



High Availability Automation for iSCSI Integrated Server and AIX 5L with Virtual Storage on System i™

James McCord
IASP, Clustering and External Storage on System i

Mike Schambureck
Integrated Server on System i

Vess Natchev
AIX 5L and SAN on System i

IBM Rochester, MN

October 15th, 2006

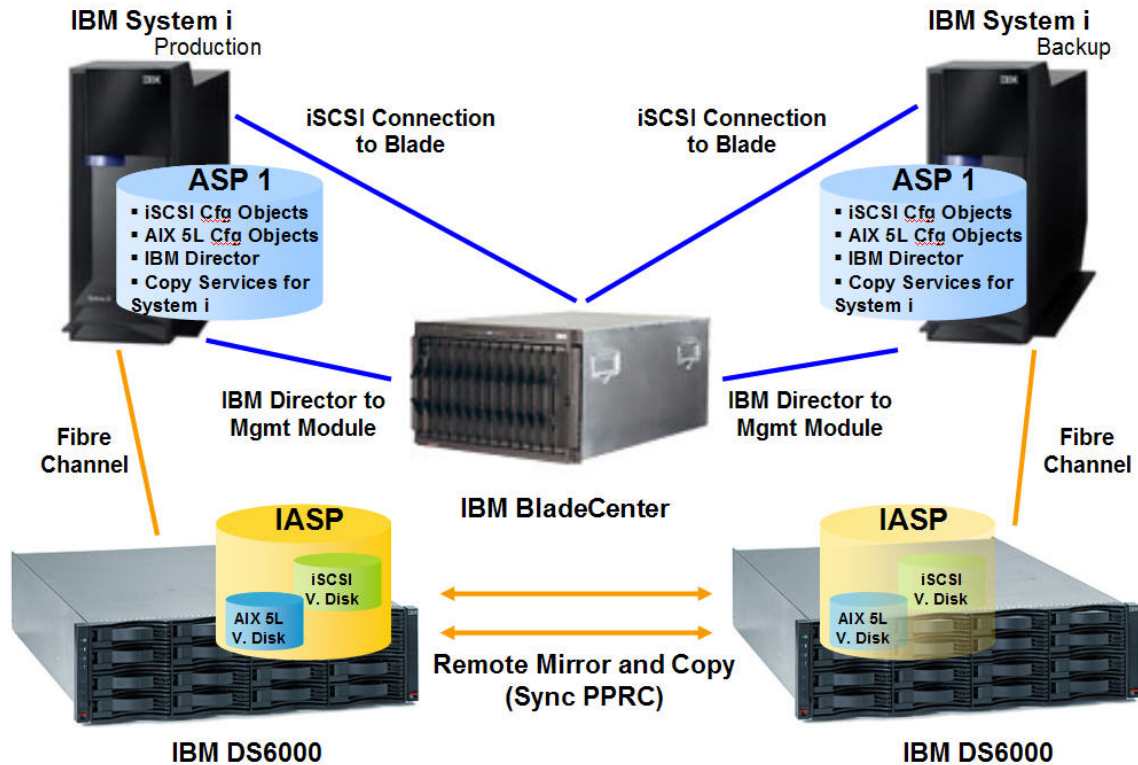


1.0 Overview: Automating High Availability for Integrated Workloads

High availability (HA) is a necessary component of a resilient IT infrastructure. Clients who wish to simplify their integrated server or UNIX environment by hosting it with i5/OS™ can further leverage existing System i™ capabilities to automate HA management for those environments. Integrating IBM BladeCenter™ and System x™ with Internet SCSI (iSCSI) or AIX 5L with i5/OS simplifies the IT environment by enabling rapid deployment of new virtual storage, while using existing skills to manage physical disk and adapter hardware. By creating virtual disk in an Independent Auxiliary Storage Pool (IASP), System i clients can incorporate a hosted environment into their core business HA strategy. i5/OS has the capability to automatically take action when an IASP becomes available or unavailable. An IASP can become unavailable on a production system and available on a backup system during an HA failover. If that IASP contains all the storage for an integrated server or AIX 5L workload, that integrated environment can be automatically stopped on the failing system and restarted on the backup system by i5/OS. This type of integrated HA environment can be further automated by using Copy Services for System i.

2.0 Statement of Purpose, Applicability and Audience

This paper will demonstrate how the combination of existing System i technologies and service offerings can provide additional value to clients by automating HA for an integrated server and/or AIX 5L environment hosted in an IASP. The example solution demonstrated here provides two of the core facets of HA: making business data available on a backup site and restarting the AIX 5L LPAR or integrated server on the backup site. This configuration can be used for both planned and unplanned outages. In both cases, the necessary switchover tasks are initiated with the Copy Services for System i interface, as described below. The switchover time is estimated at 15 to 20 minutes.



The i5/OS objects necessary to provide virtual storage to an integrated BladeCenter, System x or an AIX 5L logical partition reside in the System ASP. The virtual storage objects for both the integrated server and AIX 5L exist in the IASP. In this case, the IASP on both systems resides in an IBM System Storage DS6000. Data replication of the IASP is provided by Metro Mirror (formerly Synchronous PPRC, Peer-to-Peer Remote Copy), an IBM System Storage Remote Mirror and Copy (RMC) solution. Copy Services for System i provides i5/OS command-line management of Metro Mirror and HA switchover automation. The paper will present two example CL programs that further automate the switchover of an integrated environment.

Please note: this paper will not attempt to provide a comprehensive guide to HA on System i. Multiple storage, HA and replication alternatives exist. For example, both the Microsoft Windows server™ operation system and AIX 5L can also use direct SCSI or Fibre Channel storage. Microsoft Cluster Service™ can provide HA for Windows, while HACMP is the leading HA solution for AIX 5L. Virtual storage objects located in an IASP can also be replicated natively in i5/OS with Geographic Mirroring. While those are valid and proven solutions, the value of the approach discussed here lies in a single HA strategy for i5/OS, an iSCSI integrated server and an AIX 5L LPAR, using the same technology and leveraging existing System i features.

A note on POWER Linux on System i: this white paper discusses using existing capabilities of System i to create an integrated HA environment for i5/OS, an iSCSI integrated server and an AIX 5L LPAR. The same capabilities should also work for POWER Linux in an LPAR on the platform, since Linux storage virtualization is very similar to that for integrated servers and AIX 5L. However, at the time of the writing of this whitepaper, the ability to have a Linux partition participate in the HA scenario presented here has not been validated.

The intended audience for this paper is IBM Field Technical Sales Specialists, Systems Architects, administrators, support specialists, service technicians, IBM Business Partners, consultants, and other IT professionals who install, administer, and support a System i integrated server or AIX 5L server environment.

3.0 Technology Overview

This section will provide a technical overview of the components in this HA automation solution.

3.1 Virtual Storage, IASPs and Exit Programs

The capability to virtualize i5/OS physical storage as virtual disk for Windows and Linux has been available on System i for many releases. In i5/OS V5R3, that capability was enhanced to allow integration of UNIX workloads with virtual disk for AIX 5L. The benefits and implementation of i5/OS storage virtualization are covered extensively in several white papers and Redbooks, which can be found in the Appendix.

Similarly, the capability to store library objects in an independent database instance contained in a separate disk pool has been available since i5/OS V5R2. IASPs provide various levels of HA in many client environments in combination with other IBM technologies, such as i5/OS cluster support, switchable disk, Geographic Mirroring or Remote Mirror and Copy on IBM System Storage. The concepts and implementation of IASPs are the subjects of the several white papers and Redbooks, which can be found in the Appendix.

In the HA automation example in this paper, the IASP contains virtual storage for the integrated server and AIX 5L logical partition (LPAR). Synchronous replication of the IASP and switchover functions are managed by Copy Services for System i, while *exit points* and *exit programs* handle the automatic shutdown and start-up of both the integrated server and AIX 5L. With the concept of an exit point, i5/OS allows a user program to take action at a specific point in a system function. That user program is known as an exit program and is *registered* with the exit point. In this HA automation case, one CL program is registered with the exit point before an IASP is

varied off, while another is registered with the exit point after an IASP is varied on. Thus, these example CL programs demonstrate how to automatically and gracefully end an integrated environment before its virtual disks disappear on a production system during a switchover, and how to start the same environment on a backup system after its virtual disks have become available there.

3.2 Integrated BladeCenter and System x servers on the System i

Integrated BladeCenter and System x servers on the System i with iSCSI is new support with the V5R4 release of i5/OS. The support introduces an iSCSI connection from the integrated server to the System i. This same technology is used in the industry for access to Storage Area Networks (SANs). In this case, System i acts as the SAN. The servers contain one or more iSCSI host bus adapters (HBAs). That HBA, called the initiator, communicates with an HBA on System i, called the target. Both iSCSI and virtual Ethernet traffic use the same connection for disk and Ethernet data, respectively. The server's disk, HBA(s) and service processor are virtualized by System i5 objects that allow the System i to manage and control the server. The IBM Director server is used to check the status of the server and for power control of the server(s). All of the benefits of integrating servers (IXS/IXA) from the past continue with this new integration technology.

3.3 Copy Services for System i

The IBM Copy Services for System i is a software utility that facilitates the integration of IBM System Storage, Copy Services and the IBM System i. Due to the IBM System i single-level storage architecture and other considerations, there are strict rules and dependencies governing the use of Copy Services with the IBM System i. The steps required to control this environment are very critical and if performed incorrectly may result in an inoperable system, a lengthy recovery, or loss of data. Copy Services for System i is a software asset provided by the IBM STG (Systems and Technology Group) Lab Services organization and is packaged a part of a consulting services engagement.

Copy Services for System i leverages advanced IBM System i technologies including Clustering and IASPs to effectively support the use of Copy Services in an IBM System i environment. It hides the complexity of the overall solution by providing automated control and verification, System i5 HMC (Hardware Management Console) hardware resource manipulation, DSCLI (DS Command-line Interface) command processing, and clustering interface manipulation.

4.0 Test Environment

This section will describe the exact configuration of the HA automation test environment.

4.1 Hardware and Logical Partitioning

- Two System i model 520s, referred to as "production" and "backup."
- Each i520 contains one i5/OS hosting LPAR and one AIX 5L LPAR.
- The AIX 5L LPAR uses only virtual storage provided by i5/OS.
- The i5/OS LPAR is assigned an internal SCSI adapter and several disk drives for its System ASP, as well as a Fibre Channel adapter providing access to several logical units (LUNs) in a DS6000 for the IASP.
- The i5/OS hosting LPAR owns an iSCSI target (HBA) in order to provide virtual storage for a blade server, and an Ethernet adapter for communication with the BladeCenter Management Module (MM). This can be the same Ethernet adapter that is used to connect the System i to the intranet/internet.
- iSCSI HBAs come in two forms: initiators and targets. The adapters are similar in physical characteristics, but use different firmware. The initiator HBAs are located in the blade server or the System x server. They access their disk drives via the target HBAs located in the System i.

- One or more blade servers or System x servers: There are a number of models of each type supported. See the link off the integrated server web page in the appendix for more details. None of these servers contain local disk drives. Their disks are all virtual drives on the System i.
- 2 BladeCenter switch modules: Cisco and Dlink switches were used in this testing (other switches are supported). See the links in the appendix for more details on supported switches. If a System x server were being used, an external gigabit switch would be required.
- Two IBM DS6000 storage units, each connected to the production or backup i520
- The DS6000 controller firmware level used during this test was 5.2.0.594.
- 8.6-GB LUNs were used for the IASP on both DS6000s.

4.2 Software

- For the iSCSI environment, see the prerequisites listed in the read me first link in the appendix of this document.
- The AIX 5L LPAR ran AIX 5L v5.3 TL04. No additional licensed program products (LPPs) were installed for the purposes of this test.
- The i5/OS LPAR ran V5R4M0, and had Option 41, High Availability Switchable Resources, installed for the purpose of using a switchable IASP.
- The Copy Services for System i code.

4.3 iSCSI Objects

- Network Server Host Bus Adapter (NWSH) device descriptions: these are a new device type for V5R4 of i5/OS. They define the target HBA(s) installed in the System i.
- Network Server Configuration (NWSCFG) objects: These are also new for V5R4 and are used to virtualize the iSCSI and server hardware. There are 3 types available as follows:
 1. The Service Processor object (*SRVPRC) represents the System x service processor or the IBM BladeCenter management module. It is used by the IBM Director server for discovery, power control, checking status of the server(s) and issuing Windows shutdown commands.
 2. The Remote System object (*RMTSYS) defines the blade server or System x server and its HBA(s).
 3. The Connection Security object (*CNNSEC) can be used to secure the SCSI and virtual Ethernet data flows between i5/OS and the iSCSI attached xSeries or IBM BladeCenter server. This support is not currently there, so the object acts as a place holder for later implementation.
- Each NWSCFG object has a corresponding Validation List (VLDL) object of the same name. These are used to hold encryption information.
- Network Server Description (NWSD): defines to i5/OS which iSCSI target HBA and remote system (blade server or System x) to use. This object is automatically created when one runs the INSWNTSVR or INSLNXSVR command.
- Network Server Storage Spaces (NWSSTG): as part of either the INSWNTSVR or INSLNXSVR command, there are two storage spaces created, "system" and "install." The "system" drive is where the OS is installed. The "install" drive is used for System i integration code. Additional storage spaces can be created after the installation completes. They should be created out of the same IASP storage.
- A System i line description for the service processor connection from the System i is needed on the primary and on the secondary system. In most cases these objects already exist as the System i interface to the Internet or an intranet.

4.4 AIX 5L Objects

- A Virtual SCSI client adapter in the AIX 5L LPAR and a Virtual SCSI server adapter in the i5/OS LPAR, defined in the HMC on each i520.
- An NWSD called “AIX,” whose resource name (CTLxx) corresponds to the correct Virtual SCSI server adapter slot in i5/OS.
- A NWSSTG called “AIXHATEST” with a size of 2200 MB, linked to NWSD “AIX.”

4.5 IASP and Copy Services for System i Configuration

- An IASP called “DS6000,” which contained three 8.6-GB DS6000 LUNs.
- The vary off and on of the IASP is controlled on both production and backup systems with exit programs, which are external to Copy Services for System i. See Sections 4.6 and 4.7 for a description of the exit points and programs used.
- For hardware management and movement of the Fibre Channel adapters, Copy Services for System i made calls via SSH (Secure Shell) to the proper HMC to perform DLPAR (Dynamic Logical Partitioning) to add and remove the adapters as needed. Adequate time was given to allow the LUNs to properly report into the system. At that point the LUNS were correctly registered with System i Storage Management. If the LUNs do not properly report in, or Storage Management cannot correctly register them, the IASP vary on will fail.
- The DS CLI was used by Copy Services for System i to manipulate the LUN (volume) relationships on the DS6000s. Metro Mirror was established between LUN pairs on the two DS6000s. Once full duplex (or, a full synchronization of the data) was achieved, the DSCLI was used to control the direction of the replication, verify the volumes are in the correct state, and perform failover and failback scripts.
- Two Toolkit commands were used with the IBM System i Copy Services Toolkit:
 1. CHKPPRC (for “Check PPRC”) verified that the volume relationships were intact and data were being replicated in the correct way. CHKPPRC also verified that the hardware was in the proper state for DLPAR movement.
 2. SWPPRC (for “Switch PPRC”) was the Toolkit command that actually performed a switchover of the IASP. The switchover triggered the IASP vary off or vary on and activated the exit programs to automatically shut down or start the integrated environment.

4.6 DS6000 Hardware Replication

Metro Mirror is a function of a storage server that synchronously updates a secondary copy of a volume to match changes that i5/OS Storage Management makes to a primary volume. The primary and the secondary volumes can be on the same storage server or on separate storage servers. In the case of two storage units, the second one can be located at another site up to 300 Km away. However, there may be performance implications when using synchronous replication over this distance and it may be more practical to consider a shorter distance to minimize the performance impact.

Metro Mirror is application-independent. Because the replication occurs at the disk subsystem level in the storage unit, the application has no knowledge of its existence. The synchronous protocol guarantees that the secondary copy is up-to-date and consistent by ensuring that the primary copy will be committed only if the primary receives acknowledgment that the secondary copy has been written.

For this proof of technology, the two DS6000's were in the same room, connected to the i520s via Fibre Channel cables approximately 5 feet in length. It was not the intent of this proof of technology to benchmark performance. Therefore, no statements of performance are made or implied.

Note that while this test case used Metro Mirror between DS storage units, i5/OS can also provide replication with *Geographic Mirroring*. Geographic Mirroring is a function that keeps two identical copies of an IASP at two sites to provide high availability and disaster recovery. Geographic mirroring is logical mirroring, not physical mirroring. The two disk pools must have similar disk capacities, but the mirror copy may have different numbers and types of disk units as well as different types of disk protection. When geographic mirroring is active, changes to the production data are transmitted to the backup across TCP/IP connections. Changes can be transmitted either synchronously or asynchronously. Geographic mirroring is a sub-function of Cross-site Mirroring (XSM), which is part of i5/OS Option 41, High Availability Switchable Resources.

4.7 Exit Points and Registration

- The exit point used for taking action before the IASP is varied off was QIBM_QDC_VRYEXIT with **format PROF0100**. Exit points and their registered exit programs (if any) can be displayed with WRKREGINF. The sample CL program ENDINTENV was registered with this exit point on both systems with the following command: ADDEXITPGM EXITPNT(QIBM_QDC_VRYEXIT) FORMAT(PROF0100) PGMNBR(1) PGM(HAVIR/ENDINTENV) CRTEXITPNT(*YES) **PGMDTA(*JOB 8 DEVDASPD)**. Notice the “Program data” parameter: without the 8 bytes of data “DEVDASPD,” the exit program would not have been called when the IASP was varied off. “HAVIR” is the name of the library where the sample CL programs resided.
- The exit point used for taking action before the IASP is varied on was QIBM_QDC_VRYEXIT with **format PSON0200**. The sample CL program STRINTENV was registered with this exit point on both systems with the following command: ADDEXITPGM EXITPNT(QIBM_QDC_VRYEXIT) FORMAT(PSON0200) PGMNBR(1) PGM(HAVIR/ENDINTENV) CRTEXITPNT(*YES) **PGMDTA(*JOB 8 DEVDASPD)**.
 - Both exit programs were registered on both systems, so that they would automatically shut down and start the integrated environment, regardless of the direction of the switchover.

4.8 Data Areas

- Data area (DTAARA) “IASPNAME” contained the name of the IASP hosting the integrated server or AIX 5L LPAR virtual storage
- DTAARA “NWSA” contained the name of the NWSA managing the integrated server or AIX 5L environment
- DTAARA “NWSATYPE” contained the NWSA typed. It had one of two values:
 - “*ISCI” for an iSCSI-integrated server
 - “*AIX” for an AIX 5L LPAR
- DTAARA “NWSH” contained the name of the corresponding NWSH if the NWSA type was “*ISCI.”
- All data areas and both sample CL programs resided in library “HAVIR”, which was created for the test.

4.9 Sample CL Exit Programs

- The sample CL exit program ENDINTENV (for “End Integrated Environment”) handled the automatic and graceful shutdown of the integrated server and AIX 5L virtual storage environment *before the IASP was varied off on the production system*. ENDINTENV takes the following actions:
 - If any IASP begins a vary off, it receives information from the exit point about which object is being varied off.

- It checks if the object being varied off is a device description (that is, an IASP). It exits if the object is not an IASP.
- It compares the name of the IASP received from the exit point with a value from a data area created previously, which contains the name of the integrated environment IASP. If the IASP being varied off is not the one of interest, the program exits.
- It retrieves the name of the NWSD managing the integrated server or AIX 5L LPAR from another data area, and varies the NWSD off.
- It retrieves the type of the NWSD managing the integrated server or AIX 5L LPAR from a third data area. If the type is **not** *ISCSI, the program exits, because it has already taken all actions necessary to shut down the AIX 5L environment.
- If the NWSD type is *ISCSI, the program also shuts down the IBM Director server and varies off the corresponding NWSH, after it has retrieved its name from a fourth data area.
- For a flowchart showing the simple algorithm and the complete program, see the Appendix.
- The sample CL exit program STRINTENV (for “Start Integrated Environment”) handled the automatic and graceful start-up of the integrated server and AIX 5L virtual storage environment *after the IASP was varied on on the backup system*. STRINTENV takes the following actions:
 - After any IASP finishes a vary on, it receives information from the exit point about which object is now available.
 - It checks if the now available object is a device description (that is, an IASP). It exits if the object is not an IASP.
 - It compares the name of the IASP received from the exit point with a value from a data area created previously, which contains the name of the integrated environment IASP. If the now available IASP is not the one of interest, the program exits.
 - It retrieves the *type of the NWSD* managing the integrated server or AIX 5L LPAR from another area. If the type is **not** *ISCSI, the program retrieves the name of the NWSD managing the AIX 5L environment from a third data area, and varies the NWSD on. The program then exits, because it has already taken all actions necessary to start the AIX 5L environment.
 - If the NWSD type is *ISCSI, the program starts the IBM Director server and varies on the corresponding NWSH, after it has retrieved its name from a fourth data area.
 - The program then pauses for four minutes while it waits for IBM Director startup to complete.
 - It then retrieves the name of the NWSD managing the integrated server from its data area, and varies the NWSD on.
 - For a flowchart showing the simple algorithm and the complete program, see the Appendix.

5.0 Environment Setup

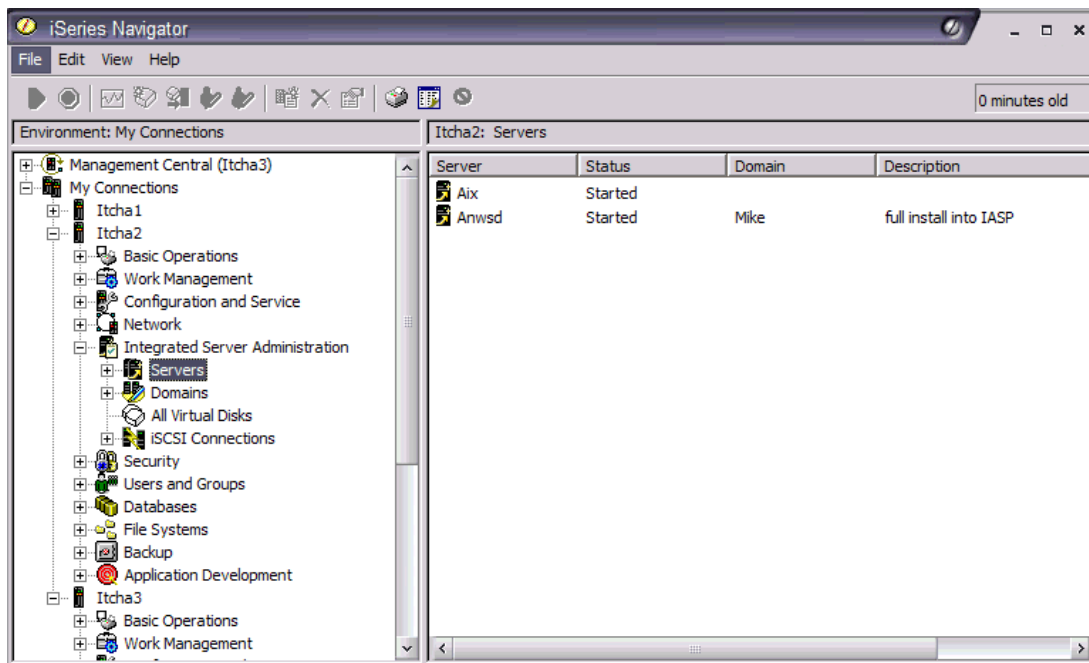
This section will describe the steps taken to create the test environment. See the Appendix for references to documentation on how to perform each of these steps

- Hardware setup and updates following the procedures from the iSCSI install read me first (see Appendix).
- Software installation on both i520s
- Software fixes installed on both systems.
- Configuration of the LUNs and host connection necessary for the IASP on both DS6000s.
- IASP creation using System i Navigator on both i520s.

- iSCSI object creation on the primary System i following the procedures from the iSCSI install read me first.
- Windows Server 2003 R2 installation on one System i, including network storage space creation directly into the IASP.
- Object saves from the primary system are restored on the secondary system. Objects include the NWSD, the virtual Ethernet line description, controller description and device description, NWSH, network server configuration objects (*RMTSYS, *SVRPRC and *CNNSEC) and validation list objects, but not the network storage spaces. The SAVCFG command saves the NWSD, line, controller and device descriptions and NWSH. The SAVOBJ command saves the network configuration objects and their associated validation list objects from the QUSRSYS library.
- AIX 5L LPAR creation on both i520s using the HMC
- Virtual SCSI client and server adapter creation for AIX 5L and i5/OS, respectively.
- AIX 5L NWSD creation in the System ASP.
- AIX 5L NWSSTG creation in the IASP.
- Standard installation of AIX 5L v5.3 TL04 on virtual storage.
- Installation and configuration of Copy Services for System i.
- Creation and compilation of the ENDINTENV and STRINTENV exit programs in library HAVIR, created previously.
- Registration of ENDINTENV and STRINTENV with exit point QIBM_QDC_VRYEXIT, with formats PROF0100 and PSON0200m respectively.

6.0 Test Scenario

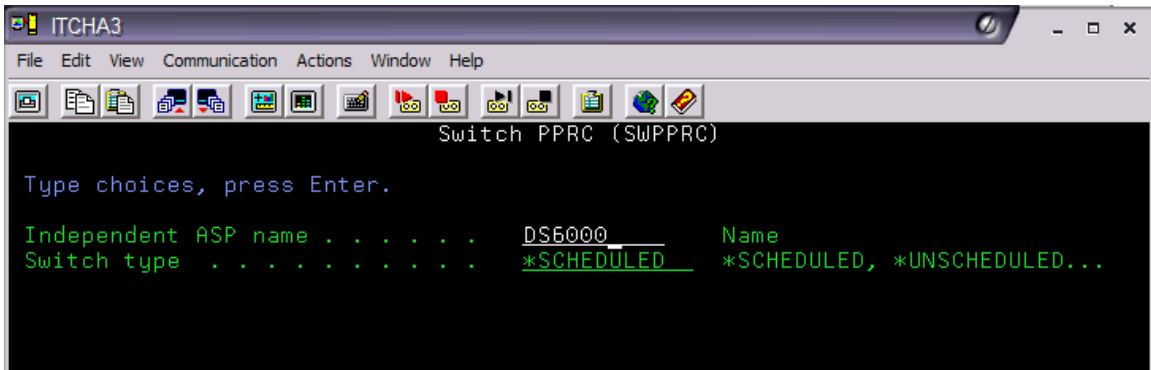
The goal of the tests was to verify that with the help of a vary off and vary on exit program, an integrated blade server or System x with iSCSI or an AIX 5L LPAR can automatically participate in an IASP switchover from a production (ITCHA2 in this case) to a backup system (ITCHA3). As described above, switching the IASP from production to backup and changing the direction of DS6000 replication was handled by Copy Services for System i. In the initial state, the IASP was available on the production system, and the iSCSI and AIX 5L environments were active:



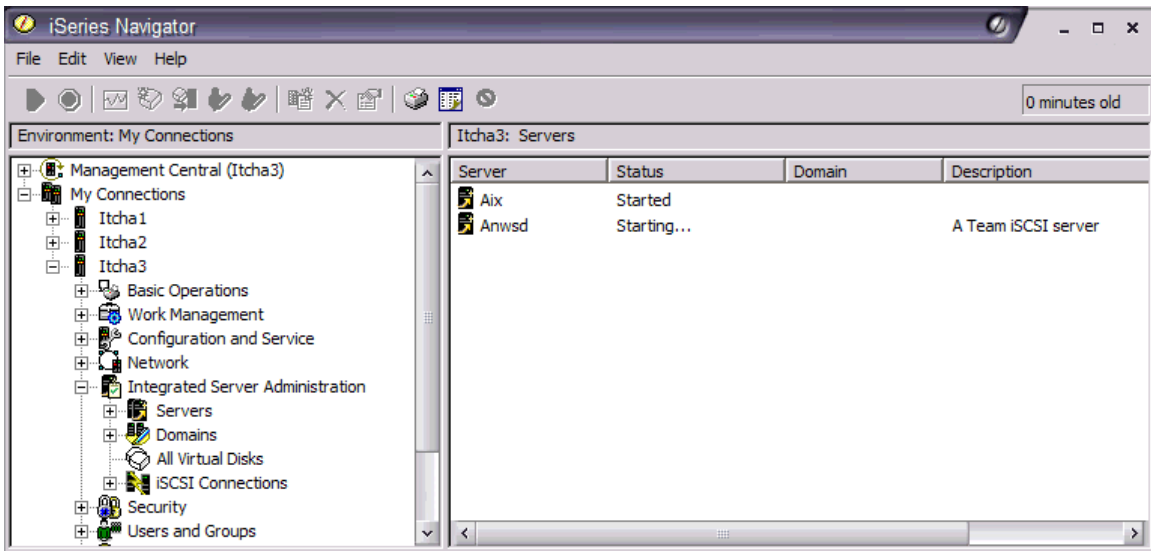
HA automation was configured by creating the appropriate data areas and registering the ENDINTENV and STRINTENV exit programs on both production and backup. Note that in this

test scenario, only one of the hosted environments – integrated server or AIX 5L – was managed by the exit programs. The reason was that the sample exit programs were written to retrieve information about a single integrated environment. This test case can easily be extended to include multiple exit programs to manage several integrated servers or AIX 5L LPARs, or still a pair of exit programs that retrieve information from multiple data areas or database files. The description of this test environment and process is intended as a starting point.

With the integrated environment running on the production system, a switchover of the IASP was initiated with the SWPPRC command from Copy Services for System i on the backup system:



Before the IASP became unavailable on the production system, the iSCSI integrated server or AIX 5L LPAR (both cases were tested) were shut down automatically. After Copy Services for System i had made the IASP available on the backup system, the iSCSI integrated server or AIX 5L LPAR were started automatically:



With the old backup system now acting as production, the hosted environments were switched back to the old production system, again using the SWPPRC command from Copy Services for System i. Switchover of the iSCSI integrated server and AIX 5L LPAR was successful in both directions, because the ENDINTENV and STRINTENV exit programs were registered on both systems.

7.0 Summary

This paper demonstrated how to use existing System i technologies and offerings to provide a level of HA and automate switchover for an iSCSI integrated server and an AIX 5L LPAR using virtual storage in i5/OS. The test case described here relied on the following System i capabilities: integration with BladeCenter and System x with iSCSI, LPAR, virtual storage, IASPs, i5/OS clustering, exit points and CL programs. Copy Service for System i and DS6000 Metro Mirror complete the solution. For System i clients who have hosted workloads and employ external storage replication for i5/OS HA, this solution can provide additional value by simplifying the HA strategy for the entire system into a cohesive plan.

Furthermore, this test case can easily be expanded for use with other hosted configurations or types of replication. Examples of similar environments include POWER Linux or Integrated xSeries Server (IXS). For replication of the IASP, i5/OS Geographic Mirroring can be used instead of Metro Mirror. The goal of this discussion was to provide one example of combining existing System i capabilities to achieve greater HA integration for multiple workloads on the same system.

APPENDIX

Resources and Links

iSCSI integrated servers:

- <http://www.ibm.com/systems/i/bladecenter/iscsi/readme>. iSCSI install read me first.
- <http://www.ibm.com/systems/i/bladecenter>. Home page for the integrated server support. See the associated links for iSCSI overview, models supported, switches for iSCSI and education offerings.
- <http://www.ibm.com/systems/i/bladecenter/library.html>. Brochures, Redbooks, white papers and Information Center links on integrated servers (IXS, IXA or iSCSI).
- INTWIN@us.ibm.com. E-mail address for integrated server general questions or inquiries on contractual support for installations and education

AIX 5L on System i:

- http://www.ibm.com/servers/eserver/series/aix/getting_started.html. Home page for technical information about AIX 5L in an LPAR on System i. Includes links to Redbooks, education offerings and other white papers.

i5/OS clustering, IASPs and Copy Services for System i:

- <http://www.ibm.com/servers/eserver/series/ha/>. Home page for System i HA and clustering. Includes links to a number of related topics, education offerings and other white papers.
- rchclst@us.ibm.com. E-mail address for general i5/OS clustering and IASP questions, or inquiries on contractual support for education and implementations.
- <http://www.ibm.com/servers/eserver/services/seriesHA.html>. Web site with information on the Copy Services for System i offering.

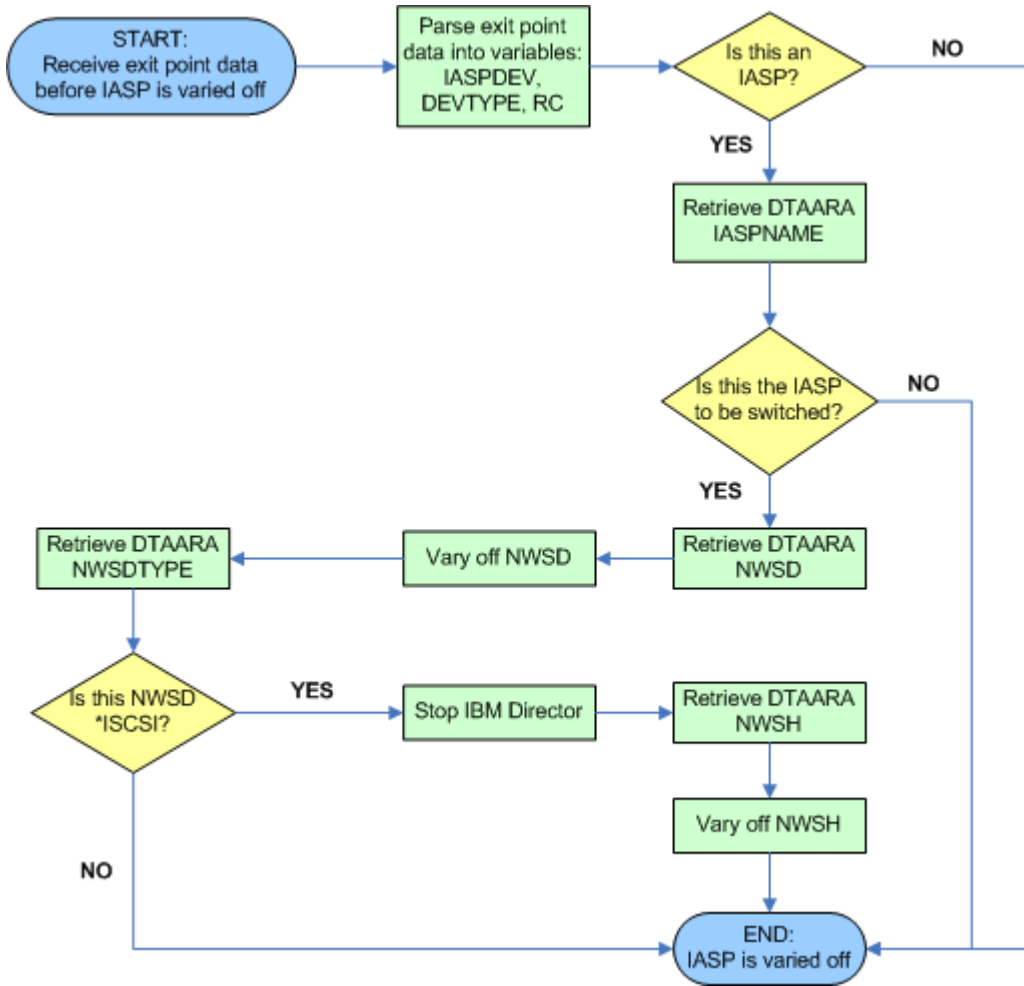
IBM DS6000 and Metro Mirror:

- <http://www.ibm.com/servers/storage/disk/ds6000/index.html>. Home page with links to features, specifications and the IBM System Storage Information Center.
- *iSeries and IBM TotalStorage: A Guide to Implementing External Disk on eServer i5, SG24-7120-00* (<http://www.redbooks.ibm.com/abstracts/sg247120.html?Open>). A Redbook specifically for using IBM System Storage in System i solutions.

Sample CL Programs

This section will present the flowcharts and CL source code for the ENDINTENV and STRINTENV exit programs. Please note that this information is provided by IBM "AS-IS" and no support or warranty is implied or should be inferred.

ENDINTENV flowchart:



ENDINTENV source code:

```

***** Beginning of data *****
0001.00 /* ***** */ 060728
0002.00 /* ***** */ 060728
0003.00 /* Example exit program to vary off an integrated server environment */ 060803
0004.00 /* that uses virtual storage located in an Independent ASP */ 060803
0005.00 /* ***** */ 060728
0006.00 /* To do pre-processing for a Vary Off action on an IASP device, register */ 060728
0007.00 /* this exit program to the QIBM_QDC_VRYEXIT exit point with format */ 060728
0008.00 /* PROF0100, specifying EXIT PROGRAM DATA of DEVDA SPD with a length of 8 */ 060728
0009.00 /* The following command can be used to register the exit program: */ 060728
0010.00 /* ***** */ 060728
0011.00 /* ADDEXITPGM EXITPNT(QIBM_QDC_VRYEXIT) FORMAT(PROF0100) */ 060728
0012.00 /* PGMNBR(1) PGM(QZRDHATLK/ENDINTENV) */ 060828
0013.00 /* TEXT('IASP Vary Off Pre-processing Exit Program') */ 060728
0014.00 /* CRTEXITPNT(*YES) PGM DTA(*JOB 8 DEVDA SPD) */ 060728
0015.00 /* ***** */ 060728
0016.00 /* ***** */ 060728
0017.00 /* ***** */ 060728
0018.00 PGM PARM(&VRYINF &RC) 060728
0019.00 PGM PARM(&VRYINF &RC) 060803
0020.00 060728
0021.00 COPYRIGHT TEXT('Copyright (C) 2006 IBM Corp. All Rights + 060728
0022.00 Reserved.') 050911
0023.00 060803
0024.00 /* Variables representing data passed to/from the program */ 060728
  
```

```

0025.00 /*      32 Character input structure from the exit point      */      060728
0026.00      DCL      VAR(&VRYINF) TYPE(*CHAR) LEN(32)                */      060728
0027.00      */      060728
0028.00 /*      Name of the device description being varried off    */      060728
0029.00      DCL      VAR(&ASPDEV) TYPE(*CHAR) LEN(10)                */      060728
0030.00      */      060728
0031.00 /*      Type of device being varried off                      */      060728
0032.00      DCL      VAR(&DEVTYPE) TYPE(*CHAR) LEN(10)              */      060728
0033.00      */      060803
0034.00 /*      Return code: it will be set to 0 to let IASP vary off */      060803
0035.00      DCL      VAR(&RC) TYPE(*CHAR) LEN(4)                    */      060822
0036.00      */      060803
0037.00 /* Variables for data retrieved from data areas              */      060803
0038.00 /*      IASP name                                                */      060803
0039.00      DCL      VAR(&IASPNAME) TYPE(*CHAR) LEN(10)              */      060803
0040.00      */      060803
0041.00 /*      NWSD name                                                */      060803
0042.00      DCL      VAR(&NWSD) TYPE(*CHAR) LEN(10)                 */      060803
0043.00      */      060803
0044.00 /*      NWSD type                                                */      060803
0045.00      DCL      VAR(&NWSDTYPE) TYPE(*CHAR) LEN(9)             */      060803
0046.00      */      060803
0047.00 /*      NWSH name, if NWSD type = *ISCSI                          */      060803
0048.00      DCL      VAR(&NWSH) TYPE(*CHAR) LEN(10)                 */      060803
0049.00      */      060728
0050.00 /*      All errors from the exit program are ignored, so don't surface */      060728
0051.00      MONMSG  MSGID(CPF0000)                                    */      050913
0052.00      */      060728
0053.00 /* Read the Input Info. in to separate variables            */      060728
0054.00      CHGVAR  VAR(&ASPDEV) VALUE(%SST(&VRYINF 1 10))          */      060728
0055.00      CHGVAR  VAR(&DEVTYPE) VALUE(%SST(&VRYINF 11 10))       */      060728
0056.00      CHGVAR  VAR(&RC) VALUE(X'00000000')                      */      060822
0057.00      */      060728
0058.00 /* Uncomment this statement for debugging or informational purposes */      060728
0059.00 /*      SNDPGMMSG MSG('Vary off information passed to exit +    */      060728
0060.00 /*          Program. Device Name: ' *TCAT +                        */      060728
0061.00 /*          &ASPDEV *TCAT ' Device Type: ' *TCAT +                */      060728
0062.00 /*          &DEVTYPE *TCAT ' Format: ' *TCAT +                    */      060728
0063.00 /*          &FORMAT *TCAT ' Force: ' *TCAT +                      */      060728
0064.00 /*          &FORCETXT *TCAT ' Return Code: ' *TCAT +            */      060728
0065.00 /*          &RCTXT) TOPGMQ(*EXT)                                  */      060728
0066.00      */      060728
0067.00      IF      COND(&DEVTYPE *NE '*DEV'D') THEN(GOTO +        */      060728
0068.00          CMDLBL(END))                                          */      060728
0069.00      */      060803
0070.00 /*      This 10-character data area is expected to contain the name */      060803
0071.00 /*      of the IASP for which we want to take action                */      060803
0072.00      RTVDTAARA DTAARA(HAVIR/IASPNAME (1 10)) RTNVAR(&IASPNAME) */      060803
0073.00      */      060728
0074.00 /*      Check to see if the IASP name from the exit point is      */      060803
0075.00 /*      the same as that from the DTAARA. Exit if not              */      060803
0076.00      */      060803
0077.00      IF      COND(&ASPDEV *EQ &IASPNAME) THEN(GOTO +        */      060803
0078.00          CMDLBL(PROCESS))                                        */      060803
0079.00      ELSE      CMD(DO)                                          */      060728
0080.00      GOTO      CMDLBL(END)                                      */      060728
0081.00      ENDDO                                                    */      060728
0082.00      */      060728
0083.00 /*      This 10-char data area is expected to contain the name    */      060803
0084.00 /*      of the NWSD for which we want to take action                */      060803
0085.00      PROCESS: RTVDTAARA DTAARA(HAVIR/NWSD (1 10)) RTNVAR(&NWSD) */      060803
0086.00      */      060803
0087.00 /*      Vary off the NWS                                          */      060803
0088.00      VRYCFG  CFGOBJ(&NWSD) CFGTYPE(*NWS) STATUS(*OFF)        */      060824
0089.00      */      060803
0090.00 /*      This 9-char data area is expected to contain the type      */      060803
0091.00 /*      of NWSD. A NWSD type of *ISCSI requires additional action */      060803
0092.00      RTVDTAARA DTAARA(HAVIR/NWSDTYPE (1 9)) RTNVAR(&NWSDTYPE) */      060803
0093.00      IF      COND(&NWSDTYPE *NE '*ISCSI') THEN(DO)            */      060803
0094.00      GOTO      CMDLBL(END)                                      */      060803

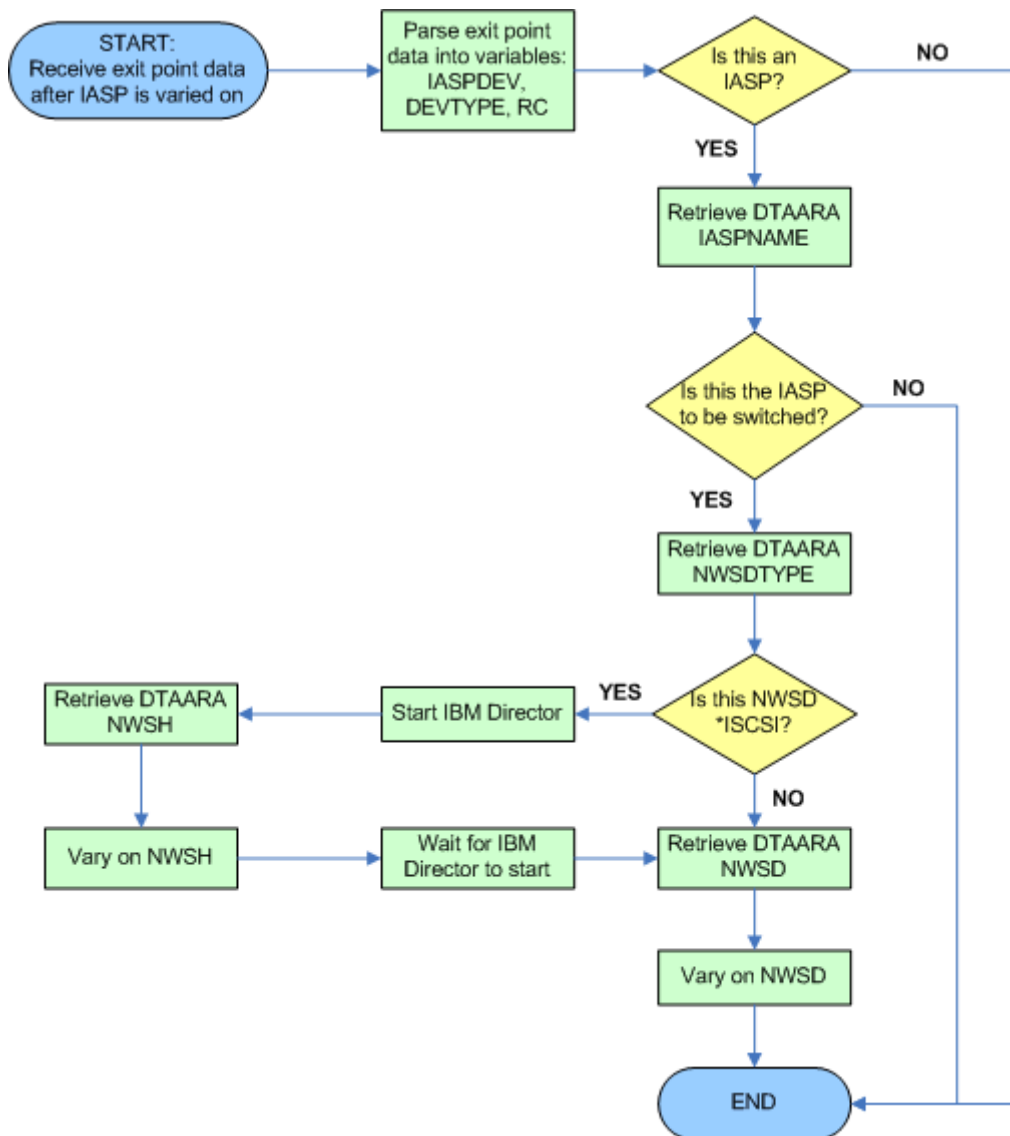
```

```

0095.00      ENDDO
0096.00      ELSE      CMD(GOTO CMDLBL(ISC SI))
0097.00
0098.00 /*      iSCSI-specific processing. Stop IBM Director and vary off */
0099.00 /*      the corresponding NWSH */
0100.00 ISCSI:  QSH      CMD('/qibm/proddata/director/bin/twgstop > +
0101.00              /dev/null')
0102.00
0103.00 /*      Read the name of the NWSH from this 10-character data area */
0104.00      RTVDTAARA  DTAARA(HAVIR/NWSH(1 10)) RTNVAR(&NWSH)
0105.00
0106.00 /*      Vary off the NWSH */
0107.00      VRYCFG    CFGOBJ(&NWSH) CFGTYPE(*DEV) STATUS(*OFF)
0108.00
0109.00 END:
0110.00      ENDPGM
***** End of data *****
060803
060803
060803
060803
060828
060828
060803
060803
060803
060803
060822
050911
060728
050912

```

STRINTENV flowchart:



STRINTENV source code:

```

***** Beginning of data *****
0001.00 /* ***** */ 060728
0002.00 /* */ 060728
0003.00 /* Example exit program to vary on an integrated server environment */ 060803
0004.00 /* that uses virtual storage located in an Independent ASP */ 060803
0005.00 /* */ 060728
0006.00 /* To do post-processing for a Vary On action on an IASP device, register */ 060828
0007.00 /* this exit program to the QIBM_QDC_VRYEXIT exit point with format */ 060728
0008.00 /* PSON0100, specifying EXIT PROGRAM DATA of DEVDasPD with a length of 8 */ 060828
0009.00 /* The following command can be used to register the exit program: */ 060728
0010.00 /* */ 060728
0011.00 /* ADDEXITPGM EXITPNT(QIBM_QDC_VRYEXIT) FORMAT(PSON0200) */ 060828
0012.00 /* PGMNBR(1) PGM(QZRDHATLK/STRINTENV) */ 060828
0013.00 /* TEXT('IASP Vary On Post-processing Exit Program') */ 060828
0014.00 /* CRTEXITPNT(*YES) PGMDTA(*JOB 8 DEVDasPD) */ 060728
0015.00 /* */ 060728
0016.00 /* */ 060728
0017.00 /* ***** */ 060728
0018.00 060728
0019.00 PGM PARM(&VRYINF &RC) 060803
0020.00 060728
0021.00 COPYRIGHT TEXT('Copyright (C) 2006 IBM Corp. All Rights + 060728
0022.00 Reserved.') 050911
0023.00 060803
0024.00 /* Variables representing data passed to/from the program */ 060728
0025.00 /* 32 Character input structure from the exit point */ 060728
0026.00 DCL VAR(&VRYINF) TYPE(*CHAR) LEN(32) 060728
0027.00 060728
0028.00 /* Name of the device description being varried off */ 060728
0029.00 DCL VAR(&ASPDEV) TYPE(*CHAR) LEN(10) 060728
0030.00 060728
0031.00 /* Type of device being varried off */ 060728
0032.00 DCL VAR(&DEVTYPE) TYPE(*CHAR) LEN(10) 060728
0033.00 060803
0034.00 /* Return code: it will be set to 0 to let IASP vary off */ 060803
0035.00 DCL VAR(&RC) TYPE(*CHAR) LEN(4) 060803
0036.00 060803
0037.00 /* Variables for data retrieved from data areas */ 060803
0038.00 /* IASP name */ 060803
0039.00 DCL VAR(&IASPNAME) TYPE(*CHAR) LEN(10) 060803
0040.00 060803
0041.00 /* NWSD name */ 060803
0042.00 DCL VAR(&NWSD) TYPE(*CHAR) LEN(10) 060803
0043.00 060803
0044.00 /* NWSD type */ 060803
0045.00 DCL VAR(&NWSDTYPE) TYPE(*CHAR) LEN(9) 060803
0046.00 060803
0047.00 /* NWSH name, if NWSD type = *ISCSI */ 060803
0048.00 DCL VAR(&NWSH) TYPE(*CHAR) LEN(10) 060803
0049.00 060728
0050.00 /* All errors from the exit program are ignored, so don't surface */ 060728
0051.00 MONMSG MSGID(CPF0000) 050913
0052.00 060728
0053.00 /* Read the Input Info. in to separate variables */ 060728
0054.00 CHGVAR VAR(&ASPDEV) VALUE(%SST(&VRYINF 1 10)) 060728
0055.00 CHGVAR VAR(&DEVTYPE) VALUE(%SST(&VRYINF 11 10)) 060728
0056.00 CHGVAR VAR(&RC) VALUE(X'00000000') 060803
0057.00 060728
0058.00 /* Uncomment this statement for debugging or informational purposes */ 060728
0059.00 /* SNDPGMMSG MSG('Vary off information passed to exit + */ 060728
0060.00 /* Program. Device Name: ' *TCAT + */ 060728
0061.00 /* &ASPDEV *TCAT ' Device Type: ' *TCAT + */ 060728
0062.00 /* &DEVTYPE *TCAT ' Format: ' *TCAT + */ 060728
0063.00 /* &FORMAT *TCAT ' Force: ' *TCAT + */ 060728
0064.00 /* &FORCETXT *TCAT ' Return Code: ' *TCAT + */ 060728
0065.00 /* &RCTXT) TOPGMQ(*EXT) */ 060728
0066.00 060728

```

```

0067.00      IF      COND(&DEVTYPE *NE '*DEVD') THEN(GOTO +          060728
0068.00          CMDLBL(END))                                         060728
0069.00
0070.00 /*    This 10-character data area is expected to contain the name    */ 060803
0071.00 /*    of the IASP for which we want to take action                    */ 060803
0072.00      RTVDTAARA DTAARA(HAVIR/IASPNAME (1 10)) RTNVAR(&IASPNAME) 060803
0073.00
0074.00 /*    Check to see if the IASP name from the exit point is          */ 060803
0075.00 /*    the same as that from the DTAARA. Exit if not                    */ 060803
0076.00
0077.00      IF      COND(&ASPDEV *EQ &IASPNAME) THEN(GOTO +          060803
0078.00          CMDLBL(PROCESS))                                         060803
0079.00      ELSE      CMD(DO)                                           060728
0080.00      GOTO      CMDLBL(END)                                       060728
0081.00      ENDDO
0082.00
0083.00 /*    This 9-char data area is expected to contain the type          */ 060803
0084.00 /*    of NWS. A NWS type of *ISCSI requires additional action          */ 060803
0085.00 PROCESS: RTVDTAARA DTAARA(HAVIR/NWSDTYPE (1 9)) RTNVAR(&NWSDTYPE) 060828
0086.00      IF      COND(&NWSDTYPE *NE '*ISCSI') THEN(DO)              060803
0087.00      GOTO      CMDLBL(OTHER)                                       060803
0088.00      ENDDO
0089.00      ELSE      CMD(GOTO CMDLBL(ISCSI))                             060803
0090.00
0091.00 /*    iSCSI-specific processing. Start IBM Director and vary on      */ 060803
0092.00 /*    the corresponding NWSH                                           */ 060803
0093.00 ISCSI:  QSH      CMD('/qibm/proddata/director/bin/twgstart > + 060828
0094.00          /dev/null)
0095.00
0096.00 /*    Read the name of the NWSH from this 10-character data area      */ 060803
0097.00      RTVDTAARA DTAARA(HAVIR/NWSH (1 10)) RTNVAR(&NWSH)          060803
0098.00
0099.00 /*    Vary on the NWSH                                               */ 060828
0100.00      VRYCFG  CFGOBJ(&NWSH) CFGTYPE(*DEV) STATUS(*ON)          060828
0101.00
0102.00 /*    Delay varying on the NWS while waiting for Director to start    */ 060828
0103.00      DLYJOB  DLY(240)
0104.00
0105.00 /*    Read the name of the NWSD, vary on the NWS                      */ 060828
0106.00 OTHER:  RTVDTAARA DTAARA(HAVIR/NWSD (1 10)) RTNVAR(&NWSD)      060828
0107.00      VRYCFG  CFGOBJ(&NWSD) CFGTYPE(*NWS) STATUS(*ON)          060828
0108.00
0109.00 END:
0110.00      ENDPGM
***** End of data *****

```

For More Information

To learn more about the IBM System i contact your IBM representative, IBM Business Partner or visit the following Web sites:

<http://www-03.ibm.com/systems/i/>
<http://www-1.ibm.com/servers/eserver/series/aix/>

©Copyright IBM Corporation 2006

IBM Systems Group
Route 100
Somers, NY 10589

IBM, the IBM logo, the e-business logo, AIX 5L, Domino, i5, i5/OS, iSeries, Lotus, Operating System/400, POWER, POWER5, Virtualization Engine, WebSphere and xSeries, System Storage, DS6000, and System i are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries or both.

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

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

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

IBM hardware products are manufactured from new parts, or new and used parts. In some cases, the hardware product may not be new and may have been previously installed. Regardless, IBM warranty terms apply.

References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates.

NOTES:

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Users of this document should verify the applicable data for their specific environment.

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