

Integrated File System: Advanced Topics

Agenda Key: 14FD
Session Number: 401194

Margaret Fenlon
mfenlon@us.ibm.com



Power is performance redefined

Deliver services faster, with higher quality
and superior economics

© 2012 IBM Corporation

Integrated File System Sessions

Understanding the Power of the Integrated File System

- 12FC 440137 Integrated File System: Fundamentals
- 14FD 401194 Integrated File System: Advanced Topics ***YOU ARE HERE***
- 26FE 460142 OPEN LAB: Integrated File System and File Servers
- 36AD 406170 IBM i NetServer: Easy Access to Data
- 45FD 440138 Integrated File System Security


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)

7 **Power is performance redefined**

© 2012 IBM Corporation


ASPs & UDFSs

User ASP




IBM i and System i Information Center:
Systems management >
Disk management >
Disk management concepts >
Disk pools

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



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

8 **Power is performance redefined**

© 2012 IBM Corporation


ASPs & 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 and System i Information Center:
Files and file systems >
Integrated file system >
Work with 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 and System 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

ASP & 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

ASP & UDFSs

User-defined File System (UDFS)

- Temporary User-defined File System (UDFS) **←New in 7.1**
 - Designed for performance for objects that do not need to remain on the system permanently
 - Available in 6.1 with PTF 5761SS1 SI34983
 - See future slide for additional information
- Solid state drive support for UDFS **←New in 7.1**
 - Available in 7.1 with PTF 5770SS1 SI39439
 - New "Preferred storage unit (UNIT)" keyword added to the Create User-Defined FS (CRTUDFS) command.
 - See PTF cover letter for additional details
- 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
 - `/dev/QASP01/tmpFiles.tmpudfs` is a *BLKSF object representing a temporary UDFS on ASP 1

ASP s & UDFSs


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 and System i Information Center at *Files and file systems > Integrated file system > File systems > User-defined file systems (UDFSs)*

ASP s & UDFSs

User-defined File System (UDFS)

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




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




ASPes & UDFSs

Step 1: Create a UDFS




ASPes & UDFSs

Step 1: Create a UDFS




ASPs & UDFSs

Step 2: Mount the UDFS



ASPs & UDFSs

Step 2: Mount the UDFS



ASPs & UDFSs

Step 3: View the UDFS Mount

Now, you can create objects under /images and they will reside on the ASP.

Folder icons highlight mount points in the IFS

Select	UDFS Name	Where Mounted	Type	Description	Ch
☐	alan.tmpudfs		TMPUDFS File		Au
☐	alan.udfs		UDFS File		Au
☐	emig.udfs		UDFS File		De
☐	emig1.udfs		UDFS File		De
☐	emig2.udfs		UDFS File		De
☐	emig3.udfs		UDFS File		De
☐	images.UDFS	/images	UDFS File		Apr

21 **Power is performance redefined**

© 2012 IBM Corporation

Performance


22 **Power is performance redefined**

© 2012 IBM Corporation

Performance

Overview

- **Recent changes for backup performance**
- **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

Recent changes for backup performance

- **ASYNCBRING now supported on Save Object (SAV) command and through BRMS**
 - Asynchronously bring objects into memory during backup to reduce the wait time for I/O
 - Supported for "root" (/), QOpenSys, and user-defined file system objects
- **Many customers have seen significant improvement in backup times for file system objects**
- **Considerations that can affect performance with ASYNCBRING**
 - Selection criteria
 - Available memory
 - Directory tree structure
 - Best performance with a well balanced tree when all objects qualify for the save
- **PTFs**
 - SAV
 - 6.1: SI44739
 - 7.1: SI44688
 - BRMS
 - See the BRMS web page for information:
http://www-03.ibm.com/systems/i/support/brms/bestPractice.html#A_support_ASYNCBRING



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
 - Internal file system caching mechanism 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.

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.

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

- IBM Systems Director Navigator for i *Properties > Storage tab*
- System i Navigator *Properties > Storage tab*
- Qp0lSetAttr() 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 your 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 and System i Information Center at: *Systems management > Performance > Applications for performance management > Performance explorer*
- IBM iDoctor for IBM i website at: https://www-912.ibm.com/i_dir/idoctor.nsf




Performance

IBM iDoctor for IBM i



Performance

IBM iDoctor for IBM i



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 and System 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.

Object name	Parent folder path	Allocated size
webnav.wab	/QIBMProdData/OS400/iSeriesNavigator	227.01MB
QPZA_QSI289365761SS10006R01M00505000030000	/QIBMProdData/OS400/iSeriesNavigator	203.01MB
com.ibm.i5OS.webnav.SystemiNavigator_1.0.0.jar	/QIBMProdData/OS400/iSeriesNavigator/plugins	189.01MB
QPZR_QSI293765761JV10006R01M00511000100000	/QIBMProdData/Java400/jdk6/lib	135.57MB
rt.jar	/QIBMProdData/Java400/jdk6/lib	135.57MB
QPZA_QSI290405761SS10006R01M00505000000000	/QIBMProdData/OS/Java/jdk/lib	113.39MB
rt.jar	/QIBMProdData/OS/Java/jdk/lib	113.39MB
rt.jar	/QIBMProdData/Java400/jdk14/lib	79.23MB
QPZA_QSI290135761JV10006R01M00510600060000	/QIBMProdData/Java400/jdk14/lib	79.14MB
setup.exe	/QIBMProdData/OS400/DirSrv/UserTools/Windows	75.01MB

'Allocated size' sizes greater than '50' megabytes Order by 'Allocated size' descending

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>'))

PTRDIRINF 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)

```

Session A - rchasacd - [27 x 132]
File ..... : QPEZDIR
Control ..... :
Find ..... :
*.....1.....2.....3.....4.....5.....6.....7.....8.....9.....0.....1.....2.....3
*.....1.....2.....3.....4.....5.....6.....7.....8.....9.....0.....1.....2.....3
5722SS1 V5R4M0 060210 Directory Space Report RCHASACD 01/23/06 10:18:58 CST
System Information
Information collected ..... : 01/23/06 10:08:21
File prefix ..... : *LAST
Library ..... : *LAST
File system ..... : "root" (/)
Customize options specified for this report:
Report type ..... : *DIR
Include Directory subtree : *ALL
Description
Total disk space on system 1,000,000 bytes
Total disk used space 120024.20
Total disk unused space 879975.80
Directory .. : /QIBM
Description
Total space used by this directory 4287.77
Total space used by objects in this directory and subdirectories 4287.77
Total space used by storage freed files 0.00
Number of Objects 7
% of Disk 3.74
Size in Megabytes 4287.77
F3=Exit F10=Cancel F19=Left F20=Right F24=More keys
03/022
1902 - Session successfully started
  
```

Directory attribute totals

Retrieve and Print Directory Information

Specific Object Information – RPTTYPE(*OBJ)

```

Session A - rchasacd - [27 x 132]
File ..... : QPEZDIR
Control ..... :
Find ..... :
*.....1.....2.....3.....4.....5.....6.....7.....8.....9.....0.....1.....2.....3
*.....1.....2.....3.....4.....5.....6.....7.....8.....9.....0.....1.....2.....3
5722SS1 V5R4M0 060210 Specific Object Information RCHASACD 01/23/06 10:18:58 CST
Object name .. : rt.jar
Object type: *STMF Object size in Kbytes: 402952.00 Last Change: 01/10/06 07:54:19
File Id: X'00000000000000018A0404600003106F' Owner: QSYS
Directory .. : /QIBM/ProdData/Java400/jdk15/Lib
Object name .. : QP2A_QS1216235722SS10005R04M00505000030000
Object type: *STMF Object size in Kbytes: 402948.00 Last Change: 11/29/05 08:18:53
File Id: X'00000000000000018A04046000028182' Owner: QSYS
Directory .. : /QIBM/ProdData/Java400/jdk15/Lib
Object name .. : rt.jar
Object type: *STMF Object size in Kbytes: 282996.00 Last Change: 01/10/06 07:53:45
File Id: X'00000000000000018A0404600003101D' Owner: QSYS
Directory .. : /QIBM/ProdData/Java400/jdk15/Lib
Object name .. : QP2A_QS1216235722SS10005R04M00505000030000
Object type: *STMF Object size in Kbytes: 282932.00 Last Change: 11/29/05 08:18:31
File Id: X'00000000000000018A040473000028106' Owner: QSYS
Specific Object Information
5722SS1 V5R4M0 060210 RCHASACD 01/23/06 10:36:18 CST
F3=Exit F12=Cancel F19=Left F20=Right F24=More keys
03/022
1902 - Session successfully started
  
```

Object specific attributes

Reclaim

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

```

Session A - rchasacd - [24 x 80]
File Edit View Communication Actions Window Help
Reclaim Storage in Progress
01/23/06 RCHASACD 09:10:35
RCLSTG:
Select/Omit/ASP device or group : *ALL *NONE IFSIASP35
Start date and time . . . . . : 01/23/06 09:09:14
Current step / total . . . . . : 7

Reclaim Storage Step      Percent  Time Elapsed  Time Remaining
Data base/library/directory recovery  100    00:00:19    00:00:00
Reading objects from disk             100    00:00:31    00:00:00
Processing data base relationships    100    00:00:15    00:00:00
File ID table recovery               100    00:00:03    00:00:00
Directory recovery                   100    00:00:03    00:00:00
Object description verification       100    00:00:03    00:00:00
Final cleanup                        0      00:00:00    00:00:00

Total . . . . . : 00:01:14    00:00:00
  
```

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



```

Session A - rchasacd - [24 x 80]
File Edit View Communication Actions Window Help
Command Entry
Request level: 1
Previous commands and messages:
> RCLLNK OBJ('/') SUBTREE(*ALL) DMGOBJOPT(*KEEP *DELETE)
27064 of 27064 object links reclaimed with 0 of 0 problems corrected.

Type command, press Enter.
==> RCLLNK OBJ('/') SUBTREE(*ALL) DMGOBJOPT(*KEEP *DELETE)

F3=Exit F4=Prompt F9=Retrieve F10=Include detailed messages
F11=Display Full F12=Cancel F13=Information Assistant F24=Help keys
13281 object links reclaimed; 0 problems found. Object /QIBM/Prod/Java/Java400/
18/007
I902 - Session successfully started
  
```

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 and System 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 and System 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 independent ASP 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 CPYxxxxxx 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 and System i Information Center: *Files and file systems > Integrated file system > Journaling objects*
 - Attend the journaling presentations

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

Helps identify "in use" problems (e.g. why can't I rename my object or mount my UDFS?)

Jobs Using Folder /rtheis - Rchasadcd.rchland.ibm.com

Job Name	User	Status	Type	Entered System
Qpadev0007	Rtheis	Running	Interactive	1/20/06 12:51:58 PM
Qpadev0003	Rtheis	Running	Interactive	1/20/06 12:42:12 PM

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

Helps identify "in use" problems from a different perspective (e.g. what objects is user JOE working with such that I can't vary off my IASP?)


File System Objects - Rchasadcd.rchland.ibm.com

Object	Path	Usage	File System Type	File System Identifier	File Identifier
/	/	Root directory (1)	Root	x'0000804000000001'	x'000000000000000018A0484560000015'
/rtheis	/	Shared (All), Read Only (1)	Root	x'0000804000000001'	x'000000000000000018A04846800031092'
/rtheis	/	Current directory (1)	Root	x'0000804000000001'	x'000000000000000018A04846800031092'

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 support for user-defined file systems**
 - Available in 7.1 with PTF 5770SS1 SI39439
 - Create User-defined FS (CRTUDFS) command
 - New “Preferred storage unit (UNIT)” keyword
 - Will create all objects of the UDFS on solid state drives if available
 - See PTF cover letter for additional details
- Alternatively
 - Configure a user ASP consisting only of solid state drives
 - Any UDFS objects on that ASP will be created on solid state drives

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

Release Enhancements

IBM i 7.1

- **ASYNCBRING now supported on Save Object (SAV) command and through BRMS**
 - Available in 6.1 and 7.1 with PTFs
 - SAV
 - 6.1: SI44739
 - 7.1: SI44688
 - BRMS
 - See the BRMS web page for information:
http://www-03.ibm.com/systems/i/support/brms/bestPractice.html#A_support_ASYNCBRING
 - Asynchronously bring objects into memory during backup to reduce the wait time for I/O
 - See PTF cover letters for additional details

References

- IBM i and System i Information Center
 - Your information source for the IBM i platform
 - <http://publib.boulder.ibm.com/eserver/ibmi.html>
- 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 and System 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...


Integrated File System Sessions

Understanding the Power of the Integrated File System

- 12FC 440137 Integrated File System: Fundamentals
- 14FD 401194 Integrated File System: Advanced Topics
- 26FE 460142 OPEN LAB: Integrated File System and File Servers
- 36AD 406170 IBM i NetServer: Easy Access to Data
- 45FD 440138 Integrated File System Security

Backup Slides

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, RMVDIR 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...**



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


ASP & UDFSs

Step 1: Create a UDFS




ASP & UDFSs

Step 1: Create a UDFS




ASPs & UDFSs

Step 2: Mount the UDFS



ASPs & UDFSs


Step 3: View the UDFS Mount



Performance

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

Performance Improvement References

■ General User Journaling Performance

- IBM i and System 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 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 and System 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 and System 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
 - Right-click on [the database that you want to work with] and select *Run SQL Scripts*
 - 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.

QEZOBJNAM	QEZCRTTIM	QEZDTASIZE	QEZOWN
JRE1.CAB	2003-10-16 08:29:44....	33074418	QSYS
Setup.exe	2003-10-06 04:07:56....	30440000	QSYS
BL147555	2002-08-05 10:35:40....	23201709	QSYS
lrvawps.war	2003-08-21 08:24:32....	19566896	QSYS

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.



QEZOBJTYPE	SUM(QEZALCSIZE)	MAX(QEZALCSIZE)
*DIR	11300864	352256
*SOCKET	0	0
*STMF	571453679	33074418
*SYMLNK	21330	198

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.QAEZD00120 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.



NUMDIRS	PATH
1	/IBM/Proddata/OS400/SeriesNavigator/SeriesNavigator.exe
1	/IBM/Proddata/OS400/SeriesNavigator/SeriesNavigator.exe
5	/IBM/Proddata/OS400/SeriesNavigator/SeriesNavigator.exe
1	/IBM/Proddata/Access/Web2/classes/com/ibm/as400/sec
1	/IBM/Proddata/GUIPlugin/IBM.Adv.JobScheduler/
2	/IBM/Proddata/OS400/ICM/msg/
2	/IBM/Proddata/OS400/SeriesNavigator/
1	/IBM/Proddata/OS400/SeriesNavigator/MRI2924/
2	/IBM/Proddata/OS400/Shell/
1	/IBM/Proddata/OS400/Eclipse/plugins/org.eclipse.update.c
1	/IBM/Proddata/OS400/security/


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.



QEZOBJNAM	QEZDIRNAM1	QEZCHGTIMD
oprcvrtv.jar	/QIBM/Proddata/Java400/	2002-08-21 10:00:20...
BL223310	/QIBM/Proddata/Java400/jdk14/lib/	2002-09-03 20:09:58...
README.txt	/QIBM/Proddata/Java400/jdk14/	2002-09-03 20:09:09...
libm	/QIBM/Proddata/Java400/jdk13/	2002-08-07 00:45:16...
IBMmisc.jar	/QIBM/Proddata/Java400/ext/	2002-08-07 00:45:04...
jdbc2_0-stdext.jar	/QIBM/Proddata/Java400/ext/	2002-08-07 00:45:04...
fta-spec1_0_1.jar	/QIBM/Proddata/Java400/ext/	2002-08-07 00:45:04...
QSI042545722D-G10005R02M00...	/QIBM/Proddata/TCM/admin/pgm/	2002-05-31 20:43:59...
ICSS	/QIBM/Proddata/	2002-12-19 09:56:12...
Cert	/QIBM/Proddata/ICSS/	2002-12-19 09:56:13...
MRI2924	/QIBM/Proddata/ICSS/Cert/	2002-12-19 09:56:12...
log4j.jar	/QIBM/Proddata/Access/Web2/lib/	2002-12-20 08:28:51...
poi.jar	/QIBM/Proddata/Access/Web2/lib/	2002-12-20 08:28:52...

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(*IFS64IO)
- 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 and System 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 and System 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 (PRTDIRINF) 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 and System 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 inquiries, 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.

Revised September 26, 2006



Special notices (cont.)

IBM, the IBM logo, ibm.com AIX, AIX (logo), AIX 5L, AIX 6 (logo), AS/400, BladeCenter, Blue Gene, 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, Active Memory, Balanced Warehouse, CacheFlow, Cool Blue, IBM Systems Director VMControl, pureScale, TurboCore, Chiphopper, Cloudscape, DB2 Universal Database, DS4000, DS6000, DS8000, EnergyScale, Enterprise Workload Manager, General Parallel 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, POWER6+, POWER7, System i, System p, System p5, System Storage, System z, TME 10, 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 full list of U.S. trademarks owned by IBM may be found at: <http://www.ibm.com/legal/copytrade.shtml>.

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

AltiVec is a trademark of Freescale Semiconductor, Inc.

AMD Opteron is a trademark of Advanced Micro Devices, Inc.

InfiniBand, InfiniBand Trade Association and the InfiniBand design marks are trademarks and/or service marks of the InfiniBand Trade Association.

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.

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

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

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.

NetBench is a registered trademark of Ziff Davis Media in the United States, other countries or both.

SPECint, SPECfp, SPECjbb, SPECweb, SPECjAppServer, SPEC OMP, SPECviewperf, SPECcapc, SPECchpc, SPECjvm, SPECmail, SPECimap and SPECsfs are trademarks of the Standard Performance Evaluation Corp (SPEC).

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.

TPC-C and TPC-H are trademarks of the Transaction Performance Processing Council (TPPC).

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

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

Revised December 2, 2010

Notes on benchmarks and values

The IBM benchmarks results shown herein were derived using particular, well configured, development-level and generally-available computer systems. Buyers should consult other sources of information to evaluate the performance of systems they are considering buying and should consider conducting application oriented testing. For additional information about the benchmarks, values and systems tested, contact your local IBM office or IBM authorized reseller or access the Web site of the benchmark consortium or benchmark vendor.

IBM benchmark results can be found in the IBM Power Systems Performance Report at http://www.ibm.com/systems/p/hardware/system_perf.html.

All performance measurements were made with AIX or AIX 5L operating systems unless otherwise indicated to have used Linux. For new and upgraded systems, the latest versions of AIX were used. All other systems used previous versions of AIX. The SPEC CPU2006, LINPACK, and Technical Computing benchmarks were compiled using IBM's high performance C, C++, and FORTRAN compilers for AIX 5L and Linux. For new and upgraded systems, the latest versions of these compilers were used: XL C for AIX v11.1, XL C/C++ for AIX v11.1, XL FORTRAN for AIX v13.1, XL C/C++ for Linux v11.1, and XL FORTRAN for Linux v13.1.

For a definition/explanation of each benchmark and the full list of detailed results, visit the Web site of the benchmark consortium or benchmark vendor.

TPC	http://www.tpc.org
SPEC	http://www.spec.org
LINPACK	http://www.netlib.org/benchmark/performance.pdf
Pro/E	http://www.proe.com
GPC	http://www.spec.org/gpc
VolanoMark	http://www.volano.com
STREAM	http://www.cs.virginia.edu/stream/
SAP	http://www.sap.com/benchmark/
Oracle, Siebel, PeopleSoft	http://www.oracle.com/apps_benchmark/
Baan	http://www.ssaglobal.com
Fluent	http://www.fluent.com/software/fluent/index.htm
TOP500 Supercomputers	http://www.top500.org/
Ideas International	http://www.ideasinternational.com/benchmark/bench.html
Storage Performance Council	http://www.storageperformance.org/results

Revised December 2, 2010

Notes on HPC benchmarks and values

The IBM benchmarks results shown herein were derived using particular, well configured, development-level and generally-available computer systems. Buyers should consult other sources of information to evaluate the performance of systems they are considering buying and should consider conducting application oriented testing. For additional information about the benchmarks, values and systems tested, contact your local IBM office or IBM authorized reseller or access the Web site of the benchmark consortium or benchmark vendor.

IBM benchmark results can be found in the IBM Power Systems Performance Report at http://www.ibm.com/systems/p/hardware/system_perf.html.

All performance measurements were made with AIX or AIX 5L operating systems unless otherwise indicated to have used Linux. For new and upgraded systems, the latest versions of AIX were used. All other systems used previous versions of AIX. The SPEC CPU2006, LINPACK, and Technical Computing benchmarks were compiled using IBM's high performance C, C++, and FORTRAN compilers for AIX 5L and Linux. For new and upgraded systems, the latest versions of these compilers were used: XL C for AIX v11.1, XL C/C++ for AIX v11.1, XL FORTRAN for AIX v13.1, XL C/C++ for Linux v11.1, and XL FORTRAN for Linux v13.1. Linpack HPC (Highly Parallel Computing) used the current versions of the IBM Engineering and Scientific Subroutine Library (ESSL). For Power7 systems, IBM Engineering and Scientific Subroutine Library (ESSL) for AIX Version 5.1 and IBM Engineering and Scientific Subroutine Library (ESSL) for Linux Version 5.1 were used.

For a definition/explanation of each benchmark and the full list of detailed results, visit the Web site of the benchmark consortium or benchmark vendor.

SPEC	http://www.spec.org
LINPACK	http://www.netlib.org/benchmark/performance.pdf
Pro/E	http://www.proe.com
GPC	http://www.spec.org/gpc
STREAM	http://www.cs.virginia.edu/stream/
Fluent	http://www.fluent.com/software/fluent/index.htm
TOP500 Supercomputers	http://www.top500.org/
AMBER	http://amber.scripps.edu/
FLUENT	http://www.fluent.com/software/fluent/fl5bench/index.htm
GAMMESS	http://www.msg.chem.iastate.edu/gamess
GAUSSIAN	http://www.gaussian.com
ANSYS	http://www.ansys.com/services/hardware-support-db.htm
ABAQUS	http://www.simulia.com/support/v68/v68_performance.php
ECLIPSE	http://www.sis.slb.com/content/software/simulation/index.asp?seq=geoquest&
MM5	http://www.mmm.ucar.edu/mm5/
MSC.NASTRAN	http://www.mscsoftware.com/support/prod%5Fsupport/nastran/performance/v04_sngl.cfm
STAR-CD	www.cd-adapco.com/products/STAR-CD/performance/320/index/html
NAMD	http://www.ks.uiuc.edu/Research/namd
HMMER	http://hmmer.janelia.org/
	http://powerdev.osuosl.org/project/hmmerAltivecGen2mod

Click on the "Benchmarks" icon on the left hand side frame to expand. Click on "Benchmark Results in a Table" icon for benchmark results.

Revised December 2, 2010

Notes on performance estimates

rPerf for AIX

rPerf (Relative Performance) is an estimate of commercial processing performance relative to other IBM UNIX systems. It is derived from an IBM analytical model which uses characteristics from IBM internal workloads, TPC and SPEC benchmarks. The rPerf model is not intended to represent any specific public benchmark results and should not be reasonably used in that way. The model simulates some of the system operations such as CPU, cache and memory. However, the model does not simulate disk or network I/O operations.

- rPerf estimates are calculated based on systems with the latest levels of AIX and other pertinent software at the time of system announcement. Actual performance will vary based on application and configuration specifics. The IBM eServer pSeries 640 is the baseline reference system and has a value of 1.0. Although rPerf may be used to approximate relative IBM UNIX commercial processing performance, actual system performance may vary and is dependent upon many factors including system hardware configuration and software design and configuration. Note that the rPerf methodology used for the POWER6 systems is identical to that used for the POWER5 systems. Variations in incremental system performance may be observed in commercial workloads due to changes in the underlying system architecture.

All performance estimates are provided "AS IS" and no warranties or guarantees are expressed or implied by IBM. Buyers should consult other sources of information, including system benchmarks, and application sizing guides to evaluate the performance of a system they are considering buying. For additional information about rPerf, contact your local IBM office or IBM authorized reseller.

=====

CPW for IBM i

Commercial Processing Workload (CPW) is a relative measure of performance of processors running the IBM i operating system. Performance in customer environments may vary. The value is based on maximum configurations. More performance information is available in the Performance Capabilities Reference at: www.ibm.com/systems/i/solutions/perfmgmt/resource.html

Revised April 2, 2007