Mel Beckman • Richard Dolewski • Dan Riehl • Debbie Saugen

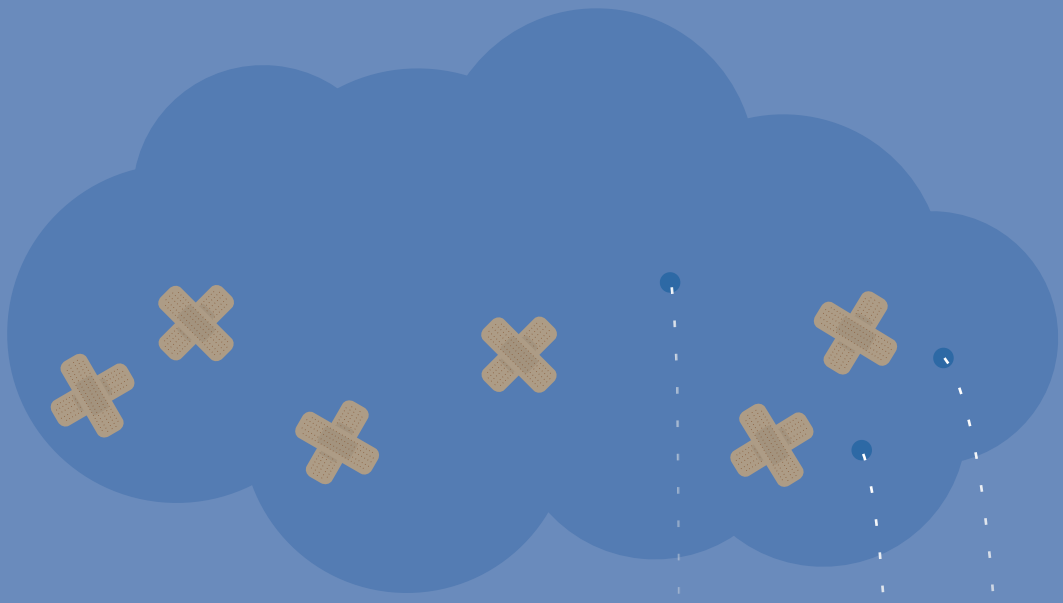

RECOVERY WITHOUT DISASTER

Disaster Recovery Planning for IBM i

iProDeveloper
Technical ADVISOR
Sponsored by
help/systems
World's Leader in Power Systems™ Software Solutions

TABLE OF CONTENTS

ABOUT THE AUTHORS

Mel Beckman is Senior Technical Director for iPro Developer and President of Beckman Software Engineering, a technical consultancy specializing in large-scale, high-bandwidth application hosting networks. He has written extensively on XML, SOA, and transaction performance tuning; and has presented seminars on computer programming and network technology throughout the United States, Europe, and Asia.

Richard Dolewski is Chief Technology Officer and Vice President of Business Continuity Services for WTS, Inc. He is a certified Systems Integration Specialist and Disaster Recovery Planner and is globally recognized as a subject matter expert for business continuity for IBM iSeries and i5 environments.

Dan Riehl is President of IT Security and Compliance Group, LLC. Dan performs IBM i security assessments and provides customized security services and software solutions for his customers. He also provides training in all aspects of IBM i security and other technical areas through his training company, The 400 School, Inc. Dan has written for iPro Developer for more than 20 years and is an iPro Developer senior technical editor. He regularly teaches classes and seminars on IBM i security and other technical topics.

Debbie Saugen is the technical owner of IBM i Backup and Recovery in the IBM Rochester, Minnesota, development lab. She is also a senior recovery architect/consultant with IBM Business Continuity and Resiliency Services. Debbie enjoys sharing her knowledge by speaking at COMMON, technical conferences, and Business Continuity and Resiliency conferences; and by writing for various magazines and websites.

ABOUT THE SPONSOR

Help/Systems improves IT operations through automated scheduling, backup and recovery, message management, system performance monitoring, and report distribution for IBM Power Systems servers running IBM i. With additional solutions for business intelligence and security, Help/Systems has been the leading independent software provider for IBM servers since 1982. See the entire collection of software solutions at www.helpsystems.com.

INTRODUCTION

By Mel Beckman

Running business IT is more than a full-time job even in the best conditions. But as IT professionals, we need to keep it going in the worst of them—from fires and floods to human-induced data and application outages. Fortunately, in the last decade we've been given tools to fight the effects of disaster: massive disk caches, virtual tape libraries, sophisticated data duplication engines, and network transport to distant locations via the cloud.

The following articles will help you put these advances to wise use in your organization's Disaster Recovery (DR) plans. Outlining best practices for strategy, products, and the scope of your program, each chapter will help you ensure that you're providing cost-effective, resilient protection that can grow with your business and adapt to technical innovations. With this knowledge under your belt, you'll be equipped to keep pace with continuing growth of IBM i mission capabilities.

CHAPTER 1: DON'T BAIL ON DISASTER RECOVERY PLANNING

By Richard Dolewski

Most IT organizations, regardless of size, are under constant pressure to reduce the cost of service delivery. Some organizations have gone as far as placing disaster recovery (DR) spending in the category of unnecessary expenses, with chief financial officers believing that reduced profits mean all costs must be cut proportionally, including DR. Neglecting DR can have detrimental consequences for organizations, making them less likely to bounce back when a bailout is needed following a disaster.

Excessive cuts may severely impact a company's ability to test and improve its disaster recovery plan. Many organizations today have a mixture of both internal and outsourced DR solutions. Some businesses may be tempted to let disaster recovery contracts expire and to postpone or cancel testing activities. But is this really an effective cost-saving measure? Paying your commercial hot-site provider every month but halting testing makes no sense. These actions may save money in the short term, but they will only heighten a company's vulnerabilities—and hence its costs—over time.

Some Common DR Shortcuts

Making cuts to disaster recovery tempts many companies when budgets are tight. Here are some of the most common shortcuts—and their consequences.

Postponing disaster recovery testing

This might be seen as a quick way to save money if travel budgets and human resources are tight.

> *Consequences*: Last year's test results won't suffice and the auditor will not show leniency because of the economic downturn. Nothing could be more dangerous for your organization and your career.
>
> *Recommended*: Always test your backups to ensure your ability to recover the infrastructure to meet recovery time objectives; and recover your data to a known, useable, and consistent state to validate recovery point objectives.

Delaying buying new tape

Sometimes this is intentional for cost-reduction purposes. More often, however, someone intended to change the backup tapes but simply never did.

Consequences: Overusing tape will affect the reliability of the tape contents.

Recommended: Follow manufacturer recommendations and always make sure your tapes are in good working order. Tape validation should be performed with an active disaster recovery test to ensure that the backup is complete, the tape media is readable, and the contents meet the requirements for full server restoration.

Scrimping on off-site storage costs

Daily pickup costs $20 to $30 per day plus a fee for the quantity of tapes in rotation. Budget officers might be tempted to switch to weekly pickup.

Consequences: Imagine having to tell your CEO that an entire week's worth of data for company sales and related business activities has been lost because you tried to save $30 for a new tape media cartridge or $400 in off-site pickup charges. Your recovery point objectives just went from 24 hours to seven days.

Recommended: You may be holding many of your archived data tapes far too long in the forgotten storage container off site. Look for your savings there.

Reducing recovery capacity to save on monthly hot-site fees

Consequences: Saving pennies will cost you many dollars later. How do you plan to recover your servers? Only recovery servers are able to handle the workload necessary to run near-normal operations. Reducing capacity is a sure recipe for a second ordeal after your initial disaster.

Recommended: You cannot randomly set up systems that are 50% of your current deployment and hope they will meet your current service levels. When testing your recovery, execute sample test scripts to benchmark performance to ensure the infrastructure will perform to meet your business needs when you need it most.

Finding a Balance

A more effective way to manage IT spending and expectations is to ensure that spending cuts provide a balance between reducing costs and meeting your recovery time and point objectives. An imbalance may compromise your ability to respond to a crisis, so it is increasingly vital to give disaster recovery the proper weight when making difficult budgetary decisions.

Ask yourself the following questions:

- Does it make sense to change my disaster recovery planning expectations for the business?
- Can the business concede extended recovery point or recovery time objectives?
- Do I outsource recovery testing or do it myself? Can we afford not to test?
- How do I assess the risks for my organization? Is the business aligned with IT objectives?
- How do I make the right technology choices for disaster recovery?
- Should I introduce new technology methodologies to recover critical business operations?

IT departments increasingly need to demonstrate their value both internally and externally. Disaster recovery is essential for ensuring business continuity, and a comprehensive plan for disaster recovery is more critical than ever. Saving pennies and risking survival dollars in a disaster does not make financial sense.

CHAPTER 2: TOP 10 "DON'T-DOS" FOR IBM i DISASTER RECOVERY

By Debbie Saugen

When you recover information from tape media, it is critical to have a complete backup strategy to ensure total recovery of your data. Whether you're testing your recovery strategy or performing a real disaster recovery, avoiding common mistakes will make sure that your recovery goes as smoothly as possible. To that end, here the top 10 mistakes I see and how to avoid them.

10 - Using Only One Set of Tape Media

You should always have at least two sets of tape in the case of media errors or incomplete saves. The best option is to make duplicate copies of your backup media and send your most current backups, with a set of your previous backups, to your recovery site.

9 - Labeling Tape Media Incorrectly

Ensure labels are correct and in the recommended format. Using improvised labels like sticky notes may damage the tape and tape device. Also make sure to position the label correctly, replace any damaged or missing labels, and remove the existing label before adding a new one.

8 - Not Using Tape Management

Without tape management, you can't know what data resides on your tape media or its location, making disaster recovery a real nightmare. Use a tape management system along with your backup to locate, track, and rotate your media according to a defined set of policies.

7 - Shipping Tape Media Improperly

Media damaged in transit due to improper packaging causes a painful experience when the media can no longer be read to perform the recovery.

Always ship tape media in original (or better) packaging. Always ship tape cartridges in a jewel case, and use recommended shipping containers that securely hold the cartridge during transportation.

Never ship tape media in a commercial envelope; always place it in a box or package. If you ship the cartridge in a cardboard or sturdy material box, ensure the following:

- Place the cartridge snugly in polyethylene plastic wrap or bags to protect it from dust, moisture, and other contaminants.
- Double-box the cartridge (place it inside a box, then put that box inside the shipping box) and add padding between the two boxes.

6 - Storing Tape Media On Site

How often you send your backup media off site depends on your recovery point objective (RPO). If the minimum amount of data you can lose is 24 hours' worth, you should move your tape media off site every 24 hours.

Do not wait to send a tape off site until it is full. If you haven't shipped a tape for a week and it's destroyed in a fire or flood, you'll lose a lot more than just a tape.

During many disasters, you will have enough warning to perform a full system backup before the disaster occurs. However, during hurricanes, many businesses have discovered that once their area is declared a mandatory evacuation zone, tape vendors will not come to pick up their tape media. Making frequent backups can minimize the amount you lose when catastrophe hits.

As another measure against natural disasters, always store your tape media off site, out of reach of potential disaster risks. Also make sure the recovery media is close to your recovery site—system recovery cannot start until the tapes arrive at that site.

5 - Not Following Recovery Procedures

Having well-documented recovery procedures is vital for any IBM i shop. You might adopt the procedures in *IBM i - Systems management - Recovering your system (IBM Redbook SC41-5304-10)*, IBM recovery scripts provided by IBM Business Continuity and Resiliency Services (BCRS), or your own custom recovery scripts.

Most important is that everyone involved in performing the recovery read and follow the procedures completely. This may seem like common sense, but recoveries fail when staff members don't thoroughly read and follow the steps, or if they rush through them. If you're unsure how to proceed, ask for the appropriate technical assistance.

4 - Failing to Perform a Complete Recovery Test

Testing the recovery of your systems does not ensure you can recover your business in the event of a disaster. And being able to restore your data does not mean that you can recover your IT environment and resume business activity in the required time frame. Recovering IBM i systems is the easy part. Harder is reestablishing your network connectivity and validating the applications and data integrity. A complete test also includes procedures for alert management, declaration, chain of command, and reporting.

One of the most difficult steps of disaster recovery is the process of deciding when to declare a disaster. Some disasters do not have obvious times to determine. You do not want an unqualified person to declare a disaster at the first sign of trouble, nor do you want any delays in starting the recovery because no one knows what to do or when to start.

3 - Executing Incomplete Saves

Only data that has been saved can be recovered. Data missing from the save strategy or objects with exclusive locks during the save processing will result in incomplete recoveries. Do not be fooled into thinking the Save System (SAVSYS) command saves everything on your system. SAVSYS saves only the Licensed Internal Code (LIC), operating system, user profiles, and device configuration objects. Save menu option 21 saves your entire system by executing the following:

- Ending all subsystems
- Saving the Licensed Internal Code
- Saving the operating system
- Saving security data
- Saving device configuration
- Saving all libraries
- Saving all documents and folders
- Saving all directories

The advantage of using an option 21 save is that it completely saves everything on your system. The disadvantage is that the system is unavailable to users during the entire save process.

If you have a very short backup window that requires a more complex backup strategy, you can use one (or a combination of) the following methods:

- Save system information in a nonrestricted state
- Save data concurrently using multiple tape devices
- Save data in parallel using multiple tape devices
- Use the Save While Active process

Before you use any of those methods, you must have a complete backup of your entire system. Now, let's take a closer look at each of these techniques.

Save system information in a nonrestricted state. The Save System Information (SAVSYSINF) command performs a cumulative save of a subset of system data and objects saved by SAVSYS without requiring the system to be in a restricted state. SAVSYSINF is not a replacement for SAVSYS and is not for use in system upgrades or migrations.

After you perform a base SAVSYS, SAVSYSINF saves the following:

- System objects: job descriptions, job queues, subsystem descriptions, and change commands
- System reply lists, service attributes, environment variables, system values required for system recovery, and network attributes
- Operating system PTFs that are copied into *SERVICE
  - Use the Change Service Attributes (CHGSRVA) command to modify your service attributes to automatically copy the PTF save files to *SERVICE when loading PTFs

For system recovery, recover the LIC and operating system from your SAVSYS media. Then, use your SAVSYSINF media and the Restore System Information (RSTSYSINF) command to restore the saved changes to system objects and PTFs.

Save data concurrently using multiple tape devices. To reduce downtime, perform save operations on more than one tape device at a time. For example, you can save libraries to one tape device, folders to a second tape device, and directories to a third tape device. Or you can save different sets of libraries, objects, folders, or directories to different tape devices.

Save data in parallel using multiple tape devices. A parallel save is intended for very large objects, libraries, or directories. With this method, the system "spreads" the data in the object, library, or directory across multiple tape devices.

Use the Save While Active process. Save While Active (SWA) can significantly reduce the amount of time your applications are unavailable and increase user access to applications and data. With SWA, users can resume activity after the save processing reaches a synchronization checkpoint.

The simplest way to use the SWA feature is to prevent user access to applications and data until the SWA checkpoint is reached. At this point, any exclusive locks are released, and users can resume their normal activity while the system continues to perform the save. Especially with large files, it takes significantly less time to reach the SWA checkpoint than to actually save the objects depending on the number, not the size, of the objects.

Starting with IBM i 6.1, the SWA function offers a single Save While Active checkpoint for multiple saves. The Start Save Synchronization (STRSAVSYNC) command ensures a single, consistent checkpoint for your library and IFS saves or for multiple concurrent library saves. If you use SWA, make sure you understand the process and monitor for any synchronization checkpoints before making your objects available for use.

2 - Using Only "Special" Backups to Test Recovery

Failing to use regular backup tapes is a huge mistake. If you perform an option 21 save only for the recovery test, you are ensuring that you have a complete backup to use to recover data. Performing a special backup to test your recovery is asking for trouble—if normal monthly, weekly, and daily backups have problems and you can't rely on them for tests, there is no way you will ever recover with these backups in a real disaster situation.

1 - No Testing Recovery Strategy

Without testing, you can't know if you can recover your systems, and limited recovery tests don't tell you what happens in reality.

If you really think your backups are good, are you confident that your organization could completely recover its system right now? If not, why? No matter how comprehensive you believe your backup is, you will never know if it works unless you actually test it. To truly verify your backup, you must test your recovery.

A Better Alternative to Tape-Based Disaster Recovery

Tapes don't let you back up continuously and have limits to physical access in disasters. Because of this, many businesses—especially those that have experienced actual disasters—have adopted logical replication, the most popular solution for both high availability (HA) and disaster recovery on the IBM i. Deployed through a high-availability ISV solution package or IBM's iCluster product for IBM i, Logical replication makes and keeps the objects on your production and backup systems identical. By journaling the objects, the transactional operations on the source system are duplicated on the target system by applying journal changes. For data that is not journaled, the changed data is saved and then written on the target. The logical replication solution provides these apply processes on the target.

The replication is near or in real-time for all journaled objects. Typically, when an object is journaled, replication is handled at a record level. For objects that aren't journaled, such as user spaces, replication is usually handled at the object level. In this case, the entire object is replicated after the completion of each set of changes to the object.

The best method for achieving efficiency and reliability using logical replication is synchronous remote journaling. With remote journaling, the IBM i operating system continuously moves data in the journal receiver to the backup system journal receiver. At this point, the selected software solution replays the journal updates, placing the updates into the object on the backup system. Once you've configured the journaling environment, you'll have two identical objects—one on the production system and one on the backup system.

In the event of a disaster, this solution lets you rapidly activate your production environment on the backup system with a role-swap operation. The biggest advantage of using logical replication for your recovery is that your data is live on a backup system, requiring minimal recovery procedures when you switch to the backup. You can also access the backup system to perform daily backups and to do read operations.

The disaster recovery failover process typically takes less than 30 minutes, and some users have achieved a 10-minute failover. Switching the direction of journal replication happens in less than a minute, and the database becomes accessible for reads even while it is serving as a backup. Less than 30 minutes for the failover is the pessimistic view of the possible lag in journal applies

that might have accumulated before the failover. The 30-minute estimate also allows you time to perform a synchronization check of the data before you give users access to the system. In summary, logical replication has the shortest failover time because the data on the backup system is already active and available for use when the applications are brought online.

An important consideration with logical replication is the possible latency of the replication. Latency is the amount of lag between the time when changes are made on the source system and the time when those changes become available on the backup system. Synchronous remote journaling mitigates this to a large extent. Regardless of the transmission mechanism you select, it is critical that you adequately project your transmission volume and properly size your communication lines and speeds to help ensure that your environment can manage replication volumes when they reach their peak, such as at month-end processing. A logical replication solution typically will fail to meet your needs when the communications sizing is inadequate for the volume of data being replicated.

CHAPTER 3: WHAT IS YOUR BACKUP LEAVING UNCOVERED?

By Dan Riehl

Many IT organizations run a backup process because it was prescribed by our software vendor or IBM business partner. But your business partners don't always know what you need backed up. Relying on cookie-cutter recommendations could cost you. How can you check to see if you are backing up the system correctly for your organization—and that you'll be able to recover after an outage?

Are You Backing Up Your User Libraries?

IBM provides a command that lets you to review the backup status of user libraries. It includes information on the last save date and whether the library has been changed since it was last saved. You can use the command Display Backup List (DSPBCKUPL) to get a listing on your display screen or a printed report.

Here's the command to get a screen listing and the resulting screen:
DSPBCKUPL BCKUPL(*LIB)

Display Library Backup List
System: MYSYSTEM
Find library . . . . . Starting characters
Type options below, then press Enter.
5=Display library contents 8=Display details

		----------Backup----------			Last
Opt	Library	Daily	Weekly	Monthly	Backup	Changed
	CBXLIB	Yes	Yes	Yes	01/31/11	No
	CLASSROOM	Yes	Yes	Yes	01/31/11	No
	CLCLASS	Yes	Yes	Yes	01/31/11	No
	CONCEPTS	Yes	Yes	Yes	01/31/11	No
	DANWORK	Yes	Yes	Yes	01/31/11	Yes
	DBU80	Yes	Yes	Yes	01/31/11	No
	F_GLT	Yes	Yes	Yes	01/31/11	No
	FB400	Yes	Yes	Yes	01/31/11	Yes
	FB400D	Yes	Yes	Yes	01/31/11	No
	GUYWORK	Yes	Yes	Yes	01/31/11	No
	GWRKSPLF	Yes	Yes	Yes	01/31/11	No

Are You Backing Up Your Folders and Documents in QDLS?

You can use the DSPBKUPL command to also list the last save dates of your Folder by specifying:

DSPBCKUPL BCKUPL(*FLR)

Display Folder Backup List

System: MYSYSTEM

Find folder . . . . . Starting characters

Type options below, then press Enter.

5=Display documents 8=Display next level

		----------Backup----------		Last	
Opt	Library	Daily	Weekly	Monthly	Backup	Changed
	*NONE	Yes	Yes	Yes		Yes
	PS	Yes	Yes	Yes		Yes
	QDIADOCS	Yes	Yes	Yes		Yes
	QFOSDIA	Yes	Yes	Yes	10/30/07	Yes
	QGA400RT	Yes	Yes	Yes		Yes
	QIWSADIM	Yes	Yes	Yes		Yes
	QJRN400	Yes	Yes	Yes		Yes
	QOTTMFLR	Yes	Yes	Yes		Yes

Do You Have a Recent Backup of the Operating System? (SAVSYS)

When you load a new release of the operating system and load PTFs, you will typically follow the IBM upgrade directions and perform a Save the System (SAVSYS) command to save a copy of your operating system. If you periodically load PTFs or make changes to the operating system, save your changes with a SAVSYS operation in case you need to re-load the operating system. But when was the last time you saved the operating system?

IBM does not mark QSYS objects saved when a SAVSYS operation is performed. Instead IBM updates the last saved date of a special data area named QSAVSYS in library QSYS. In order to determine your last SAVSYS, review the last saved date of the QSAVSYS data area, as in:

DSPOBJD OBJ(QSYS/QSAVSYS) OBJTYPE(*DTAARA)

Following that, review the last save date of the data area.

```
Display Object Description - Full
Library 1 of 1
Object . . . . . .   QSAVSYS      Attribute . . . . . .
Library . . . . . .  QSYS         Owner . . . . . .    QSYS
Library ASP device . .  *SYSBAS   Library ASP group .  *SYSBAS
Type . . . . . .     *DTAARA      Primary group . . .  *NONE

Journaling information:
Currently journaled . . . . . .   NO
Save/Restore information: . .
Save date/time . . . . . .        12/02/10  22:21:08
Restore date/time . . . . . .     12/08/10  22:18:55
Save command . . . . . .          SAVOBJ
Device type . . . . . .           Optical
Volumes . . . . . .               B2929_
File label ID . . . . . .         /Q5761SS1/Q61000M
```

When Did You Last Save Your User Profiles? (SAVSECDTA, SAVSYS)

Review the last save date of the QSAVUSRPRF data area, as in:

DSPOBJD OBJ(QSYS/QSAVUSRPRF) OBJTYPE(*DTAARA)

```
Display Object Description - Full
Library 1 of 1
Object . . . . . .   QSAVUSRPRF   Attribute . . . . . .
Library . . . . . .  QSYS         Owner . . . . . .    QSYS
Library ASP device . .  *SYSBAS   Library ASP group . .  *SYSBAS
Type . . . . . .     *DTAARA      Primary group . . . .  *NONE

Journaling information:
Currently journaled . . . . . .   NO
Save/Restore information . . . .
Save date/time . . . . . .        01/31/11  22:18:49
Restore date/time . . . . . .
Save command . . . . . .          SAVSECDTA
Device type . . . . . .           Save file
Save file . . . . . .             QTEMP/##SECDTA
```

Other QSAV* QSYS Data Areas to Determine Last Save Dates

SAVCFG	QSAVCFG
SAVLIB *ALLUSR	QSAVALLUSR
SAVLIB *IBM	QSAVIBM
SAVLIB *NONSYS	QSAVLIBALL
SAVSECDTA	QSAVUSRPRF
SAVSTG	QSAVSTG
SAVSYS	QSAVSYS, QSAVUSRPRF, QSAVCFG
SAVSYSINF	QSYSINF

IFS Directories and Files

I know of no good way to show the last save date for IFS objects. IBM recommends that you keep the output of your SAV command in a printed report; you can do this by specifying OUTPUT(*PRINT) on the SAV command. You can also place the output of SAV operations in a stream file, which you can then parse using a script or HLL program to track last save dates. See the help text on the SAV command.

Know All, See All

The worst time to discover that your backup is missing crucial objects is in the middle of a recovery operation. The only way to know you've got the goods for a future restore from backup is to create an itemized list of required objects and verify that they can be restored from your backup media. The best practice is to automate this kind of backup auditing on a continuous basis, which you can do via home-brew scripting or third party utilities. However you do it, you should start as soon as possible. The job you save will be your own.

CHAPTER 4: DO YOUR BACKUPS KEEP UP WITH OS UPDATES?

By Dan Riehl

As system administrators, one of our responsibilities is to ensure that our backups include all the information that may be needed to recover the entire system in the event of a catastrophic failure. In recent releases of the i OS, IBM has added several enhancements for such backup and recovery. But in order to take advantage of many of them, we need to update our backup policies and processes.

On a recurring schedule, my organization typically backs up

- Application and user libraries
- Security data (including user profiles)
- System configuration data
- Some third party vendor-supplied libraries, folders, and documents in the QDLS file system
- Other directories and stream files in the IFS

A basic daily backup process may save only the system changes made after the last full backup. I advise running a daily save process that backs up all user data, security data, and configuration data, instead of just those objects that have changed.

My organization usually saves the operating system (using SAVSYS) only after installing a new OS release or after significant changes have been made to the OS—for example, the application of a cumulative PFT package.

Obviously, your backup policy and processes may differ significantly depending on your unique requirements.

A Full Save

When you run SAVE Menu Option 21 (or your custom full backup process), are you really getting a full backup of your system? A more comprehensive option would be a combination of SAVSYS, SAVLIB LIB(*NONSYS), SAVDLO, and SAV to craft their own full backups.

New Capabilities

New operating systems—OS/400 5.3 and 5.4, and i OS 6.1 and 7.1—have enhanced our ability to back up information that in many cases was not available before. Take advantage of these new SAVE capabilities by updating your backup process.

Saving Access Paths

In OS/400 5.3, IBM introduced the system value QSAVACCPTH, which provides a system-wide default setting for saving access paths (e.g. indexes) when you save your database files. The system value has a shipped setting of *YES. Check your backup processes to ensure that all the SAVxxx commands (SAVLIB, SAVOBJ, SAV, etc.) specify the attribute ACCPTH as *YES or *SYSVAL. In a recovery scenario, your restore process—RSTLIB, RSTOBJ, RST, etc.—runs significantly faster if you restore the access paths instead of rebuilding them, which can be a lengthy process.

Check all of the SAVxxx commands in your backup process to ensure you are saving access paths, as in the following:

```
SAVLIB LIB(MYLIB) DEV(TAP01) ACCPTH(*SYSVAL or *YES)
```

For a full backup, make sure to save the access paths. It makes the SAVxxx process longer and uses more tape space, but if you're faced with a recovery scenario, you'll be glad you did it.

Spooled File Data

Prior to OS/400 5.4, only the description of the output queue objects was saved—not the reports in the output queue. Starting with 5.4, you can save and restore your spooled files residing in the output queues. Before 5.4, you could write your own report archive utility, but in 5.4 this capability is built into the SAVxxx and RSTxxx commands.

The following command saves all objects in the PRODLIB library; and also saves all spooled files in all output queues residing in the library:

```
SAVLIB LIB(PRODLIB) DEV(TAP01) SPLFDTA(*ALL)
```

For a full backup, make sure to back up your spooled file data. Again, this lengthens the SAVxxx process and uses more tape. But if you want to restore your reports, you have to save them.

Data Queue Data

Prior to 5.4, when a data queue (*DTAQ) object was saved, only the description of the data queue was saved—not the entries in the data queue. 5.4 added the capability to save all data queue entries.

The following command saves the PRODLIB library and also saves any data queue entries in all regular data queues. DDM data queue data cannot be saved using this value:

```
SAVLIB   LIB(PRODLIB)  DEV(TAP01)  QDTA(*DTAQ)
```

For a full backup, make sure to back up your data queue entries.

Private Authorities

With i OS 6.1, you can save private authorities with objects when you save them. Historically, private authorities would only be saved when using the commands SAVSECDTA (Save Security Data) and SAVSYS (Save System). Because private authorities are stored in user profiles, it's quite a neat trick for IBM to save and restore private authorities when simply saving and restoring the objects.

The following command saves the PRODLIB library and also saves all object private authorities.

```
SAVLIB   LIB(PRODLIB)  DEV(TAP01)  PVTAUT(*YES)
```

I addressed new 6.1 support in the article *Saving and Restoring Private Authorities*.

For a full backup, make sure to back up your private authorities using the appropriate SAVxxx command and specifying PVTAUT(*YES), or by using the SAVSECDTA command, which will save all user profiles along with their private authorities. SAVSYS can also be used to save the system (including user profiles and their private authorities).

Tomorrow's Backup Today

Right now your applications may not need to restore access paths, spool files, data queues, or private authorities. But future applications may depend on these objects retaining their state in an application or system restore. Since the latest versions of IBM i OS support saving these objects, work them into your current processes so that future applications will have what they need to resume operation after a disaster.

CONCLUSION: TAKING THE DISASTER OUT OF DISASTER RECOVERY

For many IBM i professionals, DR planning is an onerous chore that gets pushed to the back burner far too often. The result is that unexpected events that need restoration from backup can turn into unmitigated disasters. But a systematic approach to DR planning and implementation can ensure that you're backing up the right data at the right time.

FOR MORE INFORMATION

Contact Help/Systems at **1-800-328-1000** or **info@helpsystems.com** to set up a personal consultation to review your current setup and see how the Robot products can help you achieve your Disaster Recovery goals.