

Integrated File System: Fundamentals

Agenda Key: 12FC
Session Number: 440137

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 ***YOU ARE HERE*****
- **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**

Integrated File System: Fundamentals

Contents

■ Concepts

- Stream File
- Directory
- Subtree
- File System
- Integrated File System
- Path Name
- Current Working Directory
- Home Directory
- Hard Link
- Symbolic (Soft) Link
- Dot and Dot-Dot
- NLS

■ File Systems

- Overview
- "root" (/)
- QOpenSys
- User-defined File System (UDFS)
- QSYS.LIB
- QDLS
- QOPT
- QNetWare
- QNTC
- QFileSvr.400
- NFS
- Summary

■ Accessing Data


- Overview
- System i Navigator
- IBM i NetServer
- Virus Scanning
- Traditional IBM i Applications
- APIs
- CL Commands
- Qshell
- FTP
- PASE for i
- IBM Toolbox for Java

Concepts

Concepts

Stream File


- Contains a stream of data
- Can be accessed
 - At the byte level
 - Sequentially or randomly
- Examples:
 - PC file
 - UNIX™ file
 - Java™ program
 - Image
 - Audio
 - Video
 - Text
 - PDF



Concepts

Directory


- Contains a list of objects
- Each list element is a "link" (a.k.a. "hard link")
- Each link contains
 - Object name
 - Some object attributes
 - Pointer to the object
- Examples:
 - "root" (/)
 - /QIBM
 - /home



Concepts

Subtree


- A directory may contain links to other directories
- This results in a tree-like structure called a "subtree"
- The top directory is the "root" of the subtree



Concepts

File System


- A subtree having a common set of characteristics
 - ✓ Case sensitivity
 - ✓ Name length
 - ✓ Subtree depth
 - ✓ Directory implementation
 - ✓ Stream file implementation
- Examples:
 - "root" (/)
 - QOpenSys
 - QSYS.LIB
 - QDLS
 - NTFS



Concepts

Integrated File System


- A collection of file systems
- Connected together to form a single subtree
- Under a common "root"
- The set of objects you can access by name under the common "root" is known as a "namespace"



Concepts

Path Name

- A name used to locate an object in the integrated file system
- Consists of
 - ✓ A sequence of directory names, known as the "path"
 - ✓ Separated by a slash '/' (Note: Integrated file system "LG" path names provide delimiter flexibility)
 - ✓ Followed by the object name




- An **absolute path name** begins with '/'
 - Also known as the "full path name" or "fully qualified path name"
- '/' represents the "root" directory
- You specify the entire path between the "root" directory and the object

Concepts


Current Working Directory (CWD)

- A designated directory below which most of your objects are linked
- Used as the starting point for **relative path names**
- Relative path names do NOT begin with '/'
- Scoped to a job
 - All threads in a job share the same current working directory



Concepts

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



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

Concepts

Home Directory

- A path name that is used to
 - Initialize the current working directory when a job starts
 - Start path name resolution at a specific directory
- Configured in the user profile HOMEDIR parameter
 - Default value is /home/user-profile-name
 - Doesn't have to exist when the user profile is created
 - If it doesn't exist when a job starts with that user profile, the current working directory is set to '/'
- Represented by the tilde (~) character in a path name
 - Must be the first character in the path name
 - ~/ : The tilde character is replaced with the home directory path name
 - ~user-profile-name/ : The ~user-profile-name is replaced with the home directory path name defined for the specified user-profile


User profile JOHN has home directory **/home/John**, and is the initial job user profile

Specified path name	Actual path name
~/SourceCode/Hello.C	/home/John/SourceCode/Hello.C
~John/SourceCode/Hello.C	/home/John/SourceCode/Hello.C
SourceCode/Hello.C	/home/John/SourceCode/Hello.C

Concepts

Hard Link


- A connection between a directory and an object
- Created when the object is first created
- Can add multiple hard links
 - ✓ Provides multiple names for an object
 - ✓ Directory and object must be in the same file system
 - ✓ Not supported by all file systems
 - ✓ Not supported when the object is a directory
- **Link count** attribute tells how many hard links the object has



Concepts

Hard Link

- Hard links control the existence of the object
- Each hard link is equivalent
 - There is no "primary" link
- When last hard link is removed, object is destroyed
 - Unless it is open.
 Then, the object is destroyed after the last close




Concepts

EXAMPLE: Why Use Multiple Hard Links?




Benefits of Using Multiple Hard Links:

- Easier recovery from accidental deletions (similar to the "recycle bin" concept)
- Helps maintain quick and easy object access for different users



Concepts

EXAMPLE: Another Hard Link Benefit




Another Benefit of Hard Links:

- Prevent others from accessing an application's temporary objects

Concepts

Symbolic (Soft) Link


- An object containing a path name
- While resolving a path name, the path name in the symbolic link replaces the symbolic link's name in the path name being resolved
 - If the path name in the symbolic link begins with '/', resolution restarts with '/'
- Can refer to an object that
 - ✓ Is anywhere
 - Same file system
 - Different file system
 - Different system
 - ✓ May not exist when the symbolic link is created
 - ✓ Can be deleted, even though the symbolic link refers to it



Concepts

Symbolic (Soft) Link

- An object containing a path name
- While resolving a path name, the path name in the symbolic link replaces the symbolic link's name in the path name being resolved
 - If the path name in the symbolic link begins with '/', resolution restarts with '/'
- Can refer to an object that
 - ✓ Is anywhere
 - Same file system
 - Different file system
 - Different system
 - ✓ May not exist when the symbolic link is created
 - ✓ Can be deleted, even though the symbolic link refers to it




/A/B/OtherNameForY/z

/y begins with '/', so resolution restarts with '/', resulting in...

/y/z

Concepts

EXAMPLE: Why Use Symbolic (Soft) Links?




Benefits of Using Symbolic (Soft) Links:

- Move objects without impacting pathname resolution and applications
- Link across file systems and systems (Not supported by hard links)

Concepts

Dot and Dot-Dot

- An empty directory has a **link count** of 2
 - One link is the link from its parent directory




- The other link is the **dot** link, the link from itself
- Useful when you just want to operate on the current working directory
 - /a/b/c is the current working directory
 - open('.') will open the 'c' directory
 - Without the **dot** link, operating on the current working directory is more difficult

Concepts

Dot and Dot-Dot

- Every directory has a **dot-dot** link that refers to its parent directory
 - This link adds one to the parent directory's link count




- Used to back up one level from the current working directory
 - /a/b/c is the current working directory
 - open('..') will open the 'b' directory
 - Without the **dot-dot** link, resolving **up** a subtree wouldn't be easy

Concepts


National Language Support for Path Names

- Names in the "root" (/), QOpenSys, and user-defined file systems are normalized and stored in UTF-16 in the directory
- When names are extracted from the directory, they are converted to the CCSID being used by that job, application, or client
- This ensures that characters in the names are correct when accessed from clients using different CCSID
- This helps when developing global applications



Concepts

Automatic Stream File Data Conversion



File Systems

File Systems

General

- A **file system** is a subtree having a common set of characteristics
- The **integrated file system** provides access to 10 unique file systems
 - "root" (/) Directory and stream file access
 - QOpenSys Case-sensitive directory and stream file access
 - User-defined Directory and stream file access on auxiliary storage pools
 - QSYS.LIB Library and traditional IBM i object access
 - QDLS Document and folder access
 - QOPT Optical (CD) access
 - QNetWare Novell Netware Server access
 - Support removed in IBM i 6.1
 - QNTC Windows NT™ Server access
 - QFileSvr.400 IBM i to IBM i access
 - NFS Network File System - UNIX file system access

File Systems

"root" (/)

- Contains the "root" directory
- Contains all other directories and objects created in its subtree
- The anchor for all of the other file systems

Other characteristics

- 255 character object names (between the slashes of a path name)
- **Case-insensitive** names, like on a PC (a and A are the same object)
- Multiple hard links

Supports the following object types

- *DIR Directory
- *STMF Stream file (two types: *TYPE1 and *TYPE2)
- *SYMLNK Symbolic link
- *FIFO First In First Out
- *SOCKET UNIX-domain (local) socket
- *BLKSF Block special file (represents a user-defined file system)
- *CHRSF Character special file (for PASE for i devices) (Also *CHRSFC)

File Systems

QOpenSys

- Located under the **/QOpenSys** directory
- Supports the same object types as the "root" (/) file system


- Has the same characteristics as the "root" (/) file system, except

- Names are **case-sensitive**, like on a UNIX system (a and A are different objects)
- Compatible with UNIX-based open system standards, such as POSIX and XPG

File Systems

User-defined File System (UDFS)

- Physically located on auxiliary storage pools (ASPs)...system, user, or independent ASPs
- Has the same characteristics as the "root" (/) file system except names can be
 - Case-sensitive or
 - Case-insensitive...
 - Your choice
- You create them
 - You make them accessible (mount)
 - FYI: Mounting is **not** automatic. It must be redone after every IPL using the Add Mounted FS (MOUNT) command. Consider adding important mounts to an IPL startup program.
 - FYI: UDFSs provide many benefits (See Advanced presentation →)
- You make them inaccessible (unmount)




File Systems

Temporary

User-defined File System (UDFS)

- Available in 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



File Systems

QSYS.LIB

- Located under the /QSYS.LIB directory
- Provides access to objects in libraries in the system ASP and user ASPs

Names

- Start with /QSYS.LIB
- Have an object-type extension

- Symbolic links can simplify the names

Example:

/TAP01 could be a symbolic link to /QSYS.LIB/TAP01.DEVD

- A QSYS.LIB file system is also available on independent ASPs.

Supports the following object types

•LIB	Library
•FILE	Database file
•MBR	File member
•FILE	Save file
•USRSPC	User Space
•XXXXX	Other types

as a directory
as a directory
as a stream file
as a stream file
as a stream file
as "native" objects"

/ QSYS.LIB / ACCOUNTS.LIB
/ QSYS.LIB / QUSRSYS.LIB / QAEZDISK.FILE
/ QSYS.LIB / QUSRSYS.LIB / QAEZDISK.FILE / QCURRENT.MBR
/ QSYS.LIB / QGPL.LIB / SAV031203.FILE
/ QSYS.LIB / POLICIES.LIB / WORKSPC.USRSPC
/ QSYS.LIB / TAP01.DEVD


File Systems

QDLS

- Located under the /QDLS directory
- Provides access to document library objects (DLOs) and PC files in folders
 - Folders are directories
 - DLOs and PC files are stream files

Names are restricted to 8.3 format:

- / QDLS / QFPNTWE / AS4NW.IPS
- Case-insensitive




File Systems

QOPT

- Located under the **/QOPT** directory
- Provides access to an IBM i CD-ROM, DVD-ROM, DVD-RAM, or optical media library


- The volume is a sub-directory under /QOPT
- Directories and files on the optical media show up in a subtree under the volume directory



File Systems

QNTC


- Located under the **/QNTC** directory
- Provides access to directories and files managed by a CIFS server
 - On a remote PC
 - On Linux™ Samba
 - On IBM i using IBM i NetServer
 - On an IBM BladeCenter® blade and System x™ models attached to an IBM i solution via an iSCSI network



File Systems

QFileSvr.400

- Located under the /QFileSvr.400 directory
- Provides access to directories and files managed by another IBM i operating system




- The remote system name is a sub-directory under /QFileSvr.400
- The "root" (/) directory and subtree of the remote system can be accessed through the /QFileSvr.400/SystemName/path name

File Systems

NFS Client

- The IBM i can be an NFS client
- Provides access to directories and files managed by another operating system running an NFS server such as IBM i, Windows™, UNIX, or Linux




- The "RemoteData" directory can be almost any directory

File Systems

NFS Server

- The IBM i can be an NFS server
- Provides access to IBM i directories and files that have been explicitly "exported" for access by remote NFS clients
- Directories /X and /A/B/C have been made available for access by remote NFS clients ("exported")





Accessing Data


The Integrated File System

Overview

- Common IBM i interface to
 - Stream files
 - Directories
 - Subtrees
 - File systems



- IBM i "plumbing" that integrates file systems under a common interface



Stored in many kinds of file systems

Accessing Data

- PC client
 - System i Navigator
 - IBM i NetServer
 - IBM i Access
- Qshell, PASE for i, Java Toolbox applications
- Applications (ILE C/C++, RPG, COBOL, CL, etc.)
- Remote systems
 - IBM i File Server File System (QFileSvr.400)
 - Network File System (NFS)
 - FTP
 - QNTC



All ways come through the integrated file system

Accessing Data


System i Navigator

- The graphical user interface (GUI) for managing and administering your IBM i systems from your Windows desktop


Typical Tasks:

- Create and remove a directory (folder)
- Copy or move a file, directory, or subtree
- Set up security for a file or directory
- Work with properties of files, directories, and subtrees
- Check out/in a file → *Control updates to a file*
- Determine who is using a file → *Helps solve "in use" problems*
- Determine the open file system objects for a job → *Helps solve "in use" problems*
- Set up file text conversion → *Conversion ensures data can be accessed globally*
- Share a subtree with PC clients → *Helps people collaborate*
- Send a file, directory, or subtree to other systems → *Helps people collaborate*
- Start and end journaling of files or directories → *Journaling assists in object recovery*
- Create, mount, unmount, and delete user-defined file systems (UDFSs)
→ *UDFSs allow access to integrated file system objects in auxiliary storage pools (ASPs)*
- And much more...


System i Navigator - Explore the system




System i Navigator - Explore File Systems




System i Navigator - Explore the Integrated File System




System i Navigator - Explore the "root" (/) file system



System i Navigator - Explore directory /mflenon




System i Navigator - General properties



47 Power is performance redefined

© 2012 IBM Corporation


System i Navigator - Storage properties




48 Power is performance redefined

© 2012 IBM Corporation


System i Navigator - Use properties




System i Navigator - Security properties




System i Navigator – Edit file




System i Navigator - Edit file



System i Navigator – Hard Links



System i Navigator – Symbolic Links



Accessing Data

IBM i NetServer *

- IBM i NetServer allows Windows clients to access integrated file system subtrees that have been **shared**
- You create a **file share** using System i Navigator. You specify the
 - Name of the share
 - Path name of the subtree to be shared
 - Type of access you want to allow (read/write or read-only)
 - Maximum number of users you want to allow
 - Text conversion attributes
- You map a drive letter to the share using
 - System i Navigator
 - Windows Explorer
- You can also access the share directly using Windows Explorer without mapping a drive. Direct access is the most secure method.
- IBM i NetServer comes with IBM i
- No additional software needs to be installed on your PC !




* IBM i Support for Windows Network Neighborhood

55 **Power is performance redefined**

© 2012 IBM Corporation


IBM i NetServer - Create a file share




56 **Power is performance redefined**

© 2012 IBM Corporation


IBM i NetServer - Create a file share




IBM i NetServer - Create a file share




IBM i NetServer - View file shares




IBM i NetServer - Delete the file share




Accessing Data – WRKLNK prompt



Accessing Data - WRKLNK



Accessing Data – Display symbolic link



Accessing Data


Virus Scanning

- 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
- Scanning using scanner running on IBM i – **DO**
 - Creates no additional exposure
 - No network resources used
 - Faster scans



Accessing Data

IBM i Scan Enablement: How it works



Accessing Data

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>



Accessing Data

Traditional IBM i Applications

APIs



- Integrated file system APIs
 - open(), read(), write(), close() ...
- C / C++ runtime functions
 - fopen(), fread(), fwrite(), fclose() ...
- Available to ILE languages
 - C / C++
 - COBOL
 - RPG

CL Commands

- WRKLNK / DSPLNK
- EDTF / DSPF
- CRTDIR / RMVDIR
- ADDLNK / RMVLNK
- CHGATR / CHGAUT / CHGOWN
- CPY / MOV
- Wildcard support

More Information...

- IBM i and System i Information Center:
 - Files and file systems >*
 - Integrated file system >*
 - Access the Integrated File System >*
 - Accessing using CL commands -or-*
 - Accessing using APIs*



Accessing Data

APIs

UNIX-Type

- Data Access
 - creat, open, read, write, close, lseek
 - fsync, ftruncate, fclear
 - mmap, munmap
- Authority
 - access, chown, chmod, umask
- Directory
 - opendir, readdir, closedir
 - chdir, getcwd
 - mkdir, rmdir, link, unlink, rename
- Attributes
 - stat, statvfs
- Descriptor Management
 - dup, dup2
 - givedescriptor, takedescriptor
- Miscellaneous Functions
 - fcntl, ioctl, mkfifo, symlink, readlink
 - pathconf
 - utime

IBM i - Type

- Data Access
 - Qp0ISaveStgFree
- Directory
 - Qp0IGetPathFromFileID
 - Qp0IProcessSubtree
- Attributes
 - Qp0IGetAttr, Qp0ISetAttr
- Descriptor Management
 - DosSetRelMaxFH
- Journaling
 - QjoStartJournal, QjoEndJournal
 - QjoRetrieveJournalInformation
 - QjoRetrieveJournalEntries
 - QJOSJRNE (Send Journal Entry)
- National Language Enabled
 - QlgXXXXX
- Miscellaneous Functions
 - DosSetFileLocks
 - Qp0ICvtPathToSYSObjName
 - QP0LFLOP, QP0LROR, QP0LRRO
 - QP0LCHSG, QP0LRTSG

Accessing Data

CL Commands

- Data Access
 - EDTF, DSPF
 - CHKOUT, CHKIN
 - CPY, CPYFRMSTMF, CPYTOSTMF
 - MOV
- Authority
 - CHGAUD, CHGAUT, WRKAUT, DSPAUT
 - CHGOWN, CHGPGP
 - WRKOBJOWN, WRKOBJPGP
- Directory
 - WRKLNK, DSPLNK
 - ADDLNK, RMVLNK
 - CRTDIR, RMVDIR
 - CHGCURDIR, DSPCURDIR, RTVCURDIR
 - RNM
 - CVTDIR
- Attributes
 - CHGATR
 - PRTDIRINF, RTVDIRINF
- Journaling
 - STRJRN, ENDJRN, CHGJRNOBJ
 - APYJRNCHG
 - SNDJRNE, DSPJRN, WRKJRNA
- File System Functions
 - MOUNT, UNMOUNT
 - ADDMFS, RMVMFS
 - STRNFSSVR, ENDNFSSVR, CHGNFSEXP
 - RPCBIND, ENDRPCBIND
 - CRTUDFS, DLTUDFS
 - DSPUDFS
 - STATFS, DSPMFSINF
- Miscellaneous Functions
 - RLSIFSLCK
 - SAV, RST
 - RCLSTG, RCLLNK

Accessing Data

Qshell

- What is Qshell?
 - Command environment based on POSIX and X/Open standards consisting of 2 parts →
 - Part 1 → **shell interpreter** (or qsh) is a program that reads commands from an input source, interprets each command, and then runs the command using the services of the operating system (**Note:** qsh is patterned after bash but has fewer options so far...)
 - Part 2 → **utilities** (or commands) are external programs that provide additional functions and can be quite simple or very complex
- Why use Qshell?
 - Provides an **extensible command environment** nearly identical to UNIX and Linux
 - **Manage files** in any file system supported by the integrated file system
 - **Run threaded programs** that do thread-safe I/O to and from an interactive session
 - **Write shell scripts** that can be run without modification on other systems using a cross-platform command language
 - **Run powerful utilities** some of which are not provided via the GUI or CL commands
 - **Write your own utilities** to extend the functions provided by Qshell
- **More Information** found in the IBM i and System i Information Center at *Programming > Java > Qshell*

Accessing Data

Qshell

• attr	Get or set attributes for files	• mkfifo	Make FIFO special files
• basename	Return non-directory portion of pathname	• mv	Move files
• cat	Concatenate and print files	• od	Dump files in various formats
• cd	Change working directory	• pax	Portable archive interchange
• chgrp	Change file group ownership	• pr	Print files
• chmod	Change file modes	• pwd	Return working directory name
• chown	Change file ownership	• pwdx	Print working directory expanded
• compress	Compress data	• Rfile	Read or write record files
• cp	Copy files	• rm	Remove directory entries
• dirname	Return directory portion of pathname	• rmdir	Remove directories
• file	Determine file type	• setccsid	Set CCSID attribute for file
• find	Find files	• tail	Display the last part of a file
• getconf	Get configuration values	• tar	File archiver
• head	Copy the first part of files	• touch	Change file access and modification times
• ln	Link files	• umask	Get or set the file mode creation mask
• ls	List directory contents	• Uncompress	Expand compressed data
• mkdir	Make directories	• zcat	Expand and concatenate data

- More information found in the IBM i and System i Information Center at *Programming > Java > Qshell > Utilities*

Accessing Data

FTP

- Use name format 1 (NAMEFMT 1) to use path names to specify the files being transferred
- You must use stream mode (MODE S) and file structure (STRUCT F) when transferring stream files
- Can specify how CCSID conversion is to be handled (TYPE C)
- The integrated file system limits FTP support to transferring file data only. You cannot use FTP to transfer attribute data.
- QSYS.LIB file system considerations
 - FTP support is limited to physical file members, source physical file members, and save files. You cannot use FTP to transfer other object types, such as programs (*PGM).
 - However, you can save other object types to a save file, transfer the save file, and then restore the objects.
 - Recommended that you pre-create the target file when sending to the QSYS.LIB file system

Accessing Data

PASE for i

- Allows you to port AIX™ applications to the IBM i with minimal effort
- Provides an integrated runtime environment that allows you to run UNIX applications without the complexity of managing a UNIX system
- Uses the **integrated file system** for data access
- More information found in the IBM i and System i Information Center at *Programming > IBM PASE for i*

Accessing Data

IBM Toolbox for Java

- Shipped with IBM i Access for Windows
- Provides Java classes for accessing IBM i resources
 - Including the **integrated file system**
- Uses the IBM i Host Servers as access points to the system
- More information found in the IBM i and System i Information Center at *Programming > Java > IBM Toolbox for Java*

References

- IBM i and System i Information Center
 - Your IBM i and System i information source
 - <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 (PRDIRINF) 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
- **Unicode enablement of some CL commands**
 - Some CL commands are now Unicode-enabled
- **And more...**



File Systems


File Systems

QNetWare –

Support removed in i 6.1

- Located under the /QNetWare directory
- Provides access to directories and files managed by a Novell NetWare server on a remote PC


- Available only when NetWare Enhanced Integration, IBM i option 25, is installed on IBM i V5R4



Accessing Data


The Integrated File System

Who / What Accesses Data in the Integrated File System?




*PASE for i is the Portable Applications Solutions Environment


System i Navigator - Copy subtree




System i Navigator - Paste subtree




System i Navigator - Paste subtree




Accessing Data – The “Green Screen” command line




Accessing Data - WRKLNK




Accessing Data - WRKLNK




Accessing Data - WRKLNK




Accessing Data - WRKLNK



Accessing Data – Display entire name



Accessing Data - WRKLNK



Accessing Data – Display file

```

iSeries System
File Edit View Communication Actions Window Help
Browse : /QIBM/UserData/OS400/NetworkAuthentication/krb5.conf
Record : 1 of 11 by 18 Column : 1 85 by 131
Control :

.....1.....2.....3.....4.....5.....6.....7.....8.....9.....0.....1.....2.....3.
*****Beginning of data*****
??(libdefaults??)
default_keytab_name = /QIBM/UserData/OS400/NetworkAuthentication/keytab/krb5.keytab
default_realm = WINREALM.DOMAIN.COM
kdc_use_tcp = 1
??(realms??)
WINREALM.DOMAIN.COM = ??<
kdc = keyd:88
??>
??(domain_realm??)
.rchland.ibm.com = WINREALM.DOMAIN.COM
??(capaths??)
*****End of Data*****

F3=Exit F10=Display Hex F12=Cancel F15=Services F16=Repeat find F19=Left F20=Right
M8 a 03/012
[902 - Session successfully started

```

Accessing Data - WRKLNK

```

iSeries System
File Edit View Communication Actions Window Help
Work with Object Links

Directory . . . . :
Type options, press Enter.
 2=Edit  3=Copy  4=Remove  5=Display  7=Rename  8=Display attributes
11=Change current directory ...

Opt  Object link      Type      Attribute  Text
2    /                DIR


Option 2 to edit a directory

Parameters or command
==>
F3=Exit F4=Prompt F5=Refresh F9=Retrieve F12=Cancel F17=Position to
F22=Display entire field F23=More options


M8 a 10/003
[902 - Session successfully started

```


Accessing Data – Edit directory



Accessing Data – Show size information



Accessing Data – Subtree option



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.

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.mssoftware.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

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