

IBM Enterprise2013

pST588 - Planning For System Recovery

Grover Davidson – grover@us.ibm.com



Enterprise2013

Objectives

- Expand thinking to go way outside of just AIX
- Understand key points to consider in planning
- Focus on **RECOVERING** data not backups

Factors influencing backups

- Consolidated systems
- Level of application integration/ability to synchronize multiple applications
- Operations in multiple time zones
 - “Normal” operating hours have extended well outside of 8-5
- Required web presence that is integral to business
 - Business is exclusively on-line
- Downtime cannot be tolerated due to nature of business
 - Stock trading
 - Banking

Backing up to restore

- Backups should not be done just to complete a check box
- Time is a major factor when backups are used to recover business critical system
- The recovery process should be the main driving factor in backups
- Without recovery, backups are wasted resources

Special Considerations: LPARs

- Documented at:
<http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/p7hbm/copylparconfig.htm>
- Use the system planning tool to save the CEC/LPAR configuration to the local HMC
- Then use the 'System Plan' icon on the far left side to manage system plans – they can be saved to media or ftp'ed to another system
- Name XXXX.sysplan and is an *binary* file (be sure to use binary mode when ftp'ing this file)
- PC side of system planning tool is available at: <http://www-947.ibm.com/systems/support/tools/systemplanningtool/>
- Can be used to import/clone/restore LPAR configuration on a CEC
- ***Always check GUI for success/failure status***
- Also available on the IVM (Integrated Virtualization Manager)
- Can be sent to DR site ahead of time

Special Considerations: VIO Servers

- Documented at:
http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/iphb1/iphb1_vios_backup_backup.htm
- Restore process is documented here:
http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/iphb1/iphb1_vios_backup_restore.htm
- Should also perform a snap command from the padmin login and before normal backup process outlined above
- Whether you use backupios or viosbr depends on whether you are restoring to the same hardware:
 - backupios can be used to restore on different hardware
 - viosbr can ONLY be used to restore to the same hardware
- Need to keep current with HMC/VIO/Sysplan code to avoid issues
- Needs to be pulled off the VIO server and saved someplace
- Each VIO Server needs to be backed up individually
- Remember that the LPAR/CEC configuration needs to be in place for the VIO server restoration to work correctly
- **Virtual device configuration will be lost if recovery is on different hardware**

Special Considerations: AIX

- AIX OS file systems cannot be saved and restored for system recovery by normal backup processes
- AIX OS files can be saved and recovered on a file-by-file basis with almost any command
- AIX installation should account for the system it may be restored on
- Firmware levels should also be saved
- Always create a mksysb for AIX before applying APARs

Special considerations: Volume Groups

- Volume group structure data is only saved for rootvg by default
- Having volume group structure data allows for quick recovery of the volume group, logical volumes and file systems including customized attributes
- Damaged volume groups may be able to be reconstructed if maps are available
- Volume group structures can also be modified by changing the data file before restoring
- Consider using 'savevg -ipr -f OUTFILE VGNAME'
 - NOTE: this does not save any data!
 - Cannot be used on rootvg (use mksysb for rootvg)!
 - Creates a backup/restore format file
 - Attribute of the VG can then be edited before running restvg

Special considerations: Volume Groups

```
# savevg -ipr -f /tmp/k.img l3kern_vg
# restore -qf /tmp/k.img
x ./tmp/vgdata/vgdata.files43932
x ./tmp/vgdata/vgdata.files
x ./tmp/vgdata/l3kern_vg/filesystems
x
x ./tmp/vgdata/l3kern_vg/l3kern_vg.data
x ./tmp/vgdata/l3kern_vg/backup.data
# vi ./tmp/vgdata/l3kern_vg/l3kern_vg.data
# restvg -f ./tmp/vgdata/l3kern_vg/l3kern_vg.data -r hdiskX
```

Special considerations: VGDA's

- VGDA – Volume Group Descriptor Area – describes the VG (Volume Group) layout
- At least one copy on each disk in a VG
- Best to save this before running VG related commands in case of a failure:
 - Small VG:
 `dd if=/dev/hdiskX of=/tmp/hdiskX.vgda bs=128K count=17`
 - Big VG:
 `dd if=/dev/hdiskX of=/tmp/hdiskX.vgda bs=128K count=71`
 - SVG(Scalable):
 `dd if=/dev/hdiskX of=/tmp/hdiskX.vgda bs=512K count=137`

Special considerations: File systems

- Many file systems have logs that must be saved as raw devices
- File systems that share a log must be backed up together
- Some file systems have special requirements for handling the logs when snapshots are done
- File systems/backups should be tested for integrity before they are saved
- Single recovery stream should be used to reduce file fragmentation and the performance impacts of fragmented files
- See Developer Works article on JFS2 snapshots:
http://www.ibm.com/developerworks/aix/library/a-jfs2_snapshot.html

Special Considerations: Encrypted File Systems

- Save the files in /var/efs, /etc/security/user, /etc/security/group
- Keep a record of the key store encryption password
 - **Loss of this information will make recovery of encrypted file systems impossible**
- Know how the backup software will store that on the backup media – encrypted or unencrypted
- Know where all the needed filesets are to restore the system to the exact same levels
- More information is available from “AIX V6 Advanced Security Features Introduction and Configuration”:
<http://www.redbooks.ibm.com/redbooks/pdfs/sg247430.pdf>
- **TEST** your procedures!
 - Change the date forwards for restores
 - Understand key management and know how to reset keys
 - Be very careful about over-writing current keys when restoring
 - Consider using a second staging machine to protect current data
 - Use a staging machine to ensure you do not affect your current data!

Special Considerations: GPFS

- Save the /var/mmfs/gen/mmsdrfs file on each node
 - Contains all GPFS definitions
- Use the 'mmfsctl' command with suspend/resume operations before/after performing a hardware mirror for backups
- See GPFS advanced administrators guide at:
<http://publib.boulder.ibm.com/epubs/pdf/c2351824.pdf>

Special considerations: Raw Logical Volumes

- Cannot be saved/restored with file system related commands
- Have their own permissions/owners
- Striping can create 'fractured' blocks
- Any changes to the size of the LV may be problematic for the recovery software or user application

Special considerations: Databases

- Usually have logs that must be saved and restored
- Log files may need to be rotated just after the backup process starts and before actual data is saved
- Log files should also be rotated as the process completes
- Control/Configuration files may need be saved before and after the backup
- What would happen if a single copy of a single log cannot be restored?
- Consult with your DBA. Make sure all these items are addressed as part of your backup procedures

Special Considerations: System

- Build a recovery file system for each host
- Document the adapters, network configuration, file set levels, VG structures and other host unique information
- Make this near the beginning of any backup
- Allows for fast restore if needed
- Facilitates faster system recovery and host consolidation if needed

Disk Systems

- Many provide for instant point-in-time backup (snap copy) facilities
- Each snap copy can require up to 100% of the original disk space
- The snap copy can then be mounted on a different system and read for backups or other processing
- Provides a means to test the integrity of the database by starting the database on the backup system
- Many disk systems combine this with a remote mirroring option to create a duplicate storage system at a remote site
- Extreme care should be taken to prevent snapshot copies from being seen by the original host – otherwise data loss can occur (and HAS occurred in many instances)

Special Considerations: Disk systems

- Remote mirroring configuration may affect performance
- Raid configurations have different performance factors
- Connectivity/throughput between systems will affect performance

Special Considerations: Storage Architecture

- All configurations for SAN fabric, SVCs
- Document Firmware levels/model numbers for all devices
- Save firmware installation media/files if possible
- Record detailed zoning information
- Check if the FC Switches/SVCs can save their configurations to a TFTP server
- Document the process to restore the saved configuration and test the procedures

Backup infrastructure

- Can you recover your backup infrastructure itself??
- All backups need to be recoverable
- Changes in backup software may result in backup that cannot be restored
- Hardware may also need to be preserved to allow for recovery
- Need to have plans for how to recover the entire backup/recovery environment from nothing but the existing backups
- Evaluate the need for a crash kit that has all needed resources to recover from nothing
- Know where all software licensing information is and how to contact the vendors for emergency support

Designing for recovery

- Shift the focus to recovering data
- Evaluate each system to understand the impact of downtime
- Plan backup and recovery resources based on the impact of downtime
- Do not expect one plan/strategy meets all needs
- Know how long each recovery option takes
- Identify super critical information and have a plan to deal with failures of that data – archive logs might need to be on 2 different backup medias for example
- Change the backup procedures to support recovery - generate the recovery process as part of the backups

Final thoughts

- There are many software packages that save/recover data
- Multiple backup/recovery strategies may be needed
- Business function and criticality of a host will help determine what approach(es) to take
- Careful planning and implementation will result in successful recovery from multiple outages and failures
- Recover planning takes a lot of planning and testing to be successful
- Failing to be prepared for recovery can cripple or destroy your business

Checklist for recovery

- Create and save system plans for all LPARs
- Create and save storage subsystem layout backups
- Create a /recovery-\$hostname file system in rootvg for each host – this way multiple recovery filesystems can be restored without a problem
- Identify the unique data on each host and the best way to backup the data to the /recovery file system
- Make sure to save structure information for all volume groups and the maps of all LVs
- Save a snap with everything except a dump
- Save configuration files for all other products – HACMP, GPFS, db2, oracle, etc.
- Test recovering a system on a regular basis
- Be prepared to recover the backup/recover host and how to make it faster/easier

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States, other countries, or both.

Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.

Those trademarks followed by ® are registered trademarks of IBM in the United States; all others are trademarks or common law marks of IBM in the United States.

For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml:

*, AS/400®, e business (logo)®, DBE, ESCO, eServer, FICON, IBM®, IBM (logo)®, iSeries®, MVS, OS/390®, pSeries®, RS/6000®, S/30, VM/ESA®, VSE/ESA, WebSphere®, xSeries®, z/OS®, zSeries®, z/VM®, System i, System i5, System p, System p5, System x, System z, System z9®, BladeCenter®

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

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

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

* All other products may be trademarks or registered trademarks of their respective companies.

Notes:

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

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

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

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the 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.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.