

GPFS for Linux Questions and Answers



Overview

The General Parallel File System (GPFS) for Linux is a high performance shared-disk file system that can provide data access from nodes in a Linux[®] cluster environment. Parallel and serial applications can readily access shared files using standard UNIX[®] file system interfaces, and the same file can be accessed concurrently from multiple nodes. GPFS is designed to provide high availability through logging and replication, and can be configured for failover from both disk and server malfunctions.

GPFS has been widely accepted on the IBM RS/6000[®] SP[™] system to scale file system I/O which can help meet the requirements of a wide range of applications including seismic data processing, digital library file serving, and data mining in business intelligence.

GPFS for Linux is available as two separately orderable licensed programs:

- GPFS for Linux on @server[®] 325 and xSeries[®] product number 5765-G23
- GPFS for Linux on POWER[™] product number 5765-G20

It is suggested you always check GPFS common FAQ question What is the current service information for GPFS?

Updates to this FAQ include:

Date	Question
July 2005	What Linux kernel patches are provided for GPFS?
	How can I ask a more specific question about GPFS?

Questions & Answers

1. What are the latest Linux kernel serviceability suggestions?
2. How can I obtain Linux kernel support?
3. What are the latest distributions and kernel levels that GPFS has been tested with?
4. Are there any special consideration when running with the Red Hat Pro distribution?
5. What Linux kernel patches are provided for GPFS?
6. What are the mandatory GPFS updates?
7. Are there any special consideration regarding the compilation of the GPFS portability layer?
8. What levels of Reliable Scalable Cluster Technology (RSCT) are required by GPFS?
9. What are the latest bios, firmware and driver levels that GPFS has been tested with?
10. What machine models has GPFS been tested with?
11. Is GPFS supported on all IBM ServerProven @server and xSeries servers?
12. Are there any special considerations for running GPFS on the IBM @server BladeCenter or the @server 325?
13. What are the minimum machine model requirements?
14. What disk hardware has GPFS been tested with?
15. What are the current GPFS advisories?
16. What specific Linux performance tuning suggestions are there?
17. What is the largest tested file system size?
18. Why does GPFS not start automatically after rebooting nodes when using Myrinet?
19. Why aren't all the disks showing up on my disk server nodes?
20. Why can only one host successfully attach to the Fibre Channel loop and see my Fibre Channel disks?
21. What are the requirements when using Cluster System Management (CSM) and GPFS together?
22. What versions of CSM and GPFS can coexist in the same cluster?
23. Is IBM Tivoli® System Automation for Linux (SA) compatible with GPFS?
24. What is stored in the **/tmp/mmfs** directory and why does it sometimes disappear?
25. Why does my system load jump significantly during the night?
26. How can I ask a more specific question about GPFS?

Q1: What are the latest Linux kernel serviceability suggestions?

A1: If your Linux distribution does not automatically include the Linux Kernel Crash Dump (LKCD) facility, we suggest that you maintain a debug version of the Linux kernel with this facility. LKCD may be found at lkcd.sourceforge.net

Note: SUSE LINUX has recently included LKCD in SUSE LINUX Enterprise Server (SLES) 8 (IA32 version).

Should your machine crash in the kernel, IBM service may request a memory dump of the machine at the point of failure. The LKCD facility can be used to perform such a dump.

IBM service may also request that you install the latest kernel version tested with GPFS for a particular Linux distribution to ensure that a problem has not already been resolved in a more recent kernel level. We suggest that you examine the Linux kernel patches for GPFS at www-124.ibm.com/linux/patches/?project_id=132 to see if one of the patches addresses your current problem.

Q2: How can I obtain Linux kernel support?

A2: For Linux issues not related to GPFS, there is additional Linux service support that is available for a fee through IBM Global Services (IGS). This service, IBM Operational Support Services - Support Line for Linux Clusters, can assist with complex hardware and software interoperability issues in support of your Linux cluster. Support Line for Linux Clusters provides support for the underlying Linux operating systems (Red Hat and SUSE LINUX distributions) in addition to your entire cluster. For additional information, see www.ibm.com/services/e-business/linux_8.html

Q3: What are the latest distributions and kernel levels that GPFS has been tested with?

A3: While GPFS runs with many different kernel levels, it is highly suggested that customers apply the latest kernel service updates as indicated in Table 1 and Table 3.

Table 1. GPFS 2.2 for Linux on xSeries

GPFS Version	Linux Distribution	Kernel Version
2.2.0-1 – 2.2.0-3	Red Hat EL 3.0	2.4.21-4*
2.2.0-1 – 2.2.0-3	Red Hat Pro 9	2.4.20-24.9
2.2.0-1 – 2.2.0-3	SUSE LINUX ES 8.0	2.4.21 (service pack 3)
2.2.0-4 – 2.2.1-1	Red Hat EL 3.0	2.4.21-15.EL
2.2.0-4 – 2.2.1-1	Red Hat Pro 9	2.4.20-31.9
2.2.0-4 – 2.2.1-1	SUSE LINUX ES 8.0	2.4.21 (service pack 3)

Table 2. GPFS 2.2 for Linux on e325

GPFS Version	Linux Distribution	Kernel Version
2.2.1-1	SUSE LINUX ES 8.0	2.4.218 (service pack 3)

Table 3. GPFS 2.2 for Linux on POWER

GPFS Version	Linux Distribution	Kernel Version
2.2.0-1 – 2.2.0-3	SUSE LINUX ES 8.0	2.4.21 (service pack 3)
2.2.0-4 – 2.2.1-1	SUSE LINUX ES 8.0	2.4.21-218 (service pack 3)

Notes:

1. * **hugemem** kernel that ships with RHEL 3.0 is incompatible with GPFS.

Q4: **Are there any special consideration when running with the Red Hat Pro distribution?**
A4: When running GPFS on the Red Hat Pro 9.0 distribution, you must ensure the glibc level is 2.3.2-27.9 or greater.

Q5: **What Linux kernel patches are provided for GPFS?**

A5: The kernel patches linked to here are expected to correct problems that may be encountered when using GPFS with the Linux operating system. The supplied patches are currently being submitted to the Linux development community but may not be available in particular kernels. It is therefore suggested that they be appropriately applied based on your kernel version and distribution. A listing of the latest patches, along with a more complete description of these patches, can be found at the project General Parallel File System (GPFS) Linux Kernel Patches at http://sourceforge.net/tracker/?atid=719124&group_id=130828&func=browse:

- **mmap- invalidate**
- **nfsd-attr**
- **NFS lockd**

site.mcr consideration: Patches listing a **site.mcr** define have additional steps to perform:

1. Apply the patch to the Linux kernel, recompile, and install this kernel.
2. In **site.mcr** either **#define** the option or uncomment the option if already present. Consult **/usr/lpp/mmfs/src/README** for more information.
3. Recompile and reinstall the GPFS portability layer.

Q6: **What are the mandatory GPFS updates?**

A6: It is always suggested that the latest GPFS service level is installed:

- GPFS v2.2.0-0 should be updated to at least the 2.2.0-1 level
- GPFS v2.2.1-0 should be updated to at least the 2.2.1-1 level

Q7: **Are there any special consideration regarding the compilation of the GPFS portability layer?**

A7: The GPFS portability layer must be compiled for your specific Linux kernel. For directions on compiling the GPFS portability layer module, follow the directions on an installed GPFS system in the file **/usr/lpp/mmfs/src/README**.

When building the portability layer for SUSE LINUX ES, in order to prevent the failure of the compilation of the GPFS portability layer, you must change the kernel header file search path. The default search path **/usr/src/linux/include** should be changed to the **lib/modules** location:

```
KERNEL_HEADER_DIR=/lib/modules/`uname -r`/build/include
```

Periodic updates of the portability layer may be found at techsupport.services.ibm.com/server/cluster/fixes/gpfsfixhome.html.

Q8: **What levels of Reliable Scalable Cluster Technology (RSCT) are required by GPFS?**

A8: All versions of GPFS prior to 2.2 require the RSCT subsystem to be installed on all nodes. Starting with GPFS 2.2, it is possible to designate a subset of nodes to be used in calculation of node-quorum. On those nodes that are not used in quorum calculation, it is not necessary to install RSCT as a GPFS prerequisite. See the *GPFS for Clusters: Concepts, Planning, and Installation Guide* and search on *Designating quorum nodes*.

The minimum levels of RSCT required by GPFS as listed in Table 4 are shipped with the GPFS image.

Table 4. Minimum required levels of RSCT

GPFS version	Minimum RSCT version
2.2.0-1 – 2.2.0.3	2.3.2.1-0
2.2.0-4 – 2.2.1.1	2.3.3.1-0

However, since IBM products other than GPFS rely on common enabling technologies in RSCT and SRC, it is important when using GPFS and either Cluster System Management (CSM) or IBM Tivoli System Automation for Linux (SA) together that the latest RSCT and SRC RPM images be installed. Select and install or upgrade to the latest RSCT and SRC RPM images available. *Do not* uninstall RSCT and SRC then reinstall them to place the latest RPM image on your system. This will cause your CSM configuration data to be lost. Ensure you use the upgrade option. The RPMs must be installed or upgraded in the specified order and **rsct.basic**, **rsct.core**, and **rsct.core.utils** must be installed through a single command:

```
cat /tmp/gpfs.allnodes | xargs -i rsh {} rpm -Uvh /tmp/srcrsct/src-v.r.m.f-0.arch.rpm
cat /tmp/gpfs.allnodes | xargs -i rsh {} rpm -Uvh /tmp/srcrsct/rsct.core.utils-v.r.m.f-0.arch.rpm \
/tmp/srcrsct/rsct.core-v.r.m.f-0.arch.rpm \
/tmp/srcrsct/rsct.basic-v.r.m.f-0.arch.rpm
```

Additionally for GPFS for the Linux on POWER:

```
/tmp/srcrsct/rsct.64bit.SLES-v.r.m.f-0.ppc64.rpm
```

and GPFS for Linux on @server 325 distributions:

```
/tmp/srcrsct/rsct.64bit.SLES-v.r.m.f-0.x86_64.rpm
```

Q9: **What are the latest bios, firmware and driver levels that GPFS has been tested with?**

A9: The latest bios, firmware and driver levels that GPFS has been tested with are listed in Table 5.

Table 5. Latest levels of bios, firmware and driver tested with GPFS

Hardware	Level	Notes
Qlogic qla2200	Firmware 2.02.03	Single and dual adapters per node
	Driver 6.05.00-fo	
Qlogic qla2300	Firmware 3.02.16	Single and dual adapters per node
	Driver 6.06.00-fo	
Myrinet / GM	Driver 1.6.3	
Subsystem Device Driver (SDD)	IBMsdd-1.5.0.3-9	

Q10: **What machine models has GPFS been tested with?**

A10: GPFS has been tested with:

- The IBM @server xSeries (IA-32) machine models
 - 330
 - 335
 - 340
 - 342
 - 345
 - 360
 - 365
 - 440
- The IBM @server BladeCenter
 - HS20
 - HS40
 - JS20
- The IBM @server 325
- IBM @server pSeries machines models that support Linux
- The IBM @server Cluster 1300
- The IBM @server Cluster 1350

For hardware and software certification, please see the IBM ServerProven® site at www.pc.ibm.com/us/compat/eserver.html

Q11: **Is GPFS supported on all IBM ServerProven @server and xSeries servers?**

A11: GPFS is supported on all IBM ServerProven @server and xSeries servers:

1. with support distributions and kernel levels as stated in Q3: What are the latest distributions and kernel levels that GPFS has been tested with?
2. that meet the minimum machine model requirements as stated in Q13: What are the minimum machine model requirements?

Please see the IBM ServerProven site at www.pc.ibm.com/us/compat/eserver.html

Q12: **Are there any special considerations for running GPFS on the IBM @server BladeCenter or the @server 325?**

A12: GPFS support of the IBM @server BladeCenter JS20 and the @server 325 is currently limited to serving file system data through NSD servers. Externally attached NSD servers connected to the disk subsystem are required. Direct-attached disks are not supported. The NSD servers may be Linux on xSeries, Linux on pSeries, AIX on pSeries, or AIX on @server p5 nodes (with APAR IY58249 or later).

GPFS support of the IBM @server BladeCenter HS20 and HS40 includes both direct-attached disks or disks served through an NSD primary server and, if specified, a backup server.

Q13: **What are the minimum machine model requirements?**

A13: The minimum machine model requirements are:

- GPFS for Linux on xSeries: Intel Pentium 3 or newer processor, with 256 MB of memory
- GPFS for Linux on e325: AMD Opteron™ processors, with 256 MB of memory
- GPFS for Linux on POWER: IBM pSeries server, with 256 MB of memory

Note: GPFS is tested, and supported, on hardware as specified in the *Hardware requirements* section of the GPFS Sales Manual. If you are running GPFS on hardware that is not listed in the *Hardware requirements*, should problems arise, recreation of the problem by the customer on a configuration conforming to the *Hardware requirements* may be required by IBM.

Q14: **What disk hardware has GPFS been tested with?**

A14: GPFS has been tested with the IBM TotalStorage® FAStT200, FAStT500, DS4300, DS4400, and DS4500

Additionally beginning with GPFS 2.2.0.4:

- GPFS has been tested with IBM TotalStorage Enterprise Storage Server (ESS), models 2105-F20 and 2105-800, with SDD, and these Linux distributions:
 - Red Hat EL 3.0
 - SUSE LINUX ES 8.0
- GPFS for Linux on xSeries has been tested with the IBM TotalStorage DS4100, and DS4300, DS4300 Turbo, and DS4500 with IBM TotalStorage DS4000 EXP100 Storage Expansion Unit (1742-60U and 1742-60X) with Serial Advanced Technology Attachment (SATA) disk drives utilizing SM8.42 microcode and these Linux distributions:
 - Red Hat EL 3 (kernel 2.4.21-9.0.1)
 - SUSE LINUX ES 8 (kernel 2.4.19-304)

The IBM ServeRAID™ adapters are specifically not supported.

Q15: **What are the current GPFS advisories?**

A15: The current GPFS advisories include:

1. Some customers with large GPFS clusters (128 or more nodes) have experienced occasional command lockups when executing the **mmstartup -C** and the **mmshutdown -C** commands. This problem was identified as being caused by occasional failures of the Linux TCP layer to correctly handle listen queue overflows on passive TCP sockets. This behavior is detailed below along with some ways to help reduce this problem.

As a part of the normal GPFS startup procedure, many (or all) nodes in the cluster may try to communicate with the primary or secondary GPFS cluster data server node and attempt to execute a command on that node via remote shell. This procedure involves establishing a connection to the server TCP socket created by the remote shell daemon on each these nodes. If the listen queue of the server sockets is not large enough, a queue overflow will occur if too many nodes try to initiate remote shell connections simultaneously.

In a large GPFS environment, prolonged listen queue overflows can present a substantial problem. If the listen queue on the remote shell server is kept full for an extended period of time, server-side TCP starts dropping incoming connection requests without notifying clients about its decision to drop them. This leaves some of the connections in a half-opened state, thereby causing some of the remote shell client processes to hang.

We plan to release a Linux kernel patch that will help address this problem in the future. In the meantime, several preventive steps can be taken to reduce the likelihood that prolonged listen queue overflows will occur in your GPFS environment.

If your GPFS cluster is configured to use **ssh**, we suggest that you obtain the latest **ssh** sources and modify them to enlarge the listen queue of the passive socket created by the **sshd** daemon. This can be achieved by modifying the backlog parameter to the listen function called during the initialization of the **sshd** daemon. For example, in **openssh 3.1**, line 924 of **sshd.c** needs to be examined:

```
"if (listen(listen_sock, 5) < 0)"
```

This line should be changed to

```
"if (listen(listen_sock, SOMAXCONN) < 0)"
```

Additionally, please make sure that the value of **SOMAXCONN** (defined in **include/linux/socket.h**) is set to 128 on your system. After making the above changes, you should compile the ssh package and reinstall **sshd** on your primary and secondary GPFS cluster data server nodes. We suggest that you also increase the value of **MaxStartups** in **sshd_config** to at least 1024. If recompilation is not possible in your environment, or if GPFS is configured to use **rsh** (or any other remote execution command that relies on **xinetd**), we

suggest that you instead increase the number of attempts TCP makes when adding connections the server's listen queue. This can be done by modifying the value of the **tcp_synack_retries** kernel tunable on primary and secondary GPFS cluster servers and we suggest increasing this parameter to at least 10.

2. Large installations with hundreds of compute nodes and a few login or NFS exporting nodes will require tuning the GPFS parameters *maxFilesToCache* and *maxStatCache* with the **mmchconfig** command. The general recommendation is for the compute nodes to set *maxFilesToCache* to about 200. The login or NFS nodes should set this parameter much higher with *maxFilesToCache* set to 1000 and *maxStatCache* set to 50000.

This tuning is required for the GPFS token manager (file locking) which can handle approximately 1,000,000 files in memory. The default setting of *maxFilesToCache* is 1000 and the default for *maxStatCache* is $4 * \text{maxFilesToCache}$, so that by default each node holds 5000 tokens, and the token manager must keep track of a total number of tokens which equals $5000 * \text{number of nodes}$ (and this will exceed the memory limit of the token manager on large configurations).

3. The Linux virtual memory manager (VMM) exhibits undesirable behavior for low memory situations on nodes, where the processes with the largest memory usage are killed by the kernel (via OOM killer), yet no mechanism is available for prioritizing important processes that should not be initial candidates for the OOM killer. The GPFS mmfsd daemon uses a large amount of pinned memory in the pagepool for caching data and metadata, and so the mmfsd process is a likely candidate for termination if memory must be freed up.

Q16: **What specific Linux performance tuning suggestions are there?**

A16: Currently, there are no additions to or updates of the performance tuning suggestions in the *GPFS for Clusters: Concepts, Planning, and Installation Guide*.

Q17: **What is the largest tested file system size?**

A17: Based on architectural parameters (48-bits for disk and sector number) and estimates of expected memory usage for various data structures, the maximum estimated file system size possible for GPFS is about 2 PB with a 1024 KB blocksize.

The largest file system sizes that we have currently validated on Linux are listed in Table 6 (note that we have validated up to 100 TB with GPFS on AIX, so these limits are due to testing/resource constraints on Linux; larger sizes on Linux may be supported through a special bid with a joint testing agreement) :

Table 6. Largest file system size tested with GPFS

GPFS Version	Maximum File System Size Tested
2.2.x-x	95 terabytes

Q18: **Why does GPFS not start automatically after rebooting nodes when using Myrinet?**

A18: If you have enabled GPFS to autostart on reboot via options on the **mmconfig** or the **mmchconfig** commands, you should also ensure that the Myrinet driver is loaded and that the adapter is configured to run over IP when the machine boots up.

Q19: **Why aren't all the disks showing up on my disk server nodes?**

A19:

1. This may be because the **MULTI_LUN** support may be disabled in your distribution's kernel. To enable this support, perform the following steps on your disk server nodes.

On each node that has a SCSI or Fiber Channel Host Bus Adapter, modify **/etc/modules.conf** to include the following line:

```
options scsi_mod max_scsi_luns=255
```


Rename your existing initial ram disk (**initrd**) images which are located in **/boot**. The **initrd** images will have the extension **.img** as in **initrd-2.4.2-2.img**.

Make fresh copies of the **initrd** images by running the **mkinitrd** command. For example:

```
mkinitrd initrd-2.4.20-20.9.img 2.4.20-20.9
mkinitrd initrd-2.4.29-20.9smp.img 2.4.20-20.9smp
```

If using LILO bootloader, issue **/sbin/lilo** to refresh the lilo image; otherwise consult your bootloader documentation for the action to take to ensure that the updated **initrd** image will be used upon next reboot.

2. Some fibre channel host bus adapters do not reliably perform device discovery. In those cases, you must either reboot the node or unload and reload the fibre channel device module.

Q20: Why can only one host successfully attach to the Fibre Channel loop and see my Fibre Channel disks?

A20: Your host bus adapter may be configured with an enabled **Hard Loop ID** that conflicts with other host bus adapters on the same Fibre Channel loop. To see if this is the case, reboot your machine and enter the adapter BIOS via pressing Alt+Q when the Fibre Channel adapter BIOS prompt appears. Under the menu **Configuration Settings** select **Host Adapter Settings** and ensure that either the **Adapter Hard Loop ID** is disabled or assign a unique **Hard Loop ID** per machine on the Fibre Channel loop.

Q21: What are the requirements when using Cluster System Management (CSM) and GPFS together?

A21: It is important when using GPFS and CSM together that the latest RSCT and SRC RPM images are installed (with all RSCT packages at the same level). Select and install or upgrade to the latest RSCT and SRC RPM images from either the GPFS or CSM distribution. Please see the *GPFS for Clusters: Concepts, Planning, and Installation Guide* and search on *Coexistence with CSM* for full information on this procedure which lists all of the RSCT and SRC rpms that should be updated; a standard install of CSM when it has the latest level of RSCT and SRC may not install all of the rpms used by GPFS (such as **rsct.basic**), so the user should verify, and if need be, explicitly install one or more of the RSCT rpms used by GPFS from the CSM distribution. The following is an example of the case where **rsct.basic** must be explicitly installed:

1. Copy **rsct.basic** from the CSM CD-ROM to **/csminstall/Linux/RedHat/csm/1.4.0/packages/**
2. To update **rsct.basic** on the management server, issue **installms -x**
3. To update **rsct.basic** on the nodes, issue **updatenode -a**

Please see the related question, What levels of Reliable Scalable Cluster Technology (RSCT) are required by GPFS?

Q22: What versions of CSM and GPFS can coexist in the same cluster?

A22: The versions of CSM and GPFS that can coexist in the same cluster are listed in Table 7.

Table 7. Versions of CSM and GPFS which can coexist in the same cluster

	GPFS 2.2.x-x
CSM 1.1	Unsupported
CSM 1.2	Yes
CSM 1.3.x-x	Yes
CSM 1.4.x-x	Yes

Q23: Is IBM Tivoli System Automation for Linux (SA) compatible with GPFS?

A23: This depends on GPFS version. For GPFS versions prior to 2.2, GPFS and SA cannot be run on the same nodes. For version 2.2 and later, when GPFS cluster is created on top on an existing RPD domain, GPFS and SA can coexist.

Q24: **What is stored in the /tmp/mmfs directory and why does it sometimes disappear?**
A24: When GPFS encounters an internal problem, certain state information is saved in the GPFS dump directory for later analysis by IBM service. The default dump directory for GPFS is **/tmp/mmfs**.

This directory may disappear when cron runs the **/etc/cron.daily/tmpwatch** script. The **tmpwatch** script removes files and directory in **/tmp** that have not been accessed recently. Administrators wishing to use a different directory for GPFS dumps can change the directory by:

```
mmchconfig dataStructureDump=/someOtherBigFileSystem
```

Q25: **Why does my system load jump significantly during the night?**

A25: Every night cron runs the **/etc/cron.daily/slocate.cron** job, which will try to index all the files in GPFS. This will put a very large load on the GPFS token manager. You can exclude all GPFS file systems by adding **gpfs** to the **excludeFileSystemType** list in this script, or exclude specific GPFS file systems in the **excludeFileSystemType** list.

```
/usr/bin/updatedb -f "excludeFileSystemType" -e  
"excludeFileSystem"
```

If indexing GPFS file systems is desired, only one node should run the **updatedb** command and build the database in a GPFS file system. If the database is built within a GPFS file system it will be visible on all nodes after one node finishes building it.

Q26: **How can I ask a more specific question about GPFS?**

A26: Depending upon the nature of your question, you may ask it in one of several ways:

- If you have a question that can benefit other GPFS users, please post it to the GPFS technical discussion forum at www.ibm.com/developerworks/forums/dw_forum.jsp?forum=479&cat=13. The technical discussion forum is a temporary replacement for the GPFS public community mailing list. The GPFS team is looking to reinstate the mailing list at a new location. Please continue to check this FAQ and the forum for updates regarding the mailing list.
- If you have a question concerning a potential software error in GPFS and you have an IBM software maintenance contract, please contact 1-800-IBM-SERV in the United States or your local IBM Service Center in other countries. IBM Scholars Program users should notify the GPFS development team of potential software bugs through gpfs@us.ibm.com.
- If your question does not fall into the above categories, you can send a note directly to the GPFS development team at gpfs@us.ibm.com. However, this mailing list is informally monitored as time permits and should not be used for priority messages to the GPFS team.

Note: This FAQ is continually being enhanced. Please contribute questions or answers by sending them to gpfs@us.ibm.com

If you have a question that can benefit other GPFS users, then please post it to the public GPFS community mailing list (**gpfs-general@www-124.ibm.com**). You can sign up for this list at oss.software.ibm.com/developerworks/mail/?group_id=91

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's intellectual property rights may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10594-1785
USA

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not

allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Intellectual Property Law
2455 South Road, P386
Poughkeepsie, NY 12601-5400
USA

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment or a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some

measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

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

- AFS
- AIX
- AIX 5L

- @server
- Enterprise Storage Server
- HACMP
- IBM
- IBMLink
- LoadLeveler
- Netfinity
- pSeries
- POWER
- PowerPC
- RS/6000
- Redbooks
- SANergy
- Scalable POWERParallel Systems
- ServerProven
- SP
- Tivoli
- TotalStorage
- xSeries

Intel, Intel Inside (logos), MMX and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

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

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

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

Red Hat, the Red Hat "Shadow Man" logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

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

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

October 2004

© Copyright International Business Machines Corporation 2003, 2005. All rights reserved.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.