs2
         Host Group:                      GROUP2
            Host:                         BLUEJAY02
               Host type:                 AIX
               Interface type:            Fibre Channel
                  Host port identifier:  10:00:00:00:c9:89:c4:fb
                  Alias:                  BLUEJAY02_fcs3
                  Host port identifier:  10:00:00:00:c9:89:c5:00
                  Alias:                  BLUEJAY02_fcs0
                  Host port identifier:  10:00:00:00:c9:89:c4:fa
                  Alias:                  BLUEJAY02_fcs2
                  Host port identifier:  10:00:00:00:c9:89:c5:01
                  Alias:                  BLUEJAY02_fcs1
```

2. Determine the next available LUN by issuing the **`show storageSubsystem lunMappings`** command; a unique LUN is required to map the logical drive to the host group. The following command was used on the test cluster:

```
SMcli 172.22.1.130 172.22.1.131 –c 'show storageSubsystem lunMappings;'
```

The following output from the command shows the host group as GROUP1 and the next available LUN as 18 (the last LUN shown is LUN17):

```
MAPPINGS (Storage Partitioning - Enabled (2 of 8 used))-------------------
   Logical Drive Name    LUN  Controller  Accessible by      Logical Drive
status
   Access Logical Drive  31   A,B         Host Group GROUP1  Optimal
```

```
LUN03                 3    B           Host Group GROUP1  Optimal
LUN04                 4    B           Host Group GROUP1  Optimal
LUN05                 5    B           Host Group GROUP1  Optimal
Access Logical Drive  31   A,B         Host Group GROUP2  Optimal
LUN01                 1    A           Host Group GROUP2  Optimal
LUN02                 2    A           Host Group GROUP2  Optimal
LUN06                 6    B           Host Group GROUP2  Optimal
LUN07                 7    B           Host Group GROUP2  Optimal
LUN08                 8    B           Host Group GROUP2  Optimal
LUN09                 9    B           Host Group GROUP2  Optimal
LUN10                 10   A           Host Group GROUP2  Optimal
LUN11                 11   A           Host Group GROUP2  Optimal
LUN12                 12   B           Host Group GROUP2  Optimal
LUN13                 13   B           Host Group GROUP2  Optimal
LUN14                 14   A           Host Group GROUP2  Optimal
LUN15                 15   A           Host Group GROUP2  Optimal
LUN16                 16   B           Host Group GROUP2  Optimal
LUN17                 17   B           Host Group GROUP2  Optimal
Access Logical Drive  31   A,B         Storage Subsystem  Optimal
```

3. Map the logical drive to the host group by issuing the `set logicalDrive` command. The following command maps the logical drive to the host group on the test server:

```
SMcli 172.22.1.130 172.22.1.130 -c 'set logicalDrive ["BKUP_LUN_250"]
logicalUnitNumber=18 hostgroup="GROUP1";'  -p <passwd>
```

where:

- `logicalUnitNumber` specifies the next available LUN
- `hostgroup` specifies the host group on which the array with disk capacity is available
- `passwd` is the password for the storage subsystem

4. Issue the show `storageSubsystem lunMappings` command again to verify that the logical drive was mapped to the host group.

Configure the logical device on the management server

To configure the logical device:

1. Issue the following `lsdev` (list devices) command to show the hard disk devices that are currently configured:

```
lsdev | grep hdisk
```

The following output shows the configured devices on the test server:

```
hdisk0      Available 00-08-00 SAS Disk Drive
hdisk1      Available 00-08-00 SAS Disk Drive
hdisk2      Available 00-08-00 SAS Disk Drive
hdisk3      Available 00-08-00 SAS Disk Drive
hdisk4      Available 04-00-02 MPIO DS5100/5300 Disk
hdisk5      Available 04-00-02 MPIO DS5100/5300 Disk
hdisk6      Available 04-00-02 MPIO DS5100/5300 Disk
```

2. Issue the `cfgmgr` command as root on the management server. This utility discovers and configures logical drives:

```
cfgmgr -v
```

3. Show the hard disk devices that are now configured by issuing the `lsdev` command again, as follows:

```
lsdev | grep hdisk
```

The output for the test system shows that two additional disks have been configured:

```
hdisk0      Available 00-08-00 SAS Disk Drive
hdisk1      Available 00-08-00 SAS Disk Drive
hdisk2      Available 00-08-00 SAS Disk Drive
hdisk3      Available 00-08-00 SAS Disk Drive
hdisk4      Available 04-00-02 MPIO DS5100/5300 Disk
hdisk5      Available 04-00-02 MPIO DS5100/5300 Disk
hdisk6      Available 04-00-02 MPIO DS5100/5300 Disk
hdisk7      Available 04-00-02 MPIO DS5100/5300 Disk
```

You can obtain further information by using the `mpio` utility. Issue the `mpio_get_config` command as follows:

```
mpio_get_config -Av
```

If `sddpcm` devices are used instead, issue:

```
sddpcm_get_config -Av
```

The mpio_get_config command produced the following output on the test system:

```
Frame id 0:
    Storage Subsystem worldwide name: 60ab80047ac4800004a717ca0
    Controller count: 2
    Partition count: 1
    Partition 0:
    Storage Subsystem Name = 'E-Class-SAN'
        hdisk      LUN #   Ownership          User Label
        hdisk4         3   B (preferred)      LUN03
        hdisk5         4   B (preferred)      LUN04
        hdisk6         5   B (preferred)      LUN05
        hdisk7        18   B (preferred)      BKP_LUN_250
```

Make the disk space available to the management server

To make the disk space available to the management server:

1. Create a volume group by issuing the `mkvg` command. The following command creates a volume group on the test system:

   ```
   mkvg -S -y vgbackups -s 64 -P 1024 -f -V 101 hdisk7
   ```

   where:

   - `vgbackups` is the name of the new volume group
   - `-s64` sets the number of megabytes per physical partition
   - `-P 1024` specifies the total number of partitions in the volume group
   - `-V 101` specifies the volume group major number which is constant across the cluster
   - `hdisk7` is the hard disk discovered by the `cfgmgr` command

2. Create the logical volume and assign to a volume group by issuing the `mklv` command. The following command creates a logical volume on the test system:

   ```
   mklv -t jfs2 -y lvbackups vgbackups 3875 hdisk7
   ```

   where:

   - `-t jfs2` specifies the logical volume type that is recommended
   - `-y lvbackups` specifies the name of the logical volume to create
   - `3875` is the number of physical partitions to allocate to the logical volume
   - `hdisk7` is the hdisk that was discovered by the `cfgmgr` command

3. Create the directory on which backups will be stored by issuing the following command:

   ```
   mkdir -p /backups
   ```

4. Create the file system and associate it with the backup directory by issuing the following command:

   ```
   crfs -v jfs2 -d lvbackups -m /backups -a log=INLINE -A yes -t no -p rw -u
   backups
   ```

   where:

   - `-a log=INLINE`. This specifies that the jsf2 volume is used for logging.
   - `-A yes` ensures that the file system is mounted at each system restart
   - `-t no` suppresses processing by an accounting subsystem
   - `-p rw` sets the permissions for the file system

5. Mount the new file system and make it available for use:

   ```
   mount /backups
   ```

6. Verify that the new file system exists by issuing the **`following`** command and inspecting the results:

   ```
   df -g | grep backups
   ```

The command produced the following output on the test cluster:

```
/dev/lvbackups    242.19    241.65    1%       17     1% /backups
```

Configure NFS file system on all servers

To export the `backup` directory to all NFS clients, from the management server, issue the `mknfsexp` command. In the following example, `bluejay02` is the administration server, `bluejay03` and `bluejay04` are data servers and `bluejay05` is a standby server. The HMC server in the test cluster is e-hmc01.

```
mknfsexp -d /backups -B -t rw -c
'bluejay02mgt, bluejay03mgt, bluejay04mgt, bluejay05mgt, e-hmc01 '
-r
'bluejay02mgt, bluejay03mgt, bluejay04mgt, bluejay05mgt '
```

On the test system, issuing the following command from the management server, `bluejay01mgt`, mounts the `/backups` NFS file system on all servers:

```
for i in bluejay02mgt bluejay03mgt bluejay04mgt
do
  ssh $i "mkdir -p /backups"
  ssh $i "mknfsmnt -f /backups -d /backups -h bluejay01mgt -B -A -t rw -w bg"
done
```

Further reading

IBM Version History of Validate Software Stack for IBM Smart Analytics System
http://www-01.ibm.com/support/docview.wss?rs=4106&uid=swg21437888

IBM Power 550 Express (8204–E8A)
The Systems Hardware Information Center contains information related to installing, managing, using and troubleshooting Power 550 servers:
http://publib.boulder.ibm.com/infocenter/systems/scope/hw/index.jsp?topic=/iphdx/550_e8a_landing.htm

IBM System Storage DS5300
Installation, User's and Maintenance Guide - IBM System Storage DS5100 and DS5300:
https://www-947.ibm.com/systems/support/supportsite.wss/docdisplay?lndocid=MIGR-5077659&brandind=5000028

Juniper EX4200–48T Ethernet switch
Data Center LAN Connectivity Design Guide
http://www.juniper.net/us/en/local/pdf/design-guides/8020010-en.pdf

IBM System Storage SAN40B-4 Switch
IBM System Storage SAN40B-4 Installation, Service, and User's Guide
http://www-01.ibm.com/support/docview.wss?rs=1308&context=STCLLES&dc=DA400&uid=ssg1S7002302&loc=en_US&cs=utf-8&lang=en

IBM Hardware Management Console
Hardware Management Console V7 Handbook
http://www.redbooks.ibm.com/redbooks/pdfs/sg247491.pdf

IBM Remote Support Manager
http://www-03.ibm.com/systems/storage/disk/rsm/

Network Installation Management
http://publib.boulder.ibm.com/infocenter/aix/v6r1/index.jsp?topic=/com.ibm.aix.install/doc/insgdrf/basic_config_cmd.htm

IBM Fix Central
http://www-933.ibm.com/support/fixcentral/

IBM Smart Analytics System Support
http://www-947.ibm.com/support/entry/portal/Overview/Software/Information_Management/IBM_Smart_Analytics_System

IBM Tivoli Storage Manager
http://publib.boulder.ibm.com/infocenter/tsminfo/v6/index.jsp?topic=/com.ibm.itsm.srv.doc/c_tsmintro.html/

IBM Developerworks Best Practices papers
http://www.ibm.com/developerworks/data/bestpractices/

Contributors

Darren Rogers

IBM System Verification Testing

Julie Craft

IBM Smart Analytics and AIX Development

Sermsak Sukjirawat

WW Data Warehouse and Smart Analytics System Practices

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

Without limiting the above disclaimers, IBM provides no representations or warranties regarding the accuracy, reliability or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The use of this information or the implementation of any recommendations or techniques herein is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Anyone attempting to adapt these techniques to their own environment do so at their own risk.

This document and the information contained herein may be used solely in connection with the IBM products discussed in this document.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

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

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Trademarks

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

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