

Integrated File System: Advanced Topics

S4i Systems presentation



Power your planet.

© 2010 IBM Corporation

Overview of What's New in the...



IBM i 6.1

- **Collecting and analyzing folder attributes with System i Navigator**
 - Alternative to Retrieve Directory Information (RTVDIRINF) and Print Directory Information (PRTDIRINF) commands
- **Support additional Unicode characters**
 - Automatic conversion of directories that are not case sensitive to support additional characters
- **Support IPv6 connections**
 - QNTC, QFileSvr.400, NFS, and IBM i NetServer support both IPv4 and IPv6 addresses
- **QNetWare file system**
 - QNetWare file system is no longer supported



Overview of What's New in the... (continued)

IBM i 6.1

- **Unicode enablement of some CL commands**
 - Some CL commands are now Unicode-enabled
- **Subtree support added to more commands**
 - CHKIN, CHKOUT, RMDIR now have the SUBTREE parameter
 - System i Navigator updated to check in or check out all objects in a folder and subfolders
- **Changed copy related commands for more flexibility of permissions for created objects**
 - CPY, CPYFRMSTMF, CPYTOSTMF, CPYFRMIMPF, CPYTOIMPF
- **And more...**



Integrated File System: Advanced Topics

Contents

▪ Topics Covered

- Overview
- Auxiliary Storage Pools and User-defined File Systems
- Performance
- Retrieve and Print Directory Information
- Reclaim
- Virus Scanning
- Copying
- Journaling
- Other Useful Features
- Release Enhancements

Are There Other Topics of Interest?



Overview

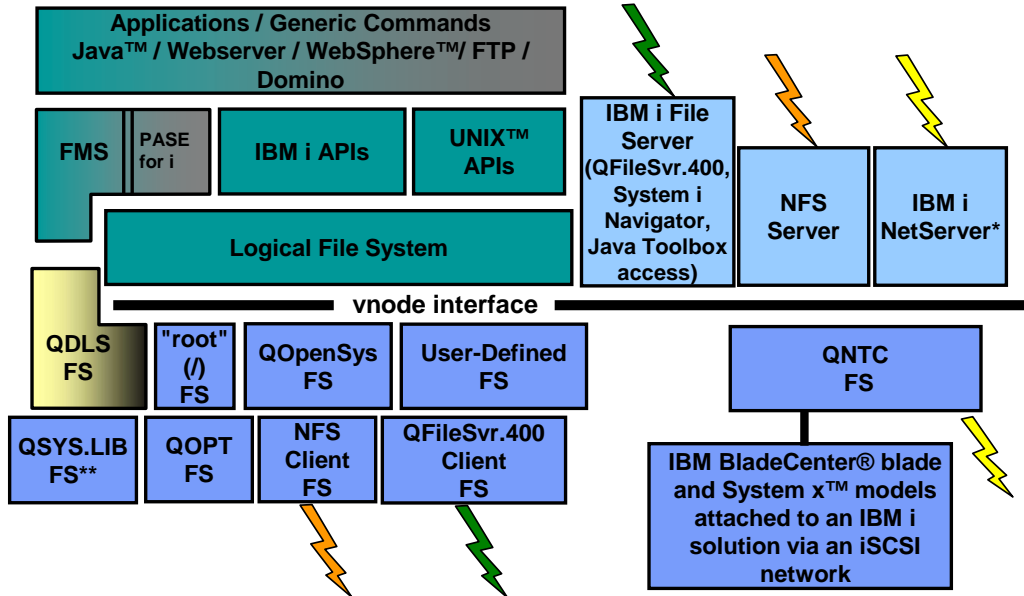
Overview

Integrated File System

- The Integrated File System is a...
 - **Common IBM i interface** that allows...
 - **Users** to access...
 - **Data** stored in various locations.



Under Your Kitchen Sink: The IFS “Plumbing”



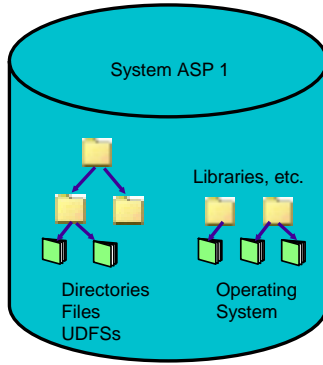
*IBM i Support for Windows™ Network Neighborhood and Samba

**Includes independent ASP QSYS.LIB file system

Auxiliary Storage Pools (ASPs) and User-defined File Systems (UDFSs)

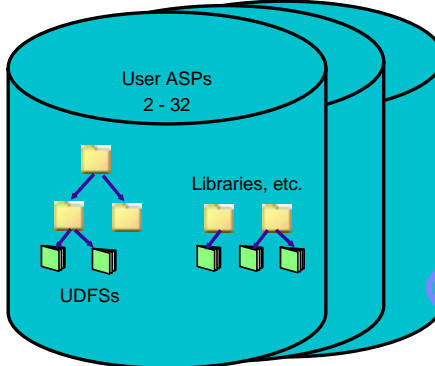
ASP & UDFSs

User ASP



System ASP

- Number 1
- Multiple disk units
- Always exists
- Contains operating system and user data



User ASPs

- Numbers 2 - 32
- One or more disk units
- Optional
- Contains Libraries and UDFSs
- Can overflow into system ASP

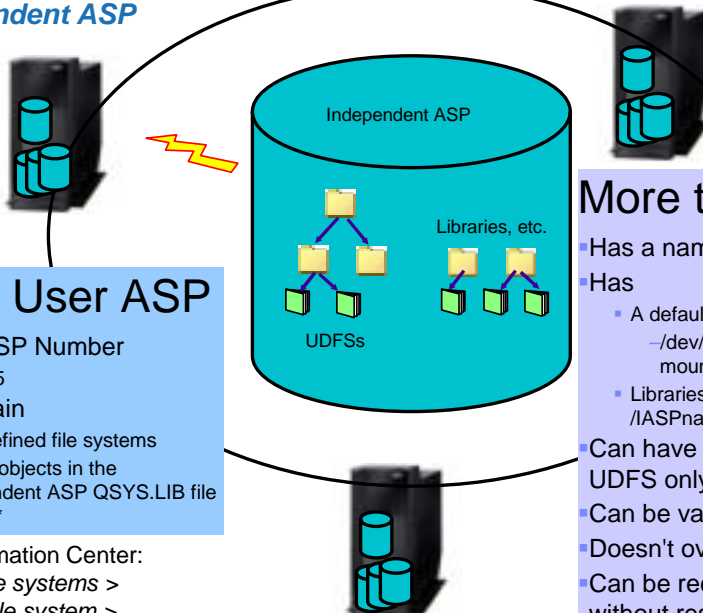
- Can improve performance***
 - Dedicated disks

IBM i Information Center:
 Systems management >
 Disk management >
 Disk pools

*Measured in lab environments. Actual improvements in customer environments may differ.

ASP & UDFSs

Independent ASP



Like a User ASP

- Has an ASP Number
 - 33 - 255
- Can contain
 - User-defined file systems
 - Library objects in the independent ASP QSYS.LIB file system*

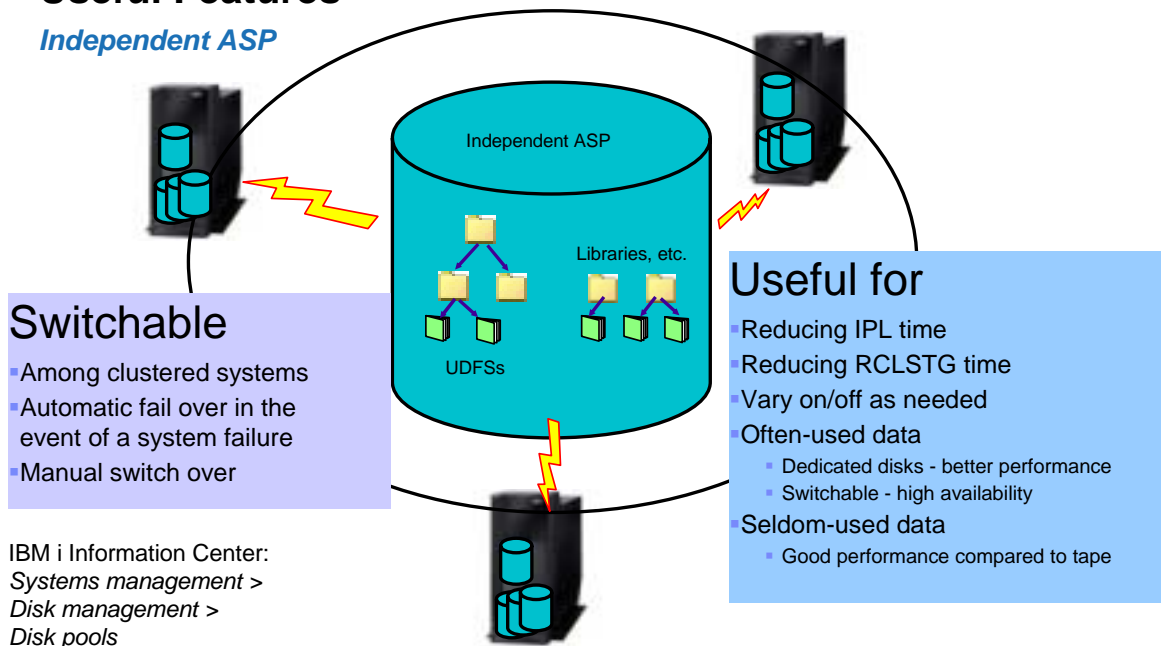
*IBM i Information Center:
 Files and file systems >
 Integrated file system >
 File systems >
 Independent ASP QSYS.LIB

More than a User ASP

- Has a name and a device description
- Has
 - A default user-defined file system
 - /dev/IASPname/QDEFAULT.UDFS is mounted at /IASPname
 - Libraries that are mounted at /IASPname/QSYS.LIB
- Can have Primary & Secondary or UDFS only
- Can be varied on and off
- Doesn't overflow to system ASP
- Can be reclaimed in parallel (RCLSTG) without restricted state
 - But the independent ASP can't be in use

Useful Features

Independent ASP



Switchable

- Among clustered systems
- Automatic fail over in the event of a system failure
- Manual switch over

IBM i Information Center:
Systems management >
Disk management >
Disk pools

Useful for

- Reducing IPL time
- Reducing RCLSTG time
- Vary on/off as needed
- Often-used data
 - Dedicated disks - better performance
 - Switchable - high availability
- Seldom-used data
 - Good performance compared to tape

ASPs & UDFSs

User-defined File System (UDFS)

- The only way to access directories and stream files in a user ASP or independent ASP
- Created by
 - You...as many as you want (Limit of 4,112 per independent ASP)
 - The system...it creates one called QDEFAULT.UDFS for each independent ASP
- Anchored under the /dev/AspName directory, where AspName is
 - "QASP01" ... "QASP32", for user ASPs
 - The name of the independent ASP, for independent ASPs
 - The AspName directory is created by the system for each user ASP that is defined, and for each independent ASP that is varied on (made available)
- Represented by a *BLKSF object (block special file)
 - This is a place-holder object that simply represents the UDFS
- Examples:
 - /dev/QASP02/images.udfs is a *BLKSF object representing a UDFS on ASP 2
 - /dev/ACCOUNTS/QDEFAULT.UDFS is a *BLKSF object representing the system-created UDFS on the independent ASP named ACCOUNTS
 - /dev/ACCOUNTS/active.udfs is a *BLKSF object representing a UDFS on the independent ASP named ACCOUNTS

ASP s & UDF S s

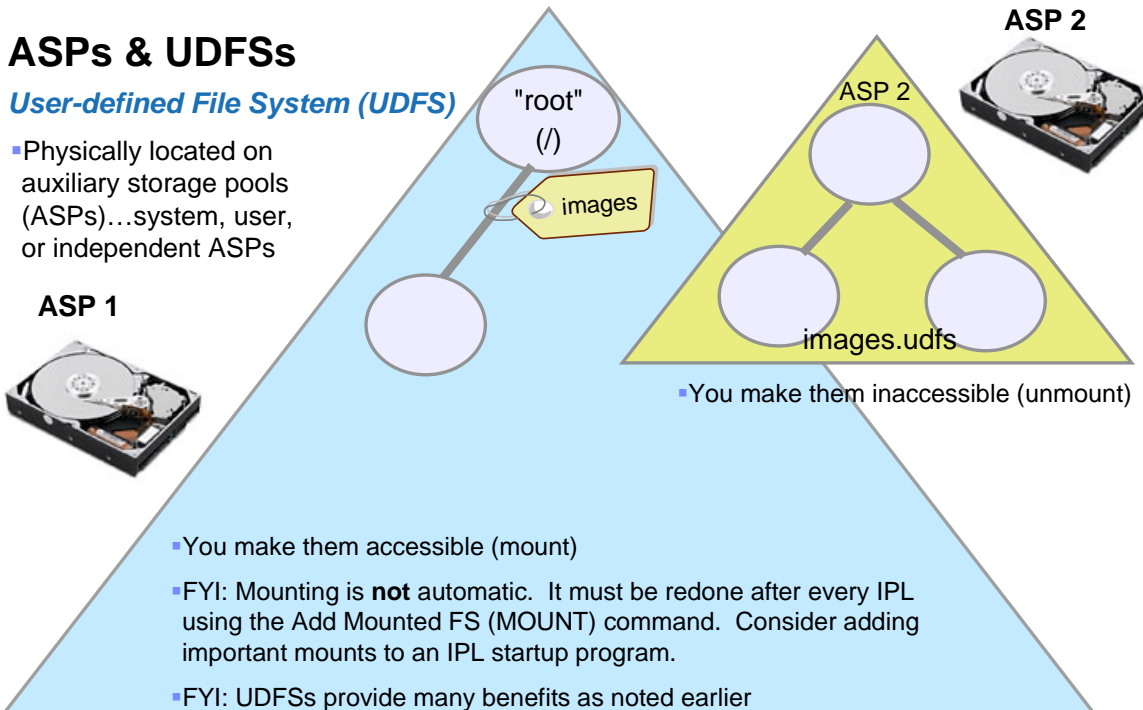
User-defined File System (UDFS)

- Before objects can be accessed in a UDFS, the UDFS must be "mounted"
- Mounting associates the UDFS with a particular directory
- Objects in the UDFS are accessed through that directory
- If an IPL of the system is performed, the UDFS must be mounted again.
- Examples:
 - When you mount `/dev/QASP02/images.udfs` over directory `/images`, you can access objects in that UDFS using the `/images/...` path name
 - When an independent ASP is varied on, the `/dev/AspName/QDEFAULT.UDFS` UDFS is automatically mounted over the `/AspName` directory
- You can unmount a UDFS to make the objects within it inaccessible
 - A UDFS can be unmounted only if the objects within it are not in use
- More information found in the IBM i Information Center at :
 - Files and file systems >*
 - Integrated file system >*
 - File systems >*
 - User-defined file systems (UDFSs)*

ASP s & UDF S s

User-defined File System (UDFS)

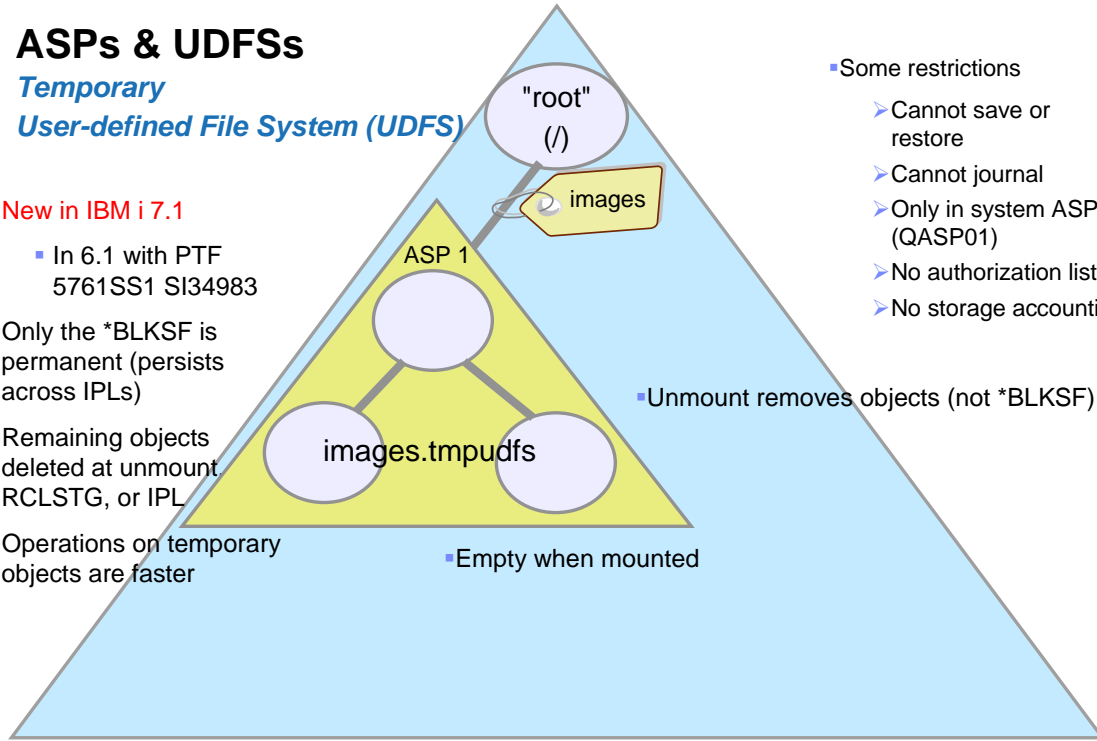
- Physically located on auxiliary storage pools (ASP s)...system, user, or independent ASP s



ASPs & UDFSs

Temporary
User-defined File System (UDFS)

- **New in IBM i 7.1**
 - In 6.1 with PTF 5761SS1 SI34983
- Only the *BLKSF is permanent (persists across IPLs)
- Remaining objects deleted at unmount RCLSTG, or IPL
- Operations on temporary objects are faster



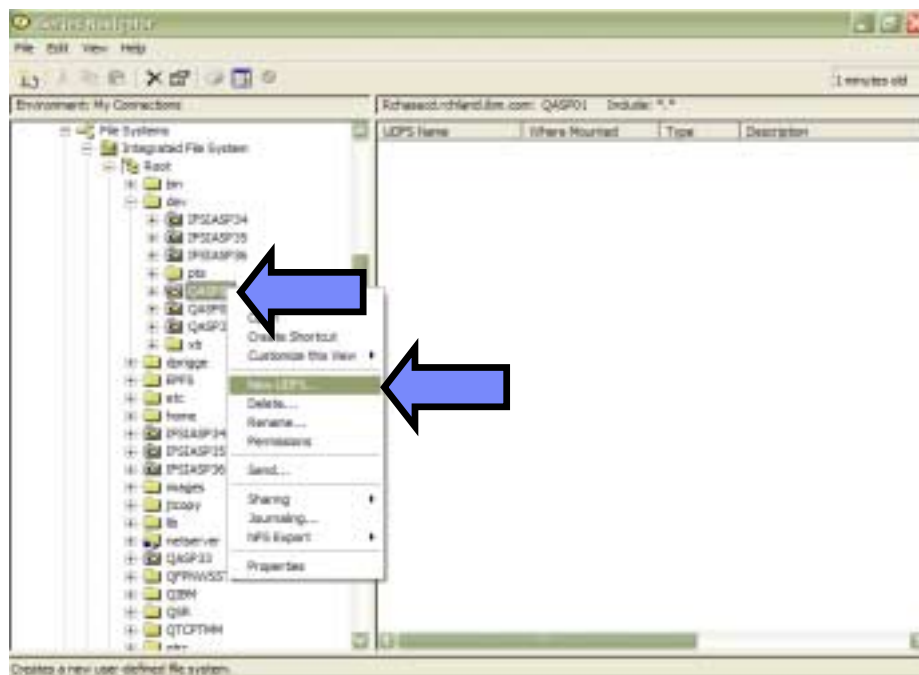
- Some restrictions
 - Cannot save or restore
 - Cannot journal
 - Only in system ASP (QASP01)
 - No authorization lists
 - No storage accounting

▪ Unmount removes objects (not *BLKSF)

▪ Empty when mounted

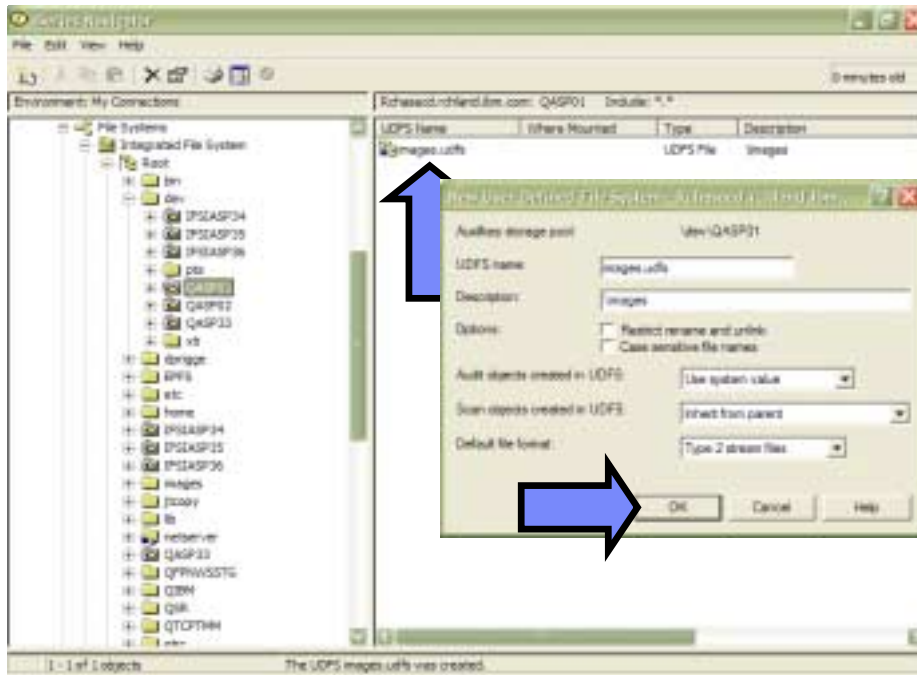
ASPs & UDFSs

Step 1: Create a UDFS



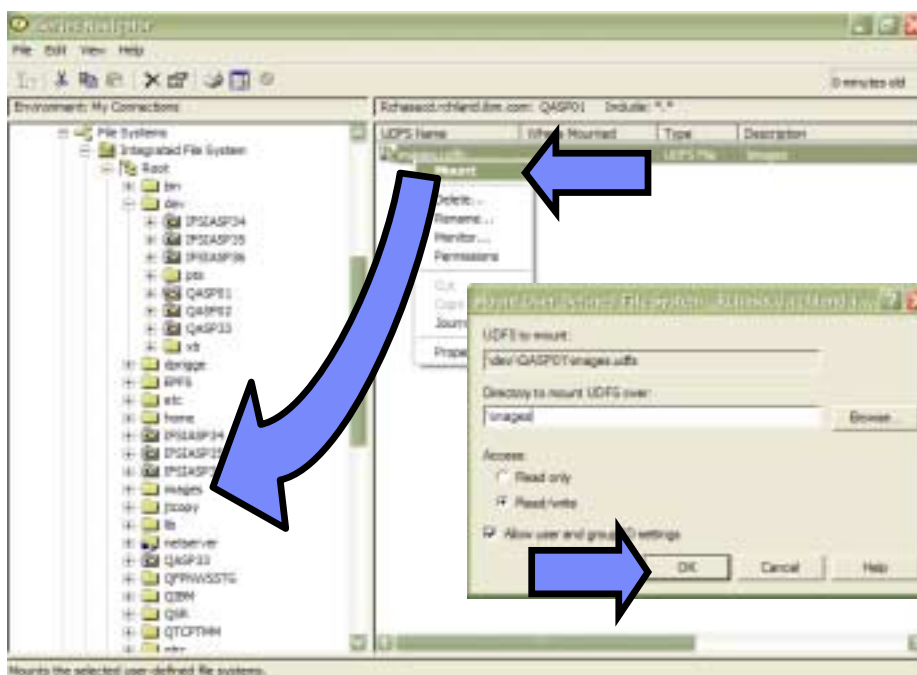
ASPs & UDFSs

Step 1: Create a UDFS



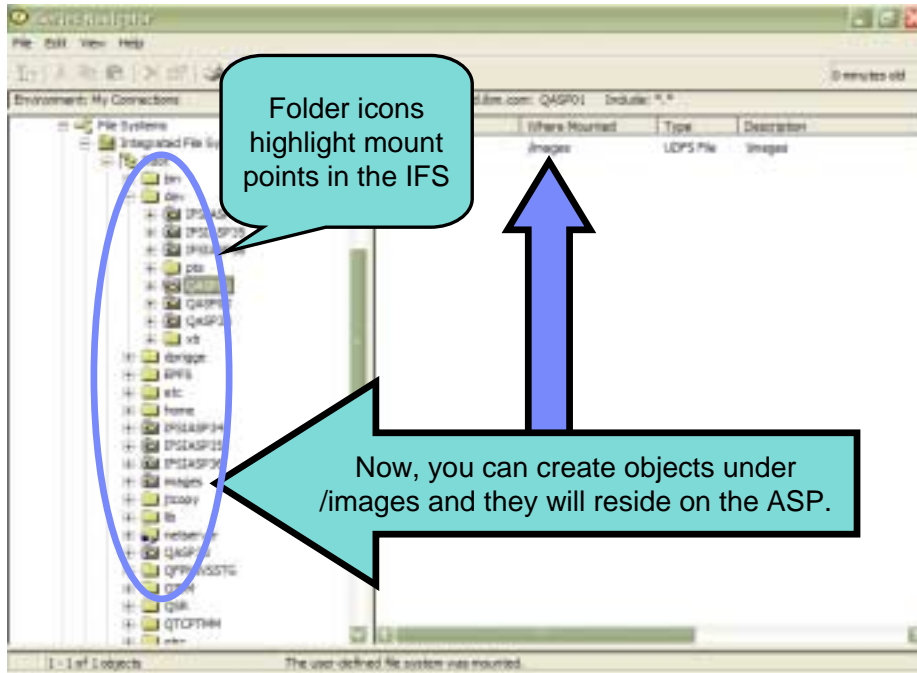
ASPs & UDFSs

Step 2: Mount the UDFS



ASPs & UDFSs

Step 3: View the UDFS Mount



Performance

Performance

Overview

- **Performance Improvement Tips**
 - General Tips
 - Directory Specific Tips
 - Stream File Specific Tips
- **Performance Improvement References**
- **Performance Analysis Topics**
 - Performance Explorer Data
 - IBM iDoctor for IBM i



Performance

General Performance Improvement Tips

- **Utilize UDFSs and ASPs** (See previous slides for performance details)
 - Consider temporary UDFS, if appropriate
- **Minimize File System Resource Contention**
 - More information on the following slides →
- **Stream-line User Journaling and Auditing Performance**
 - Omit (if applicable) the open, close, and fsync journal entries (OMTJRNE(*OPNCLOSYN)) on the Start Journal (STRJRN) command
 - Only user journal necessary data and only audit operations of interest since there is some overhead involved in both
 - More information in Performance Improvement References →
- **Streamline Save and Restore (S/R) Performance**
 - Minimize or eliminate auditing and scanning during S/R, if possible
 - More information in Performance Improvement References →



Performance

Minimize File System Resource Contention

What is it?

- Occurs when two or more threads request the same internal file system resource at the same time
- Often occurs when threads access the same file system data (links, directories, stream files, etc.)

Example on the following slide →


How to prevent it?

- Ensure multiple threads don't operate on the same file system data at the same time
- May require restructuring of data and/or applications



Performance

EXAMPLE: Minimize File System Resource Contention



Performance

Directory Specific Performance Improvement Tips

- **Use Short Link and Path Names**
 - Short link name (e.g. 'image_050107.jpg') is <= 16 characters
 - Short path name (e.g. '/archive/img/001.jpg') is <= 20 characters
 - Two Internal file system caching mechanisms designed for short names
 - Improve path name resolution performance by as much as 15%*
- **Use a Current Working Directory (CWD) and Avoid Deep Subtrees**
 - Both minimize path name resolution (i.e. get to your objects faster)*
 - Example on the following slide →



*Measured in lab environments. Actual improvements in customer environments may differ.

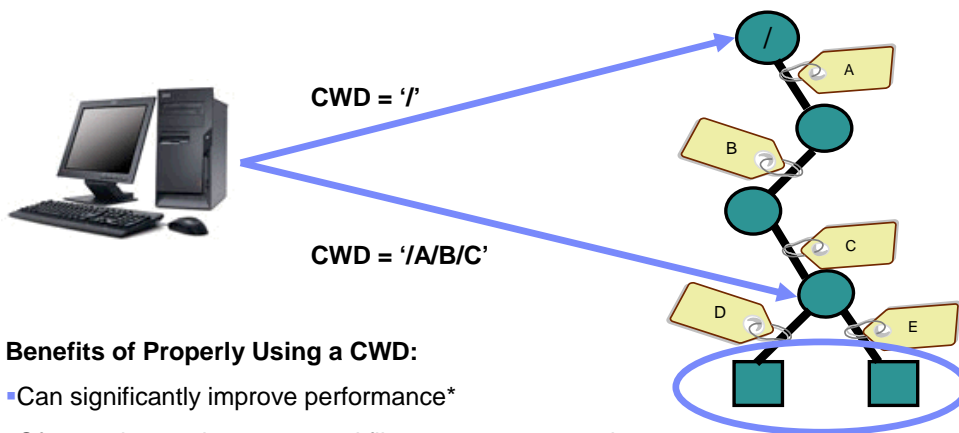
25

Power your planet.

© 2010 IBM Corporation

Performance

EXAMPLE: Why Use a Current Working Directory (CWD)?



Benefits of Properly Using a CWD:

- Can significantly improve performance*
- Often makes maintenance and file system usage easier

*Measured in lab environments. Actual improvements in customer environments may differ.

26

Power your planet.

© 2010 IBM Corporation

Performance

Stream File Specific Performance Improvement Tips

- **Use Memory Mapping APIs When Applicable**
 - Establishes a mapping between memory and a stream file for reads and writes
 - Improve read and write performance by as much as 4x*
 - Restriction: cannot perform memory mapping and user journaling on same file
- **Read and write data on a page boundary**
 - Current page size obtained via the `f_bsize` field returned on the `statvfs()` API
 - Can improve performance by minimizing the number of pages accessed*
- **Utilize the Disk and Main Storage Options**
 - More information on the following slides →
- **Minimize the Amount of Data Conversion Performed**
 - Data conversion is a valuable tool but it can degrade read and write performance
- **Optimize “On-Access” Virus Scanning**
 - More information on the following slides →



*Measured in lab environments. Actual improvements in customer environments may differ.

Performance

Disk and Main Storage Options



- **Example:** Using the main storage option (memory allocation) value of *MINIMIZE (minimize) for Domino databases on memory constrained systems has been shown in the lab to reduce page faults and improve response time*



*Measured in lab environments. Actual improvements in customer environments may differ.

Performance

Disk and Main Storage Options

▪ Disk Storage Allocation (*TYPE2 Stream Files Only)

- Determines how disk storage is allocated for a stream file
- Normal: Extend the file in increasingly larger blocks to reduce I/O. May take up more disk storage than is actually used.
- Minimize: Allocate what is needed. Conserves disk storage, but may cause more I/O and disk fragmentation.
- Dynamic: Determine how to allocate disk storage based on how the file is being used.

▪ Main Storage Allocation

- Determines how main storage is allocated as a stream file is being accessed
- Normal: Use as much main storage as is needed to cache what is being used to reduce I/Os.
- Minimize: Allocate what is needed. Conserves main storage, but may cause more I/O because less data is cached.
- Dynamic: Determine how to allocate main storage based on system activity and main storage contention.

▪ Changeable attributes

- System i Navigator *Properties > Storage tab*
- Qp0ISetAttr() API and Change Attribute (CHGATR) command



Performance

Optimize “On-Access” Virus Scanning

▪ Reduce the Number of Scans Required

- Set a stream file object's scanning attribute to “no scan” or “scan only if changed”
- Set the scan file system control system value to “scan accesses through file server only”
- WARNING: Reducing the number of scans can impact security

▪ Perform Scanning During Off Shift Hours

- Download new virus definitions and perform a system-wide scan
- Fewer scans during normal hours



Performance

Performance Explorer Data

Performance Explorer (PEX)

- Create via the Add PEX Definition (ADDPEXDFN) command
- Start collection data via the Start Performance Explorer (STRPEX) command
- <Do your stuff (e.g. Run you application)>
- End data collection via the End Performance Explorer (ENDPEX) command
- Print the data collected via the Print PEX Report (PRTPEXRPT) command
- View data collect manually or via iDoctor →

Operating System Event (OSEVT) parameter on ADDPEXDFN

- *IFSOPEN (open, create, close)
- *IFSIO (read, write, dup, lseek, fcntl, and more...)
- *IFSCOMP (file system kernel operations)

More information

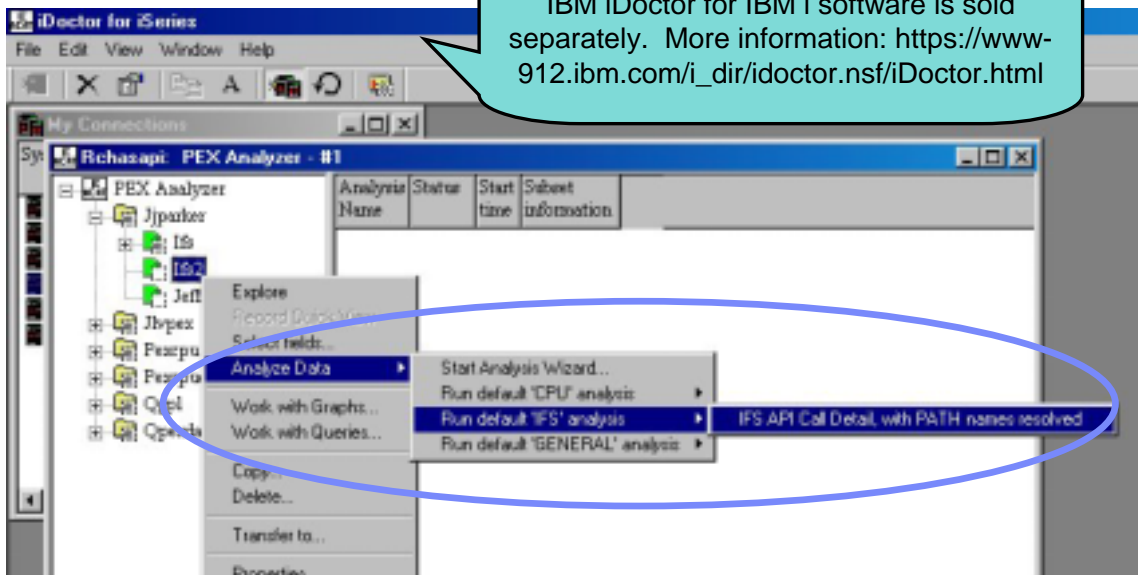
- IBM i Information Center at: *Systems management > Performance > Applications for performance management > Performance data collectors > Performance explorer*
- IBM iDoctor for IBM i website at: https://www-912.ibm.com/i_dir/idoctor.nsf



Performance

IBM iDoctor for IBM i

IBM iDoctor for IBM i software is sold separately. More information: https://www-912.ibm.com/i_dir/idoctor.nsf/iDoctor.html



Performance

IBM iDoctor for IBM i

The screenshot shows the IBM iDoctor for IBM i interface. A window titled 'Data Viewer [Richasapi] - #1' displays a table of IFS API call details. A callout bubble highlights the table with the text: "Data collected can help identify performance bottlenecks...which may be fixed by incorporating some of the performance tips discussed earlier."

IFS API File Descriptor	IFS API New File Descriptor	IFS API Requested New File Descriptor	IFS API Name	Error Number	IFS Path Name	IFS Path CCSID	IFS Path Language	IFS Path Country	Conversion Id	Number of Bytes Requested	Number of Bytes Received	Imp Off
0	0	0	open	0	/jiparker/peo/pefile	37	ENU	US	0	0	0	
0	0	0	write	0	/jiparker/peo/pefile	0			0	100	100	
0	0	0	write	0	/jiparker/peo/pefile	0			0	0	0	

Collect and Analyze Directory Information

Collect and analyze directory attributes

▪ Collect Overview

- Collect detailed attribute information for integrated file system objects which can be used to:
 - Assist in managing integrated file system security
 - Detailed monitoring of storage usage
 - Identify recently used objects
 - Much more...
- Information can be collected about objects in a specific directory subtree
- Information collected is placed by default into two database tables in QUSRSYS library
 - QAEZDxxxxO – attribute information for all object types
 - QAEZDxxxxD – information for parent directories
- Lengthy process, recommend not running against “root” (/), instead run against subdirectories

▪ Analysis Overview

- Information can be printed using the Print Directory Information (PRTDIRINF) CL command
- Object reports can be filtered (owner, size, etc.) and can be sorted by various attributes
- Attributes that are collected described in the IBM i Information Center: *Files and file systems >Integrated file system>Accessing the integrated file system >Accessing using CL commands >Working with output of the RTVDIRINF and PRTDIRINF commands*

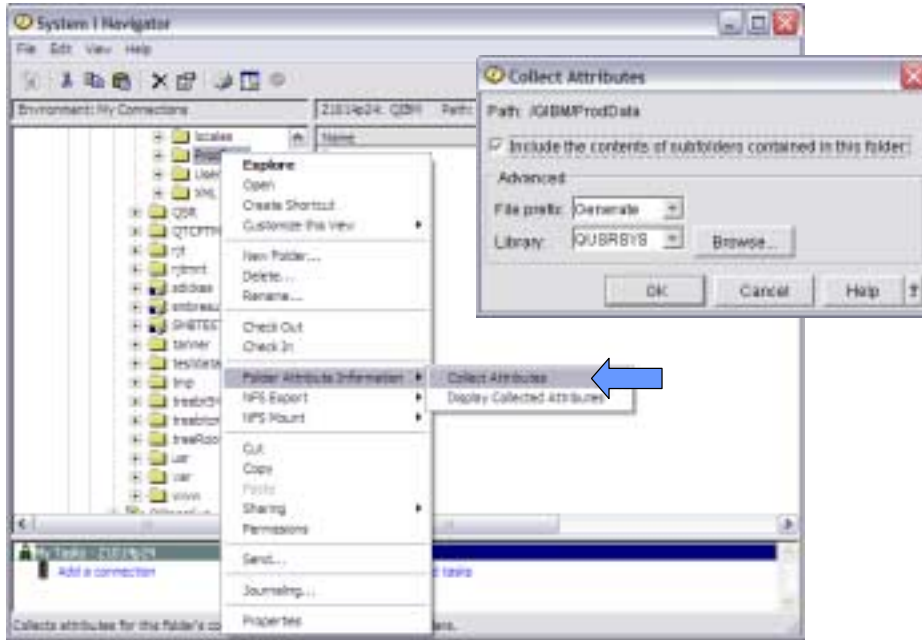
Collect/analyze directory information

- *Using System i Navigator (available starting with IBM i 6.1)*

- **GUI alternative to RTVDIRINF command**
- **Interfaces to:**
 - Collect attributes
 - Analyze information
 - Delete data when output is no longer needed
- **Information stored in database files on IBM i**
- **PRTDIRINF can be used to print reports**
- **Analyze Information**
 - Select fields to display
 - Filter out data
 - Order data

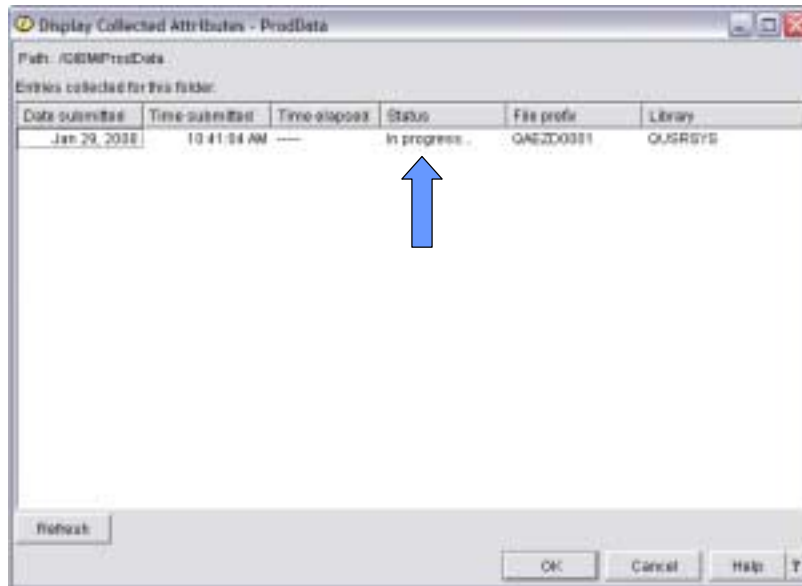
Collect/analyze directory information

System i Navigator collect attributes (available starting with IBM i 6.1)



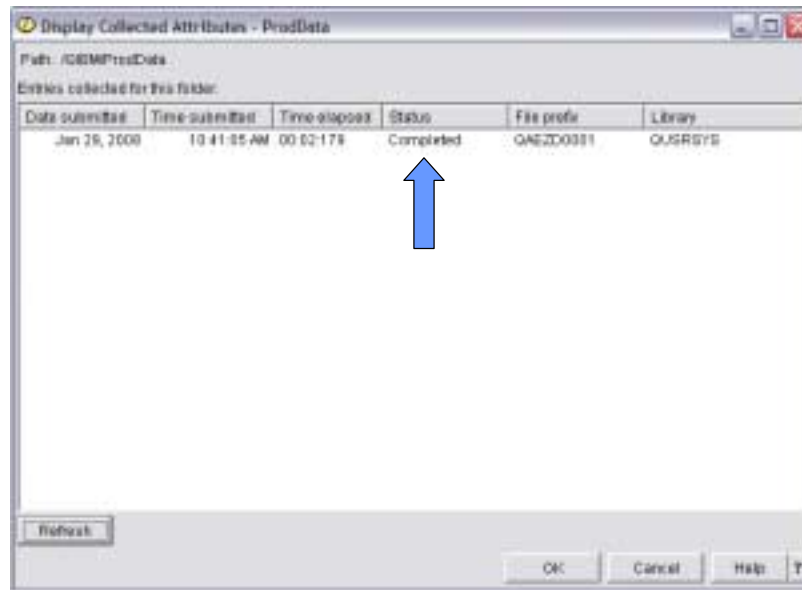
Collect/analyze directory information

System i Navigator collect attributes (available starting with IBM i 6.1)



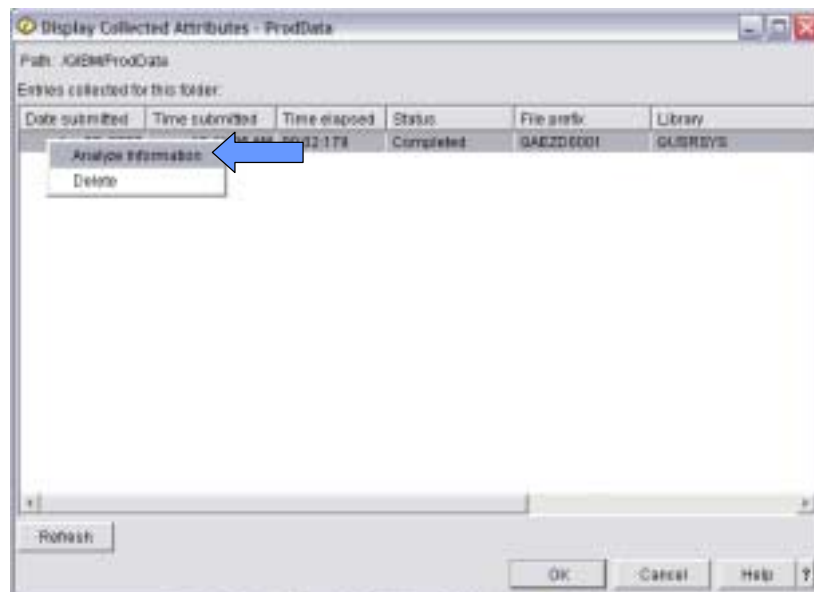
Collect/analyze directory information

System i Navigator collect attributes (available starting with IBM i 6.1)



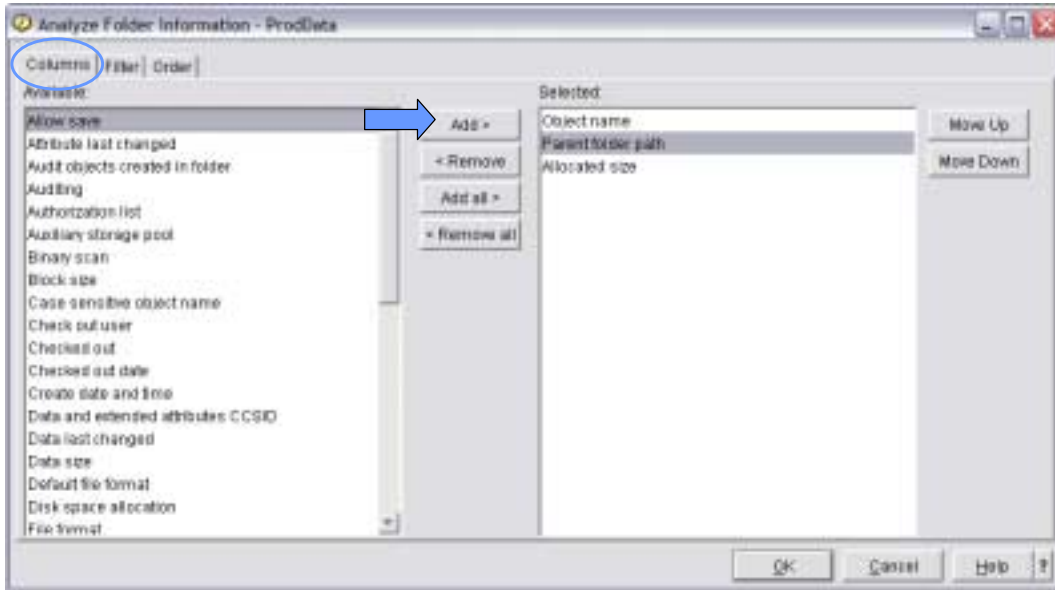
Collect/analyze directory information

System i Navigator analyze information (available starting with IBM i 6.1)



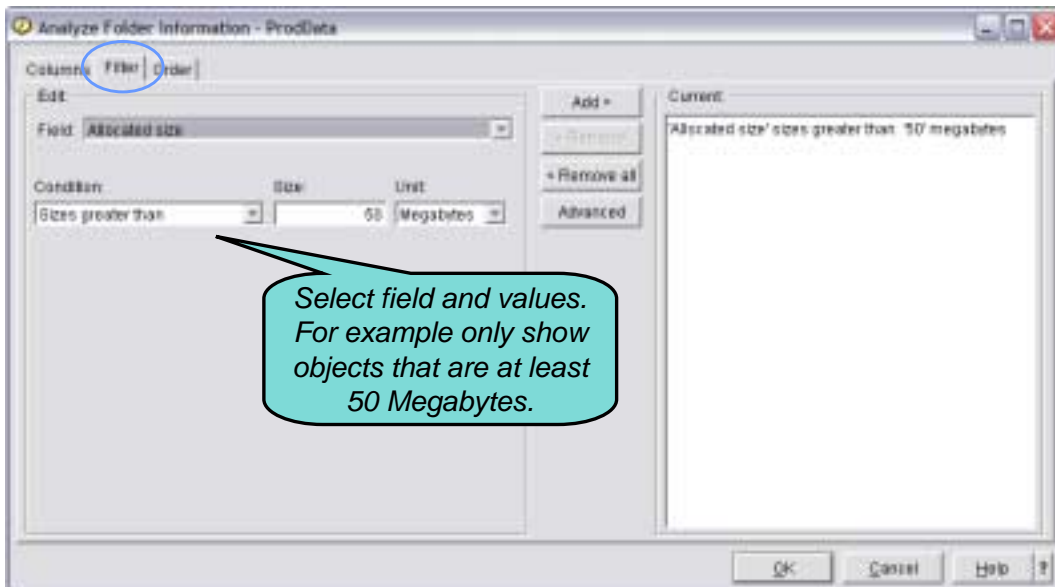
Collect/analyze directory information

System i Navigator analyze information (available starting with IBM i 6.1)



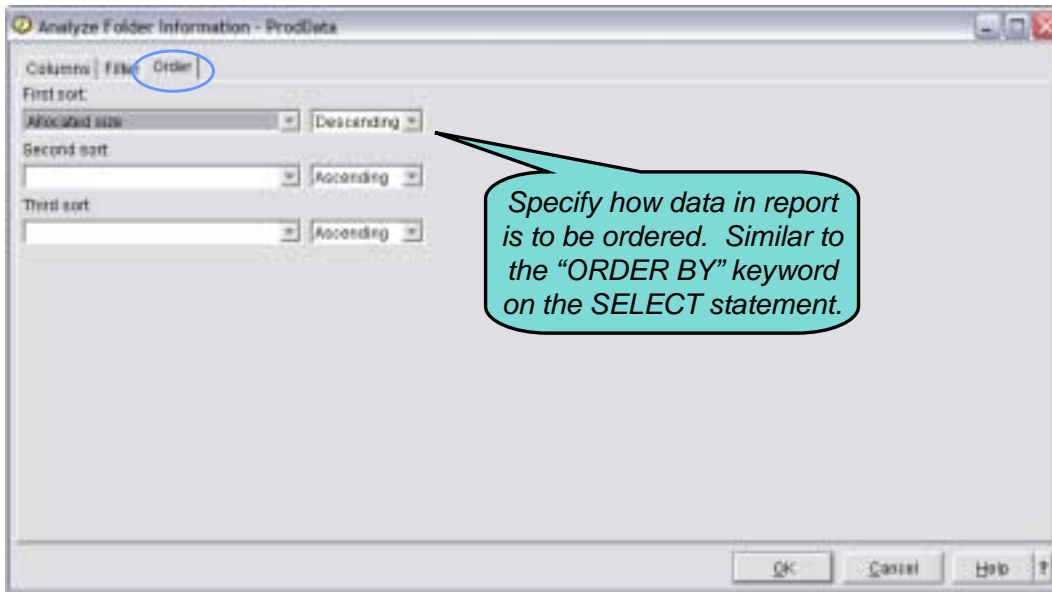
Collect/analyze directory information

System i Navigator analyze information (available starting with IBM i 6.1)



Collect/analyze directory information

System i Navigator analyze information (available starting with IBM i 6.1)

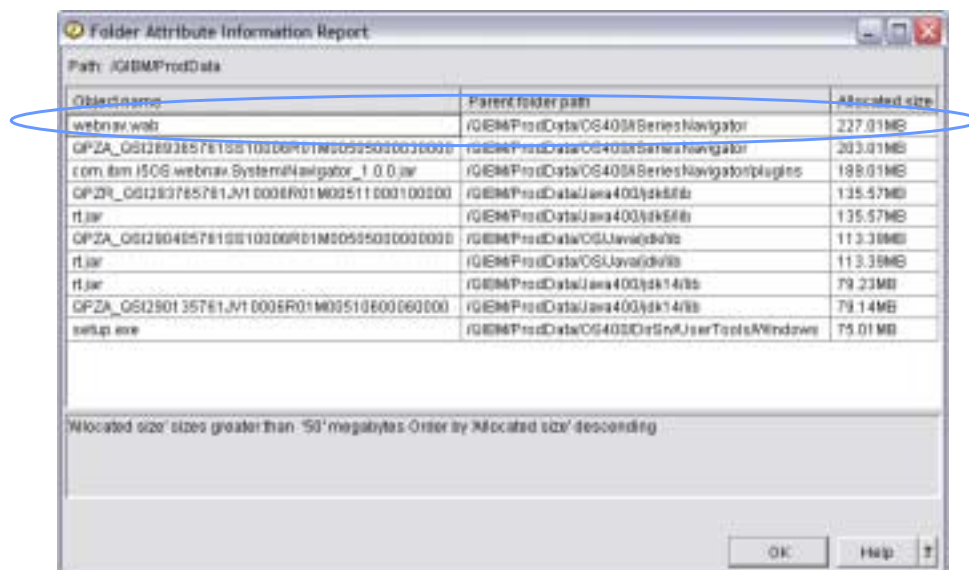


Collect/analyze directory information

System i Navigator analyze information (available starting with IBM i 6.1)

Report showing storage usage.

In this example webnav.wab is the largest file.



Retrieve and Print Directory Information

RTVDIRINF and PTRDIRINF Commands

RTVDIRINF

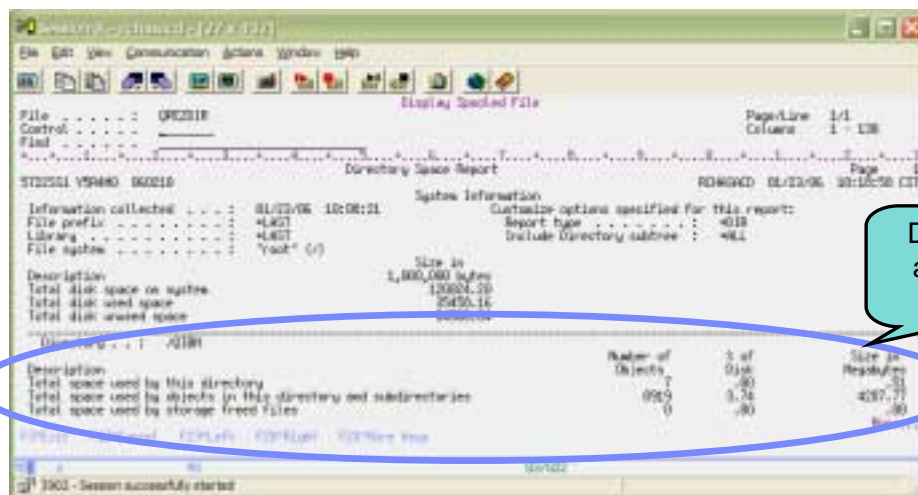
- Command interface to collect attribute information for objects in the IFS
- The output files can be used by System i Navigator analyze functions
- Each run produces a new file which is given a sequentially numbered name (represented by xxxx above.) resulting is a maximum of 9999 result files possible in any one library
- Best to run as batch job: SBMJOB CMD(RTVDIRINF DIR('<target>'))

PRTDIRINF Overview

- Summarize and print information collected by RTVDIRINF
- Object reports can be filtered (owner, size, etc.) and can be sorted by various attributes

Retrieve and Print Directory Information

Directory Space Report – RPTTYPE(*DIR)



Directory attribute totals

Reclaim

Reclaim Storage (RCLSTG) Command

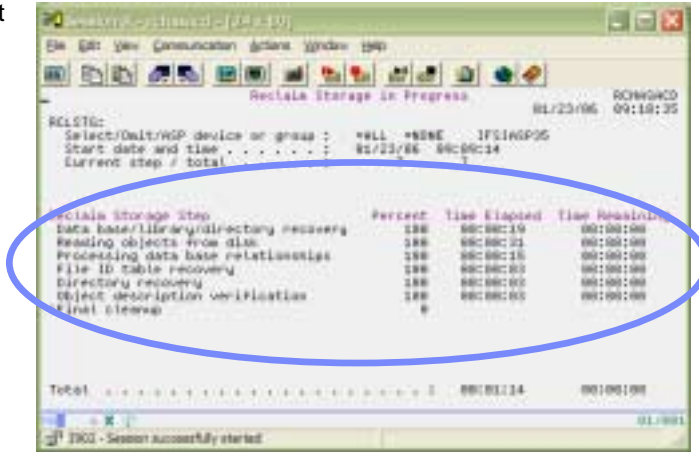
RCLSTG Overview

- Reclaims lost storage and lost objects
- Repairs problems in the file system, database, etc.
- Requires restricted environment



RCLSTG Features

- SELECT(*DIR) option
Perform directory (i.e. file system) reclaim only
- ESTIMATE parameter
Estimate reclaim duration
- Status screen
Monitor reclaim progress



Reclaim

Reclaim Object Links (RCLLNK) Command

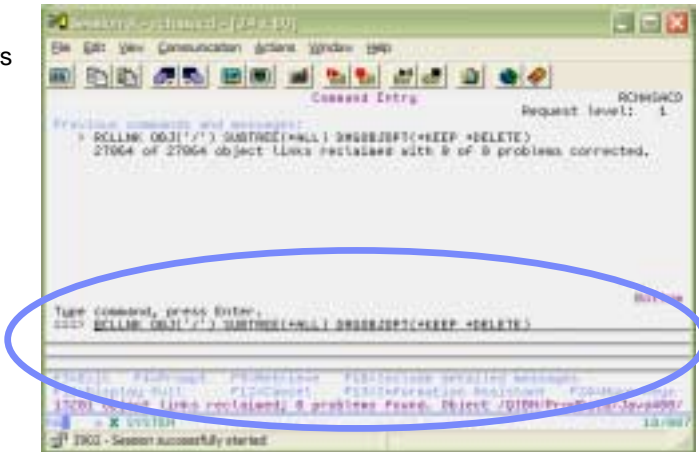
RCLLNK Overview

- Repairs "root" (/), QOpenSys, and user-defined file system problems
- Restricted environment **NOT** required



RCLLNK Features

- OBJ and SUBTREE parameters
Reclaim a single object -OR- Reclaim a group of objects
- DMGOBJOPT parameter
Specify how damaged objects are handled
- Status message
Monitor reclaim progress



Reclaim

Reclaim Command Comparison

Which Command Should I Use?

- In general and when applicable, use the RCLLNK command
- Use the RCLSTG command only when necessary
- The following comparison chart should help you answer this question →



RCLLNK	RCLSTG ASPDEV(*SYSBAS)
NO restricted state	Restricted state required
Identifies and corrects MOST known file system problems	Identifies and corrects ALL known file system problems
Objects are reclaimed on an individual or subtree basis	Objects are reclaimed on an system-wide basis
ALL file systems are usable throughout the duration of the reclaim	NO file systems are usable throughout the duration of the reclaim
Lost objects are NOT found	Finds and re-links lost objects

- More information in the IBM i Information Center at:
Files and file systems > Integrated file system > Reclaim the "root" (/), QOpenSys, and user-defined file systems

Virus Scanning

Virus Scanning

Overview


- Can IBM i objects be infected? – **YES!**
 - Objects writable by a PC can be infected and/or damaged
 - Any object can be deleted
 - Mapping a drive with all object authority exposes IBM i to attacks by PC viruses
- Scanning using IBM i NetServer – **DON'T!**
 - Uses up network resources
 - Moves data across network in the “clear”
 - Scanners can go into infinite loops
 - Updates file access timestamps
 - Usually scans everything, every time
- Scanning using scanner running on IBM i – **DO!**
 - No network resources used
 - Creates no additional exposure
 - No side effects
 - Faster scan
 - System “remembers” scan results, only re-scans when necessary



Virus Scanning

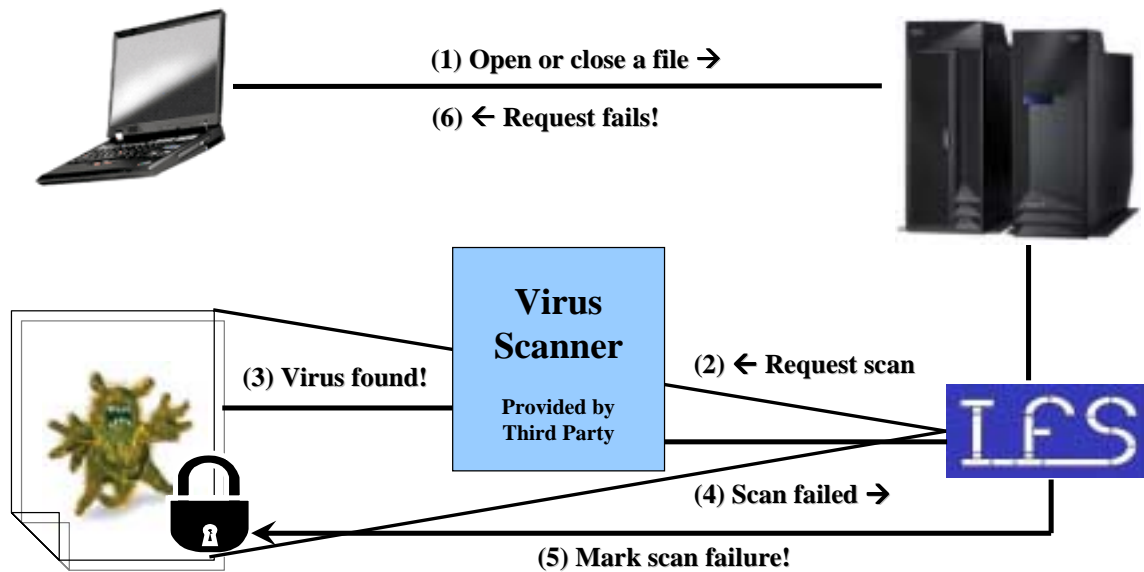
IBM i Scan Enablement: Overview

- IBM i enablement support to optimize virus scanning
 - System provided exit points on opens and closes of objects in the “root” (/), QOpenSys, and user-defined file systems
 - QIBM_QP0L_SCAN_OPEN
 - QIBM_QP0L_SCAN_CLOSE
 - System values to control the above exit point scanning
 - Scan file systems (QSCANFS)
 - Scan file systems control (QSCANFSCTL)
 - Object attribute to control whether objects are scanned or not. This attribute can be inherited from a parent directory’s setting.
 - Scan Yes
 - Scan No
 - Scan Only when Object is Changed
 - System “remembers” scan history and only re-scans if necessary
- This is only enablement!
 - An application is required to take advantage of this support
 - See Bytware, Raz-Lee.



Virus Scanning

IBM i Scan Enablement: How it works



Virus Scanning

IBM i Scan Enablement: Virus Scanning Products

■ StandGuard Anti-Virus from Bytware

- McAfee based
- More Information: <http://www.bytware.com/products/av/index.html>

■ iSecurity anti-virus product from Raz-Lee

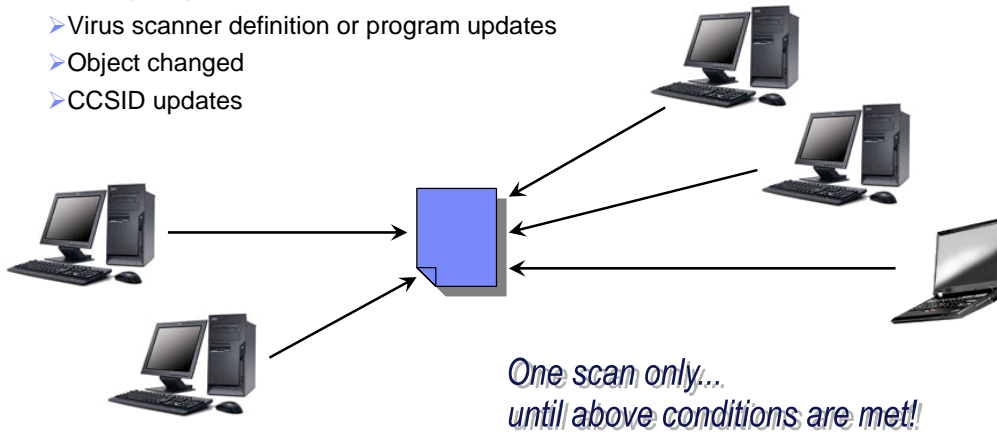
- Open-source based
- More Information: <http://www.razlee.com>



Virus Scanning


IBM i Scan Enablement: Object Scanning Attribute

- Scan object no
- Scan object only when changed
- Scan object yes
 - Virus scanner definition or program updates
 - Object changed
 - CCSID updates



Virus Scanning

IBM i Scan Enablement: Scan Security Policy



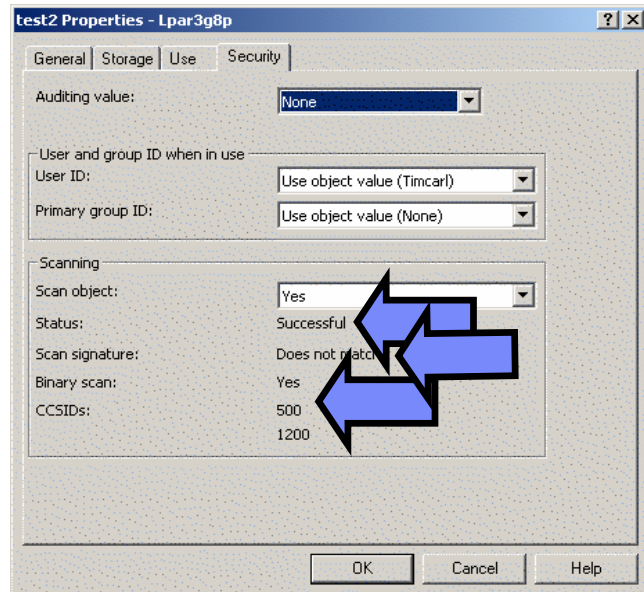
Found under: Security > Policies > Security Policy > Scan Tab

Virus Scanning

IBM i Scan Enablement: Scanning Properties and Status

Found under **Security** tab

- Object scanned
 - Successfully
 - In binary
 - In CCSIDs 500 and 1200
- Scan signature different
 - Object needs to be rescanned
 - Indicates a scanning update such as new virus definitions loaded
- More information found in the IBM i Information Center at *Files and file systems > Integrated file system > Integrated file system concepts > Scanning support*



Virus Scanning

IBM i Scan Enablement: Special Considerations

- **Save Considerations**
 - Option to scan objects during a save
 - Option to not save objects that have failed a scan
 - Default is SCAN(*NO *NOSAVFAILED)
- **Restore Considerations**
 - NO option to scan objects during a restore
 - NO option to not restore objects that have failed a scan
 - Use the QSCANFCTL system value to control scanning after a restore - *NOPOSTRST
- **File Server Considerations**
 - Many viruses written for Windows PCs
 - Use the QSCANFCTL system value to only scan file server accesses *FSVONLY
- **IASP Considerations**
 - Maintain their own level of scanning
 - No re-scan penalty if IASP scanning level matches system scanning level




Copying

Copying

Copy Commands

Command	Description
CPY	Duplicates an object or set of objects as closely as possible
CPYTOSTMF CPYFRMSTMF	Transfers data between a program-described physical file member, source file, or save file and a stream file
CPYTOIMPF CPYFRMIMPF	Transfers data between an externally-described physical file member and a stream file, program-described physical file member, or source file



Copying

Copy Command Details

Feature	CPY	CPYTOSTMF CPYFRMSTMF	CPYTOIMPF CPYFRMIMPF
Source-> Target file system	Any->Any	QSYS.LIB->Any Any->QSYS.LIB	QSYS.LIB->Any Any->QSYS.LIB
Data conversion			
- CCSID based	Yes	Yes (IBM i 6.1)	Yes
- Codepage based	Yes	Yes	No
- Table based	No	Yes	No
Subtree support	Yes	No	No
Save file support	Yes	Yes	No
Choice of end-of-line char	No	Yes	Yes
Retain original owner	Yes	No	No
Replace target	Yes	Yes	Yes
Append data to target	No	Yes	Yes

- Understand full data conversion concerns. Test various sets of data to avoid problems that may occur with special data sets. Don't just use A-Z and 0-9.

Copying

Copy Functions

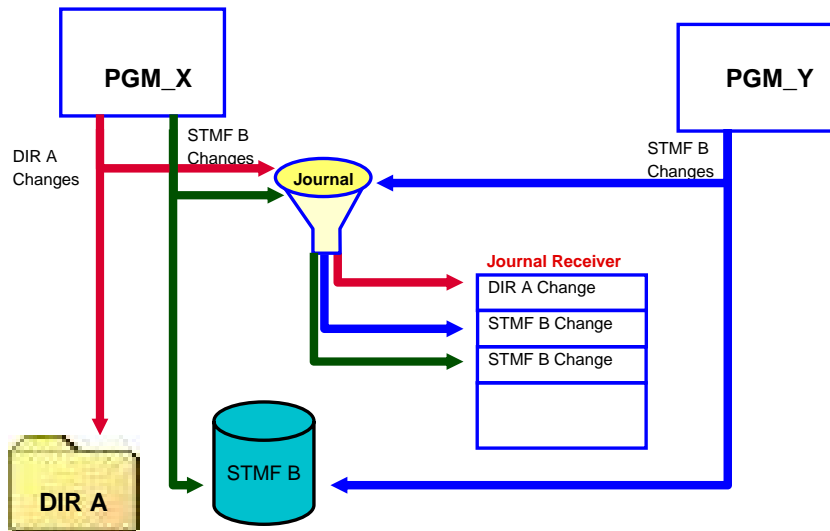
- Be aware that the CPYxxxxx commands may not copy all of the attributes and authorities
 - CPY does a good job for objects in the "root" (/), QOpenSys, and User-defined file systems
 - When copying within a file system, don't forget that there are still specific copy commands available for some file systems such as QDLS (CPYDOC), QSYS.LIB (CRTDUPOBJ), and QOPT (CPYOPT)
 - **IBM i 6.1** - CPY, CPYTOSTMF, CPYTOIMPF have added new parameters to indicate how authority information should be set for the copied objects.
 - See command help and Memo To Users document for IBM i 6.1
- In addition to the IBM i copy functions, copying data can be done using
 - System i Navigator and IBM i NetServer (drag & drop, copy & paste)
 - Depending on the source and destination file systems, these copies may be accomplished by the IBM i file server in response to a single command from the PC client (data movement is all on the IBM i)
 - On the other hand, if the PC client does not recognize that the server can do the copy, the copy may involve data movement to and from the PC
 - FTP

Journaling

Journaling

Example

- Works in the "root" (/), QOpenSys, user-defined file systems
- Works on directories, stream files, symbolic links (plus database files, data areas, and data queues)



Journaling

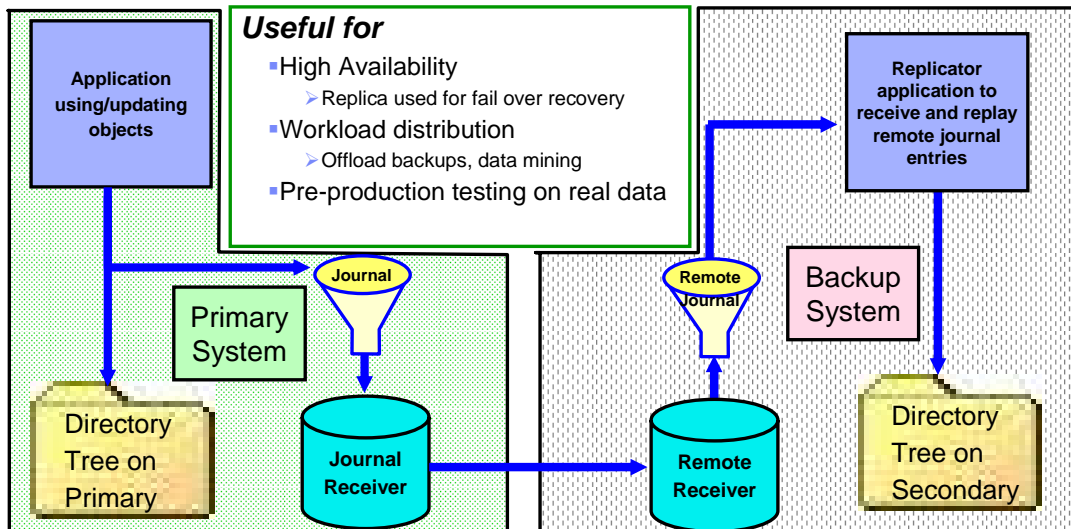
Advantages

- Crash Recovery
 - Automatic recovery after system crash
 - In-flight changes are completed or backed out
 - No need to restore last saved version after the crash
- Single Object Recovery
 - Recover a saved object to a known state
 - Recovers from application or disk failure, or site disaster
- Reduce Backup Time and Volume
 - Save objects weekly
 - Save journal receivers daily
 - Less data to save
- File System Integrity
- More Information
 - IBM i Information Center: *Files and file systems > Integrated file system > Journaling objects*

Journaling

Replication

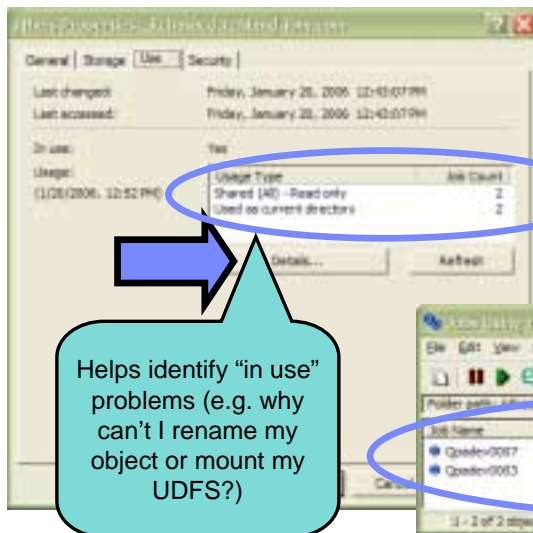
- Transfer journal entries to the replica system's journal receiver
- Replicator application receives journal entries and applies them on the replica
- Replica can be in "standby mode" for quick startup after a fail over



Other Useful Features


Other Useful Features

Object Usage Tracking for an Object in System i Navigator



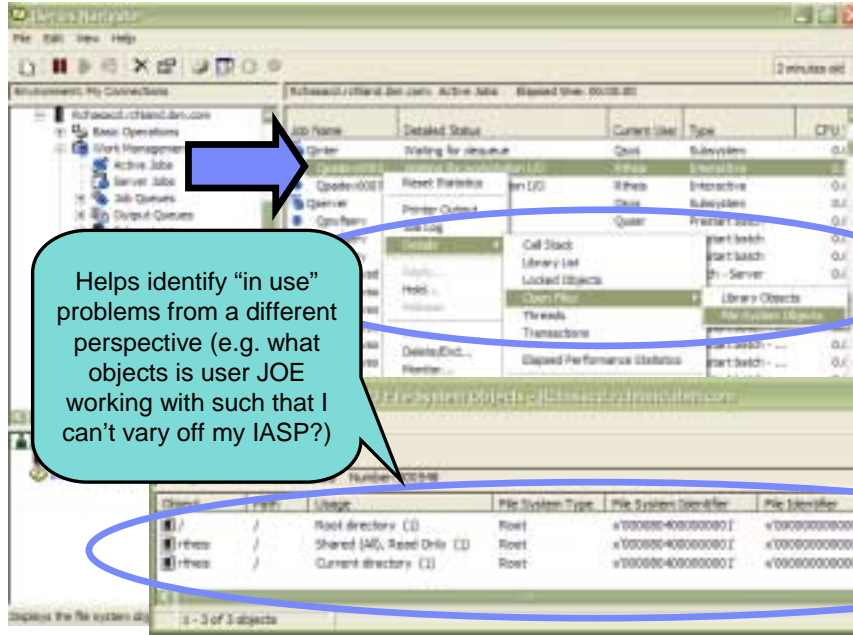
Shows individual object usage

- How many
- Type
- Jobs



Other Useful Features

Object Usage Tracking for a Job in System i Navigator



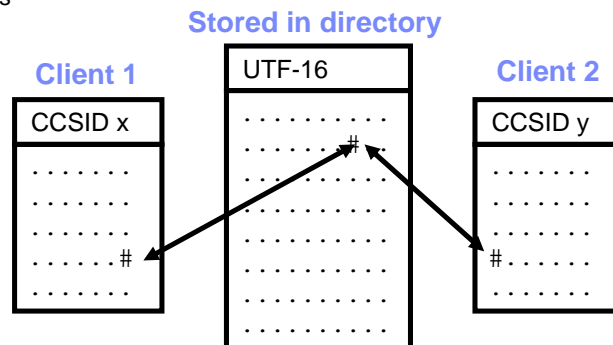
Shows object usage by job

- How many
- Type
- Path name
- File system ID
- File ID

Other Useful Features

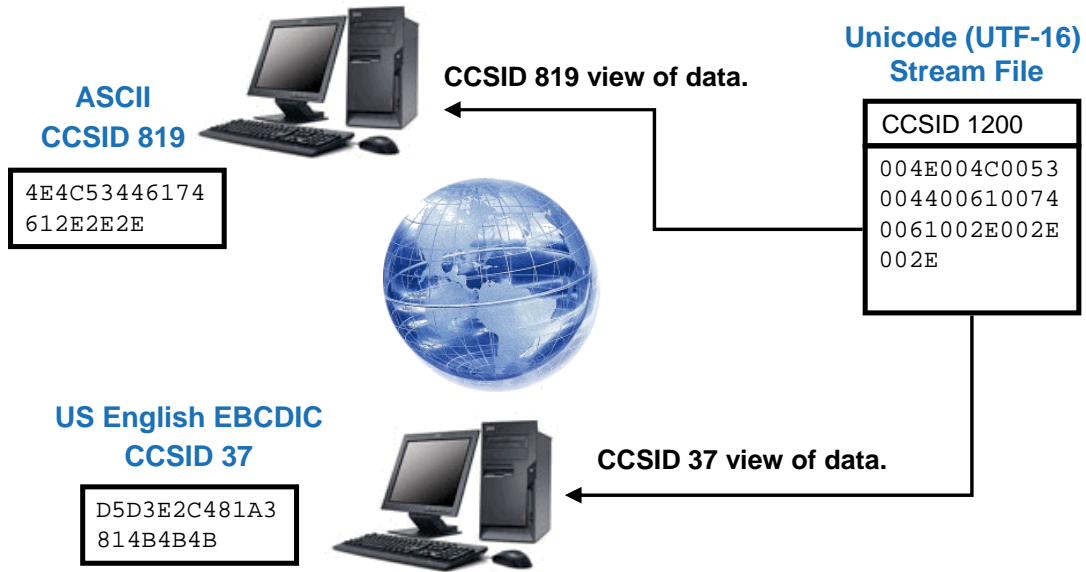
National Language Support

- Multi-language support for IFS APIs
 - Can specify pathnames in any CCSID, language, and country
- Name continuity in the "root" (/), QOpenSys, and user-defined file systems
 - Object names can be specified or retrieved in any CCSID
 - Stored in normalized form (UTF-16)
 - Preserves the name when accessed from jobs, application, and clients using different CCSIDs
 - Helps when developing global applications
- Automatic data conversion using IFS APIs
 - Files are tagged with CCSIDs
 - Conversion occurs to and from the job CCSID or the specified CCSID
- Automatic data conversion using IBM i NetServer File Shares
 - Files are tagged with CCSIDs
 - Conversion occurs to and from the specified CCSID for files with specific extensions



Concepts

Automatic Stream File Data Conversion



Release Enhancements

Release Enhancements

IBM i 7.1

▪ Temporary user-defined file systems

- Available in 6.1 with PTF 5761SS1 SI34983
- User-defined file system
 - Created with:
 - Create User-defined FS (CRTUDFS) command
 - Systems Director Navigator for i (IBM i 7.1 only)
 - Extension must be 'tmpudfs'
- Improved performance with create and delete of stream files*
- Not intended for critical data
- Restrictions
 - Cannot save or restore
 - Cannot journal
 - Only in system ASP (QASP01)
 - No authorization lists
 - No storage accounting

*Measured in lab environments. Actual improvements in customer environments may differ.

Release Enhancements

IBM i 7.1

▪ Solid state drive (SSD) support for user-defined file systems

- Available in 7.1 with PTF 5770SS1 SI39439
- New parameter on the Create User-defined FS (CRTUDFS) command
 - "Preferred storage unit (UNIT)"
 - *ANY – no preferred storage media
 - *SSD – storage for objects in the UDFS should be allocated from SSD
- ASP containing only SSD
 - Storage for objects will be allocated from SSD
 - Attribute will only be set on objects if UNIT(*SSD) indicated when UDFS created

References

- IBM i Information Center
 - Your information source for the IBM i platform
 - <http://publib.boulder.ibm.com/series/>
- IBM i Resource Center
 - The essential destination for the IBM i community. It gives you access to a wealth of information to help you improve your IT department...and your business.
 - <http://www-03.ibm.com/systems/i/resources/>
- IBM i Information Center Integrated file system Topic (*Files and file systems > Integrated file system*)
 - File system information (*Work with file systems*)
 - NLS (*Programming support > Naming and international support*)
 - Integrated file system and System i Navigator (*Working with files and folders using System i Navigator*)
 - And much more...


Backup Slides

***TYPE2 Directories**

***TYPE2 Directories**

Boost Your Integrated File System Performance

- An enhanced implementation of the directory object in the following integrated file systems:
 - "root" (/)
 - QOpenSys
 - UDFSs in basic and independent ASPs
- Built for improved ("Better Plumbing")
 - Performance
 - Reliability
 - Functionality
 - Size and Capacity
- How can you get *TYPE2 directories?
 - New systems come with *TYPE2 directories
 - On existing systems , the conversion is automatic



- More information found in the IBM i Information Center at
*Files and file systems > Integrated file system > Convert directories from *TYPE1 to *TYPE2*

*TYPE2 Directories

Benefits

- **Performance**

- **Reliability**

- Superior recovery
 - from system crashes
- reduces the need for reclaim storage and support calls

- **Functionality**

- System i Navigator automatic name sorting
- Scan enablement

- **Size / Capacity**

- Small directories (< 350 objects) are significantly smaller (empty directory is 4 – 5 x smaller *)
- Large directories (> 350 objects) are 10 percent larger on average
- Increased capacity to store more object attributes
- Supports up to 1 million links (subdirectories in a single directory) (vs. 32765)

- **Compatibility**

- Still a *DIR object - Properties, WRKLNK option 8, and QlgGetAttr API show the directory format
- SAV / RST to and from *TYPE1 works with no conversion needed

*Measured in lab environments. Actual improvements in customer environments may differ.

Operation	Improvement Factor*
Create directory	up to 12 - 13 x
Delete directory	up to 3 - 4 x
Read directory	up to 4 x
Open directory	up to 2 x
Reclaim Storage (RCLSTG)	up to 1.5 - 2 x
Save / Restore	up to 2x (in many-directory workload)

*TYPE2 Directories

Automatic Conversion While Active

- **CVTDIR OPTION(*CHECK)** tells you

- What type of directories you have and which file system is currently being converted
- Progress of the conversion (e.g. Information provided tells you if CWA is done)
- Progress also found in the QFILESYS1 job log and the QSYSOPR message queue

- **CVTDIR OPTION(*CHGPTY) RUNPTY(xx)** lets you tune the priority of the conversion thread

- Default priority is 99 (lowest possible)

- Converts “root” (/), QOpenSys, and UDFSs in basic ASPs 1-32, in that order

- Independent ASPs are still converted, if necessary, when varied on

- After conversion begins on a file system, new directories created in that file system are *TYPE2

- **NOT** “first touch” conversion. Instead, it converts directories that aren’t being used, and interlocks with system activity on directories that are being used. Thus, all directories eventually get converted, not just those that are touched.

- No restrictions on file system or application usage, except

- Reclaim of directories is restricted until fully converted (must can OMIT(*DIR) on RCLSTG)
- Scanning doesn’t occur until the file system is fully converted

Performance

Performance

Performance Improvement References

■ **General User Journaling Performance**

- IBM i Information Center: *Systems Management > Journal management > Local journal management > Journal management concepts > Journal management and system performance*

■ **Integrated File System Save and Restore Performance**

- Chapter 9 in the following manual.
“IBM i 7.1 Performance Capabilities Reference”
<http://www-03.ibm.com/systems/i/solutions/perfmgmt/resource.html>
- Chapter 15 in the following manuals.
“V6R1 Performance Capabilities Reference”
“V5R4 Performance Capabilities Reference”
<http://www-03.ibm.com/systems/i/solutions/perfmgmt/resource.html>
- IBM i Information Center: *Files and file systems > Integrated file system > Related information for integrated file system > Experience reports > Backing up the integrated file system*

Retrieve and Print Directory Information

Retrieve and Print Directory Information

RTVDIRINF and PTRDIRINF Commands

Additional Command Information

- In IBM i 5.4, internal updates have been made to improve performance by as much as 7x*
- Note: In V5R3, the CCSID of the job from which the commands are run cannot be 65535. To work around this, use SBMJOB CMD(RTVDIRINF DIR('<target>')) CCSID(xxx), where xxx is the appropriate CCSID for your country or region. In V5R4, the commands default the CCSID to 37.
- More information in the IBM i Information Center at *Files and file systems > Integrated file system > Access the Integrated File System > Access using CL commands > Work with output of the RTVDIRINF and PRTDIRINF commands.*

Run Queries to Obtain Useful Information

- In System i Navigator, do the following to prepare to run queries
 - 1) Right-click on [the database that you want to work with] and select *Run SQL Scripts*
 - 2) *The Run SQL Scripts* window will be displayed
- Run any of the following example queries by entering them into the window and selecting the *Run* menu followed by *All* →
- Can also be done on the “green screen” using native SQL and Query

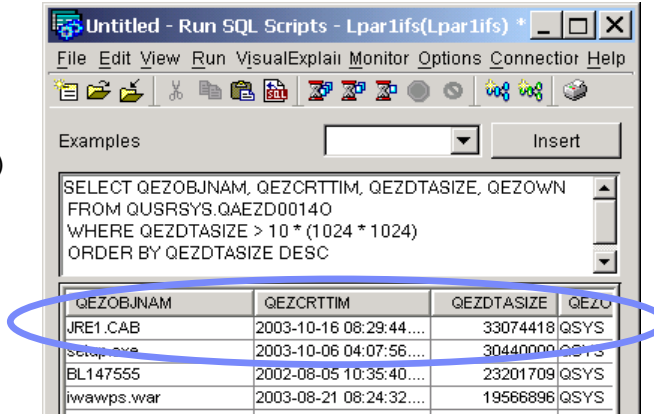
*Measured in lab environments. Actual improvements in customer environments may differ.

Retrieve and Print Directory Information

Find Large Objects

```
SELECT QEZOBJNAM, QEZCRTTIM,
       QEZDTASIZE, QEZOWN
FROM QUSRSYS.QAEZD00140
WHERE QEZDTASIZE > 10 * (1024 * 1024)
ORDER BY QEZDTASIZE DESC
```

QEZOBJNAM is the object name.
 QEZCRTTIM is the object creation time.
 QEZDTASIZE is the data size
 QEZOWN is the owner of the object.



Lists all objects which are **greater than 10 MB in size** in order from largest to smallest.

There is a file JRE1.CAB created on November 16, 2003 which is owned by QSYS and uses approximately 33 MB of storage.

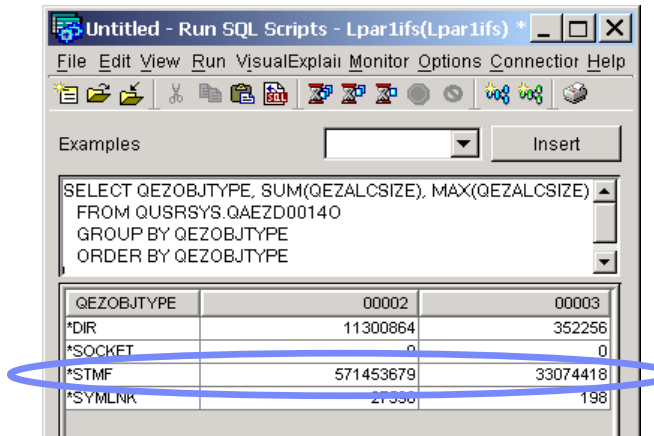
Retrieve and Print Directory Information

Get Storage Usage Summary

```
SELECT QEZOBJTYPE,
       SUM(QEZALCSIZE), MAX(QEZALCSIZE)
FROM QUSRSYS.QAEZD00140
GROUP BY QEZOBJTYPE
ORDER BY QEZOBJTYPE
```

QEZOBJTYPE is the object type.
 QEZALCSIZE is the physical size of the object.

Calculates the **total amount of storage used** for each object type and finds the **largest object** of each type.



Stream files take up 570 MB of storage, including a 33 MB file.

Retrieve and Print Directory Information

Join Tables to Get Detailed Information About Directory Contents

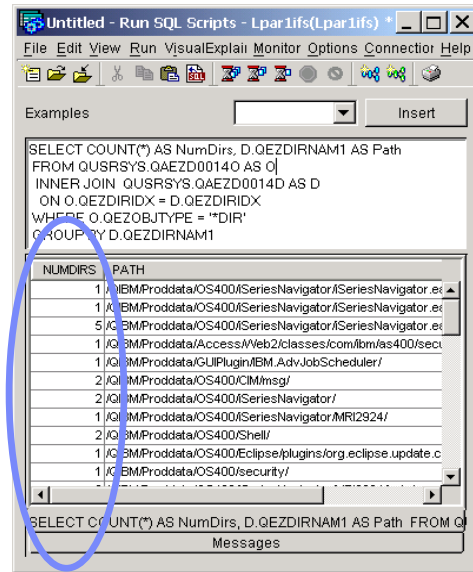
```

SELECT COUNT(*) AS NumDirs,
       D.QEZDIRNAM1 AS Path
FROM QUSRSYS.QAEZD0012O AS O
     INNER JOIN QUSRSYS.QAEZD0012D AS D
       ON O.QEZDIRIDX = D.QEZDIRIDX
WHERE O.QEZOBJTYPE = '*DIR'
GROUP BY D.QEZDIRNAM1

```

QEZDIRNAM1 is the pathname of the directory.

Find the **number of subdirectories** in each directory.



Retrieve and Print Directory Information

Find Path Name Information for Objects of interest

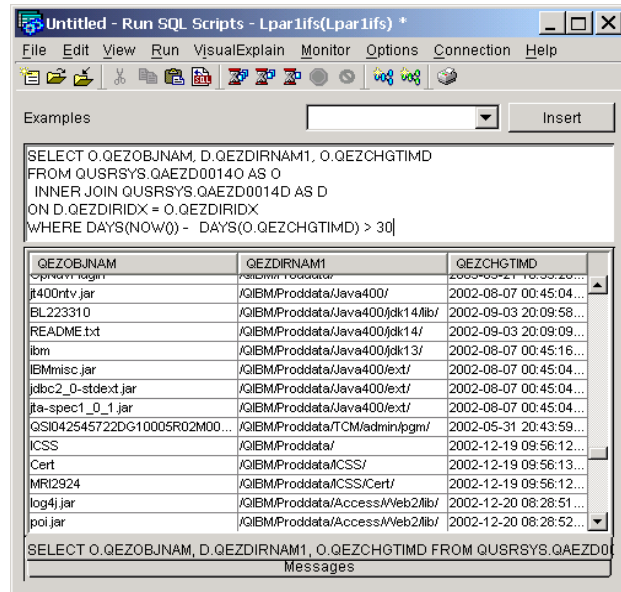
```

SELECT O.QEZOBJNAM, D.QEZDIRNAM1,
       O.QEZCHGTIMD
FROM QUSRSYS.QAEZD0014O AS O
     INNER JOIN QUSRSYS.QAEZD0014D AS D
       ON D.QEZDIRIDX = O.QEZDIRIDX
WHERE DAYS(NOW()) -
       DAYS(O.QEZCHGTIMD) > 30

```

Finds object names and paths for all objects which have **not been modified within the past month**.

NOTE: Information is only current at the time that it is collected by the RTVDIRINF command.



Other Useful Features

Other Useful Features

Powerful IFS APIs

■ Large-File

- Regular APIs support files up to 2 gigabytes
- Large-file APIs expand support to approximately 1 terabyte
- Can map existing APIs to large-file APIs, or can use large-file APIs directly
- Controlled by compile flags
- C-runtime large-file support is controlled by SYSIFCOPT(*IFS64I0)

■ Memory Mapping

- Obtain a process-local pointer to the file...access it directly for reading or writing
- Private or shared mapping
- System value QSHRMEMCTL limits write access

■ Process Subtree

- Sweeps a tree, selects objects, and passes them to a user-exit function for processing
- It handles object selection and sweeping the tree...you just do your thing

Other Useful Features

Powerful IFS APIs

- `pread()`, `pwrite()`
 - Most efficient way to perform random I/O
 - Combines `lseek()` with `read()/write()`
- `accessx()`, `faccessx()`, `QlgAccessx()`
 - Determines whether an object can be accessed by a class of users in a particular way
 - Classes of users include: current thread, others besides the owner, or everyone
- `O_SYNC`, `D_SYNC`, `R_SYNC`
 - `open()`-specified controls over the integrity of your file's data and/or attributes
 - Sync's data and attributes (`O_SYNC`) or just data (`D_SYNC`) on write()
 - On `read()` too, when `R_SYNC` is also specified
- Retrieve Object References
 - Find and show references (how many, type, jobs) to an object
 - Available on System i Navigator Properties/Usage tab

Other Useful Features

Powerful IFS APIs

- Retrieve Referenced Objects
 - Find and show references to integrated file system objects by a specific job
 - Available in System i Navigator via Details > Open Files > File System Objects on a specific job
- `fclear()`, `fclear64()`
 - Most efficient way to clear a file
 - Writes binary zeros
- Change Scan Signature
 - Changes the scan key signature associated with a specific scan key
 - Scan enablement
- Retrieve Scan Signature
 - Retrieve the scan key signature associated with a specific scan key
 - Scan enablement
- More information found in the IBM i Information Center at *Programming > Application programming interfaces > APIs by category > UNIX-Type > Integrated File System APIs*

Release Enhancements

Release Enhancements

IBM i 6.1

- **Support additional Unicode characters**
 - File systems store names in Unicode
 - File systems that are not case sensitive are affected by changes in characters and casing rules for a specific Unicode standard
 - “root” (/), and user-defined file systems (UDFS) created with CASE(*MONO) support Unicode Standard 4.0 as of IBM i 6.1
 - Automatic conversion for “root” (/) and UDFSs in basic user ASPs (1-32)
 - Starts shortly after IBM i 6.1 operating system is installed
 - Runs in a low priority thread of the QFILESYS1 system job
 - Resumes after each IPL until “root” (/) and all UDFSs in basic user ASPs conversion is complete

Release Enhancements

IBM i 6.1

- **Support additional Unicode characters (continued)**
 - Automatic conversion for UDFSs in independent ASPs
 - During vary on of the independent ASP, QFSYS##### system job is started for independent ASP group #####.
 - Automatic conversion runs in a low priority thread of this system job
 - Not started until name conversion for “root” (/) and UDFSs in basic user ASPs is complete
 - Resumes each time the independent ASP is varied on until the UDFSs conversion is complete
 - Only one independent ASP conversion is allowed to be active at any given time. When one completes, another is allowed to become active.
 - Status of name conversion
 - General messages sent to job log of job performing the conversion
 - Error messages sent to QHST

Release Enhancements

IBM i 6.1

- **Support additional Unicode characters (continued)**
 - Analyze Object Conversion (ANZOBJCVN) in previous release
 - Identify names with at least one character affected by new Unicode characters and casing rules
 - Potential for objects to be renamed by the automatic conversion when names have affected characters
 - For additional information
 - Information APAR II14306
 - Redpaper REDP-4293, “i5/OS Program Conversion: Getting Ready for i5/OS V6R1”
at <http://www.redbooks.ibm.com>
 - IBM i Information Center: *Files and file systems > Integrated file system > Integrated file system conversion > Converting names to support additional characters*

Release Enhancements

IBM i 6.1

- **Collecting and analyzing folder attributes with System i Navigator**
 - GUI alternative to Retrieve Directory Information (RTVDIRINF) and Print Directory Information (PRDIRINF) commands. See examples in earlier slides
- **Support IPv6 connections**
 - QNTC, QFileSvr.400, NFS, and IBM i NetServer support both IPv4 and IPv6 addresses
- **QNetWare file system**
 - QNetWare file system is no longer supported
- **Unicode enablement of some CL commands**
 - Some CL commands are now Unicode-enabled
 - IBM i information center: *Files and file systems > Integrated file system > Accessing the integrated file system > Accessing using CL Commands*

Release Enhancements

IBM i 6.1

- **Subtree support added to more commands**
 - CHKIN, CHKOUT, RMVDIR now have the SUBTREE parameter
 - System i Navigator updated to check in or check out all objects in a folder and subfolders
- **Copy related command changes**
 - New parameter for more flexibility of permissions for created objects
 - AUT parameter added to CPY, CPYTOSTMF
 - STMFAUT parameter added to CPYTOIMPF
 - Stream file CCSID parameter added to CPYTOSTMF and CPYFRMSTMF
 - Memo to Users for IBM i 6.1 and command help for additional information

Additional information

■ “Guide to IFS Reclaim”

- <http://www.systeminetwork.com/>, System iNEWS, August 2007

■ “Slicing and Dicing RTVDIRINF Output”

- www.mcpressonline.com, MC Mag Online, March, 2006

■ “Real-Time” Virus Protection

- www.mcpressonline.com, MC Mag Online, August 9, 2004

Special notices

This document was developed for IBM offerings in the United States as of the date of publication. IBM may not make these offerings available in other countries, and the information is subject to change without notice. Consult your local IBM business contact for information on the IBM offerings available in your area.

Information in this document concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

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

The information contained in this document has not been submitted to any formal IBM test and is provided "AS IS" with no warranties or guarantees either expressed or implied.

All examples cited or described in this document are presented as illustrations of the manner in which some IBM products can be used and the results that may be achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions.

IBM Global Financing offerings are provided through IBM Credit Corporation in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates are based on a client's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice.

IBM is not responsible for printing errors in this document that result in pricing or information inaccuracies.

All prices shown are IBM's United States suggested list prices and are subject to change without notice; reseller prices may vary.

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

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. Some measurements quoted in this document may have been estimated through extrapolation. Users of this document should verify the applicable data for their specific environment.

Special notices (cont.)

IBM, the IBM logo, ibm.com AIX, AIX (logo), AIX 6 (logo), AS/400, Active Memory, BladeCenter, Blue Gene, CacheFlow, ClusterProven, DB2, ESCON, i5/OS, i5/OS (logo), IBM Business Partner (logo), IntelliStation, LoadLeveler, Lotus, Lotus Notes, Notes, Operating System/400, OS/400, PartnerLink, PartnerWorld, PowerPC, pSeries, Rational, RISC System/6000, RS/6000, THINK, Tivoli, Tivoli (logo), Tivoli Management Environment, WebSphere, xSeries, z/OS, zSeries, AIX 5L, Chiphopper, Chipkill, Cloudscape, DB2 Universal Database, DS4000, DS6000, DS8000, EnergyScale, Enterprise Workload Manager, General Purpose File System, , GPFS, HACMP, HACMP/6000, HASM, IBM Systems Director Active Energy Manager, iSeries, Micro-Partitioning, POWER, PowerExecutive, PowerVM, PowerVM (logo), PowerHA, Power Architecture, Power Everywhere, Power Family, POWER Hypervisor, Power Systems, Power Systems (logo), Power Systems Software, Power Systems Software (logo), POWER2, POWER3, POWER4, POWER4+, POWER5, POWER5+, POWER6, POWER7, pureScale, System i, System p, System p5, System Storage, System z, Tivoli Enterprise, TME 10, TurboCore, Workload Partitions Manager and X-Architecture are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml

The Power Architecture and Power.org wordmarks and the Power and Power.org logos and related marks are trademarks and service marks licensed by Power.org. UNIX is a registered trademark of The Open Group in the United States, other countries or both. Linux is a registered trademark of Linus Torvalds in the United States, other countries or both. Microsoft, Windows and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries or both. Intel, Itanium, Pentium are registered trademarks and Xeon is a trademark of Intel Corporation or its subsidiaries in the United States, other countries or both. AMD Opteron is a trademark of Advanced Micro Devices, Inc. Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries or both. TPC-C and TPC-H are trademarks of the Transaction Performance Processing Council (TPPC). SPECint, SPECfp, SPECjbb, SPECweb, SPECjAppServer, SPEC OMP, SPECviewperf, SPECcapc, SPECchpc, SPECjvm, SPECmail, SPECimap and SPECsfs are trademarks of the Standard Performance Evaluation Corp (SPEC). NetBench is a registered trademark of Ziff Davis Media in the United States, other countries or both. Altivec is a trademark of Freescale Semiconductor, Inc. Cell Broadband Engine is a trademark of Sony Computer Entertainment Inc. InfiniBand, InfiniBand Trade Association and the InfiniBand design marks are trademarks and/or service marks of the InfiniBand Trade Association. Other company, product and service names may be trademarks or service marks of others.

Revised February 9, 2010

Trademarks and Disclaimers

Adobe, Acrobat, PostScript and all Adobe-based trademarks are either registered trademarks or trademarks of Adobe Systems Incorporated 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.

Linux is a registered trademark of Linus Torvalds 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.

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

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.

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

Cell Broadband Engine and Cell/B.E. are trademarks of Sony Computer Entertainment, Inc., in the United States, other countries, or both and are used under license therefrom.

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

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

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

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

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

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

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

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

Prices are suggested U.S. list prices and are subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.