Make Your System Highly Available and Recoverable Using PureApplication System

IBM Redbooks Solution Guide

The goal of any enterprise is to operate systems uninterrupted for a long time and to reconstruct a data center's operations in the event of disaster.

This IBM® Redbooks® publication Solution Guide describes how to use the strategies and the capabilities of IBM PureApplication® System to provide system high availability (HA) and disaster recovery (DR). PureApplication System provides a high availability framework to eliminate single points of failure. PureApplication System also provides technologies to aid in designing disaster recovery scenarios. Figure 1 shows a simple disaster recovery strategy.

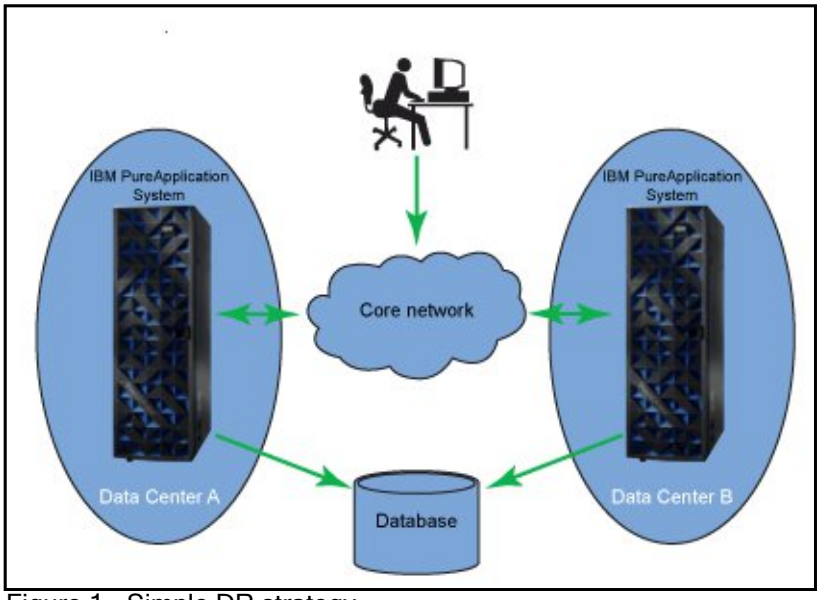

Figure 1 - Simple DR strategy

Did you know?

System administration teams need solid tools and processes to provide system HA and system DR. IBM PureApplication System provides the technology to designhigh availability solutions and disaster recovery scenarios.

Business value

PureApplication System provides a flexible platform for running a wide range of application workloads in a cloud infrastructure. Its design helps to eliminate single points of failure with the goal of allowing applications that are running in the rack to achieve high levels of availability. Businesses that strive for high levels of resiliency must consider how to run workloads across multiple systems and even across geographically distributed data centers. This strategy alleviates problems by resolving them proactively outside of any single piece of hardware, a local network, or power supply. Therefore, a failure of network or equipment does not result in a potentially lengthy outage for that service.

Solution overview

The HADR solutions provided by PureApplication System are described as HADR scenarios in *Data Recovery Strategies and Implementation for PureApplication System*, SG24-8246-00. The HADR scenarios describe how to build HADR solutions using a combination of patterns and HADR capabilities (both hardware and software) built into the PureApplication System.

Solution architecture: PureApplication System hardware HA

PureApplication System's hardware redundancy is designed for HA. Compute nodes, network controllers, management nodes, virtualization nodes, storage controllers, and storage are all redundant and contribute to a highly available environment (Figure 2).

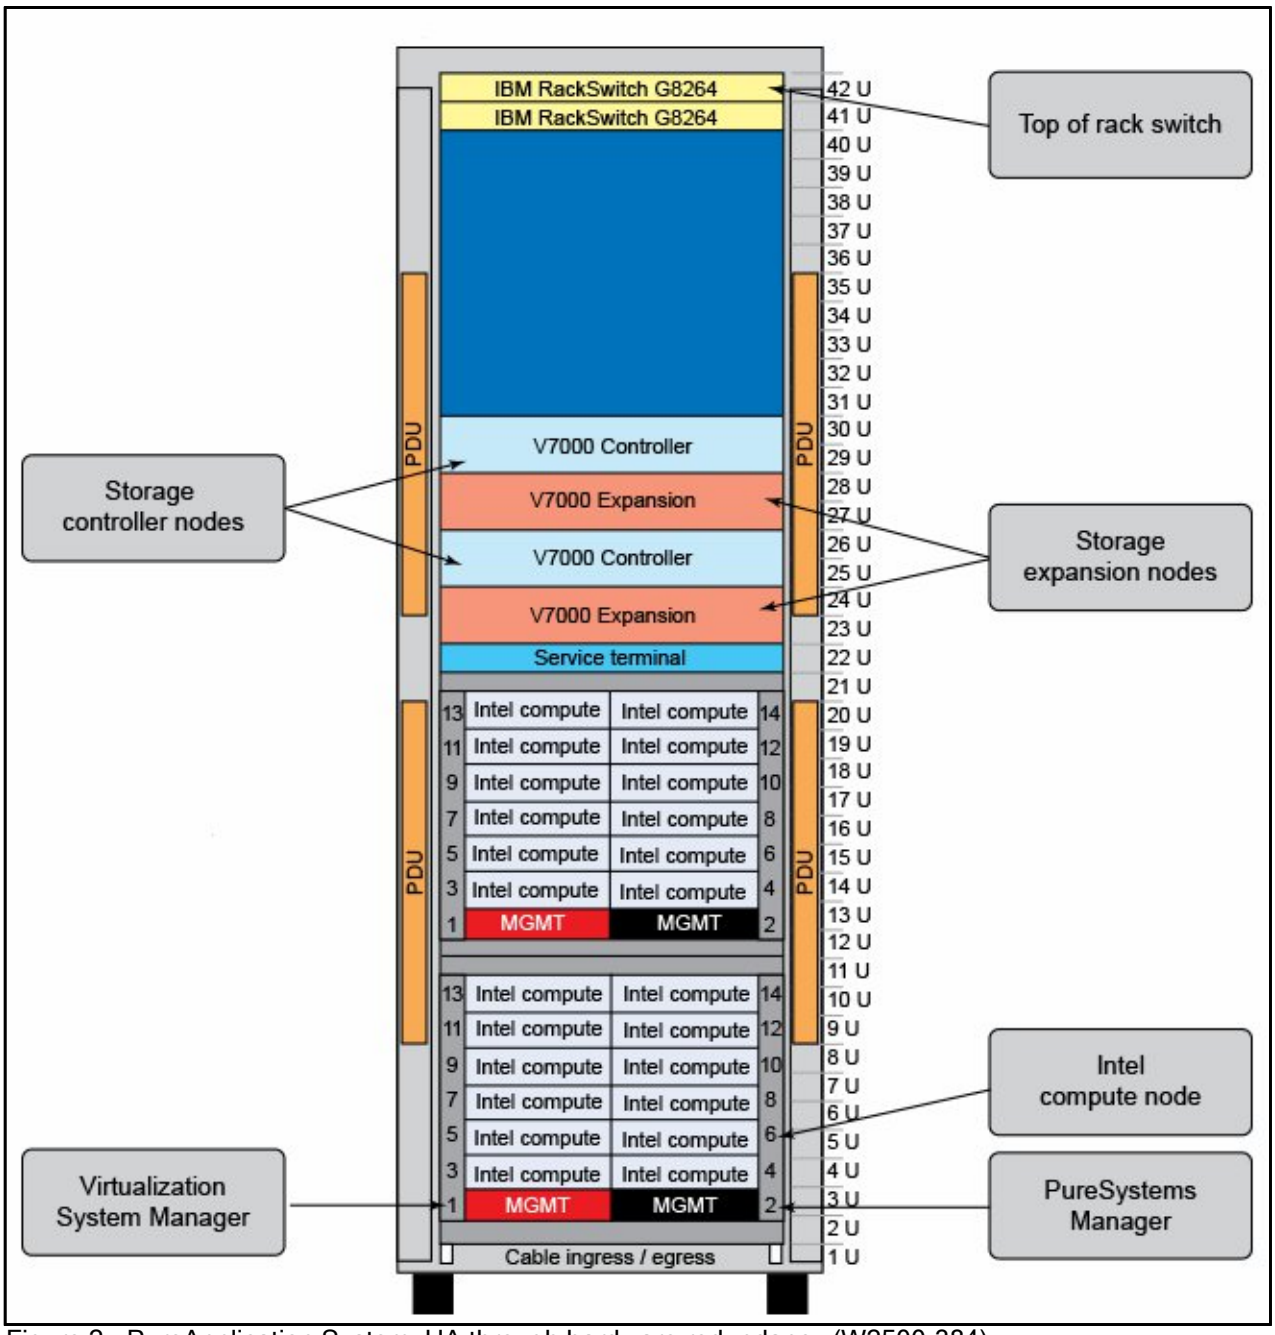

Figure 2 - PureApplication System, HA through hardware redundancy (W2500-384)

The following list notes the component level details available on a single rack system:

- Compute nodes: The management system automatically routes around failed cores. If an entire node fails, the system tries to move the VM to another compute node in its cloud group (if free space is available).
.
- Network controllers: Cabling and switches are redundant. The failure of one of the controllers will reduce bandwidth. Service is continuous.
- Management nodes: Node has a backup server. A floating IP address is assigned to the active management node (workload deployer).
- Virtualization nodes: Node has a backup server.
- Storage controllers: Each controller has two canisters that service all of the traffic. If one fails, the other handles all traffic.
- Storage: SSD and HDD storage is configured in RAID5 plus spares. Storage is designed to tolerate two concurrent failures without data loss (after the spares are in use).

Solution architecture: PureApplication System capabilities

PureApplication System provides storage capabilities and patterns that maximize storage capabilities to aid in HADR solutions. The following list describes those storage capabilities:

- *Storage volumes* are additional storage that can be attached to virtual machines. For example, adding storage volumes for DB2® to host large databases.
.
- *Block storage* is a new type of storage volume that was introduced with PureApplication System V2.0. Block storage is declared with the storage area network (SAN) found in the PureApplication System. Data is stored in volumes, which are also referred to as blocks. Block storage allows administrators to create volumes, of type block, that allow better access to the underlying SAN. Block storage provides access to multiple blocks through the file system.
- Configure *block storage replication* for disaster recovery at the individual storage volume level instead of replicating the entire system configuration for disaster recovery. Replication for each storage volume can be configured between local and remote (recovery) systems in either direction. Replication is configured in volume replication pairs.
- Configure *external storage* to one or more PureApplication System racks. To do so, make a connection from an external storage volume controller to the PureApplication System rack's SAN connection. External storage volumes can then be assigned to cloud groups as either block or block shared storage volume types.

Solution architecture: PureApplication System patterns

Several IBM software products, provided as patterns with PureApplication System, have high availability characteristics built into them:

- *DB2 10.5 Virtual System Pattern*: This pattern includes the HADR feature that provides a high availability solution for both partial and complete site failures. This pattern supports up to three remote standby servers.

- *WebSphere® Application Server patterns*: These patterns compose cell configurations with deployment managers, custom nodes, web servers, and built-in script packages.

- *WebSphere MQ*: There are several HADR options possible with WebSphere MQ (such as MQ Clustering and Multi-Queue Instance Manager (MQM) using GPFS™ as the mechanism to share the file system across the different WebSphere MQ instances).

- Load balancers such as the *WebSphere Application Server On Demand Router and DataPower® XI52 Virtual Appliance*.

- *IBM Pattern for GPFS*: The IBM Pattern for GPFS can be deployed as a GPFS cluster. The GPFS cluster provides high-performance, highly available storage to the middleware and applications that are deployed in the IBM PureApplication System environment. GPFS allows the same file to be accessed from multiple different clients. GPFS is built to be redundant so that the file system remains active even if the host nodes fail.

  Note: More information about patterns can be found in the PureSystems® Centre at: http://www.ibm.com/ibm/puresystems/us/en/puresystemscentre.html

Solution architecture: PureApplication System DR

Three aspects of a PureApplication System must be moved from one system to another to restore functionality of a system's previous configuration. These aspects are noted in the following list and shown in Figure 3:

- Replicate management data: These data include workload components such as patterns, virtual images, pattern types, script packages, and plug-ins. These data can be synchronized automatically through the multi-target deployment features of PureApplication System V2.0.

- Replicate application data: These data include logs, message queues, configuration data, and databases.

- Redirect network traffic: Traffic from the primary system to the backup system.

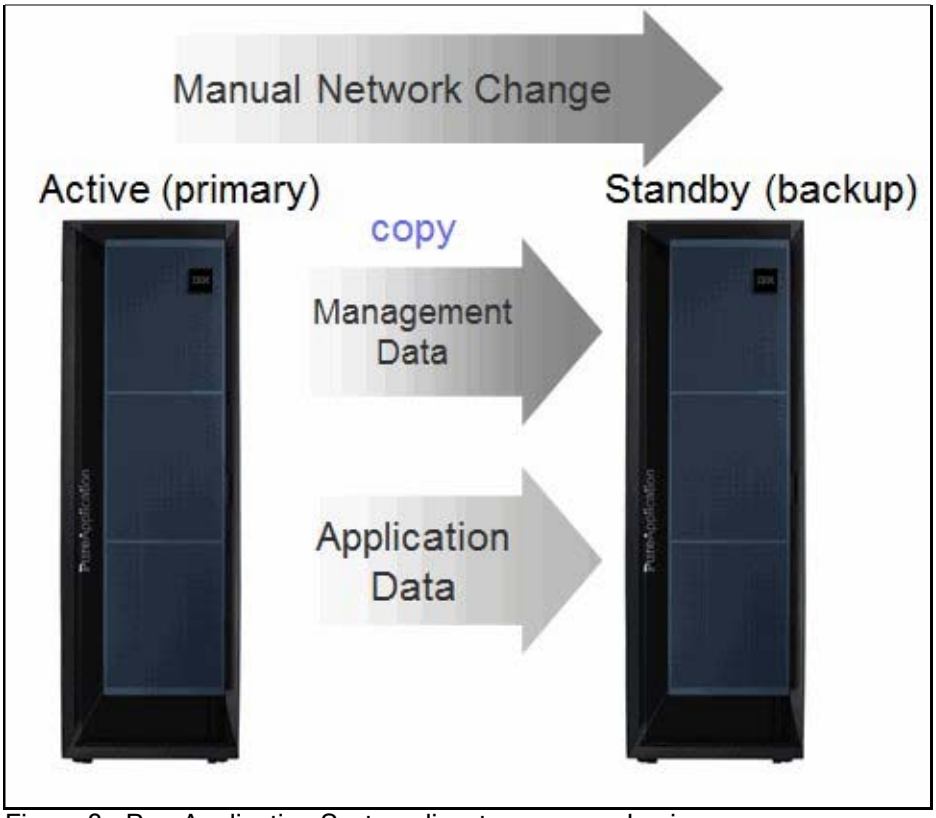

Figure 3 - PureApplication System disaster recovery basics

Supported HADR scenarios

The scenarios listed in the tables in this section are discussed in detail in *Data Recovery strategies and implementation for PureApplication System*, SG24-8246-00. Although there are many different HADR scenarios possible, the focus of that book is on HADR scenarios across IBM middleware products. Several IBM middleware products, available as patterns with PureApplication System V2, include those in the following list and Table 1 through Table 4:

- WebSphere Application Server V8.5.5, see Table 1.

Table 1. WebSphere Application Server scenarios

Scenario	**Description**
Cell in a single rack in the primary data center	Cell in the same rack with transaction log stored in GPFS or database and shared by multiple nodes
Cell across two racks in the primary data center	Transaction logs stored in GPFS or database. Multirack deployment of single pattern across two racks
Two identical cells across the primary data center and secondary data center	Active/Passive. Transaction logs stored in GPFS or database

- DB2 V10.5, see Table 2.

Table 2. DB2 HADR scenarios

Scenario	Description
DB2 HADR Active/Standby in single rack in the primary data center	Single DB2 HADR pattern deployed on a single rack
DB2 HADR Active/Standby across two racks in the primary data center	Single DB2 HADR pattern deployed using multi-rack deployment
DB2 HADR Active/Standby across two racks in the primary data center	DB2 HADR Active and Standby in two patterns
DB2 HADR Active/Standby in primary data center and same Active/Passive DR set-up in secondary data center	Uses block storage replication

- WebSphere MQ V7.5 (Hypervisor Image), see Table 3.

Table 3. WebSphere MQ scenarios

Scenario	**Description**
Two WebSphere MQ instances (primary and standby) with MIQM - in same rack	Same pattern, Active/Standby in primary data center
Two WebSphere MQ instances (primary and standby) with MIQM - across two racks	Active/Standby in primary data center. Split pattern for MQ.

- WebSphere Application Server and DB2, see Table 4.

Table 4. WebSphere MQ scenarios

Scenario	**Description**
WebSphere Application Server cell and DB2 HADR in a single rack in the primary data center.	Transactions stored in GPFS
WebSphere Application Server cell and DB2 HADR in single rack in the primary data center.	Transactions stored in database
WebSphere Application Server Active/Active across two racks, split cell and DB2 HADR in two separate racks in the primary data center.	Single pattern using multi-rack deployment. Transactions stored in database
WebSphere Application Server Active/Active across two racks, split cell and DB2 HADR in two separate racks in the primary data center.	Split pattern across two racks. Transactions stored in database
WebSphere Application Server Active/Passive, stray node in the primary data center .	Transactions stored in database
WebSphere Application Server Active/Passive, two identical cells across primary data center. DB2 environment replicated on disaster recovery site.	Transactions stored in GPFS
WebSphere Application Server Active/Passive, two identical cells across primary data center. DB2 environment replicated on disaster recovery site.	Transactions stored in database

Integration

IBM PureApplication System is an integrated highly scalable system that is based on IBM hardware and middleware products. The integration of hardware and middleware provides a robust environment for configuring HA solutions and DR scenarios. Refer to the following sections in this solution guide:

- “Solution architecture: PureApplication System hardware HA” on page 5
- “Solution architecture: PureApplication System capabilities” on page 6
- “Solution architecture: PureApplication System patterns” on page 7
- “Solution architecture: PureApplication System DR” on page 7

Supported platforms

Each of the systems (whether based on Intel or Power and shown in Table 5 through Table 7) completely supports the full range of high availability and disaster recovery technologies outlined in the Integration section.

Table 5. Intel - Ivy Bridge x86 based Mini systems

Model	Compute	Memory per core	Storage
W2500- 32m	32 cores	16 GB	2.4 TB SSD 24 TB HDD
W2500- 64m	64 cores	16 GB	2.4 TB SSD 24 TB HDD
W2500- 96m	96 cores	16 GB	2.4 TB SSD 24 TB HDD
W2500- 128m	128 cores	16 GB	2.4 TB SSD 24 TB HDD

Table 6. Intel Ivy Bridge x86 Enterprise systems (part 1 of 2)

Model	Compute	Memory per core	Storage
W2500- 32	32 cores	32 GB	6.4 TB 48 TB HDD
W2500- 64	64 cores	32 GB	6.4 TB 48 TB HDD
W2500- 96	96 cores	32 GB	6.4 TB 48 TB HDD
W2500- 128	128 cores	32 GB	6.4 TB 48 TB HDD
W2500- 160	160 cores	32 GB	6.4 TB 48 TB HDD

Table 6. Intel Ivy Bridge x86 Enterprise systems (part 2 of 2)

Model	Compute	Memory per core	Storage
W2500- 192	192 cores	32 GB	6.4 TB 48 TB HDD
W2500- 224	224 cores	32 GB	6.4 TB 48 TB HDD
W2500- 320	320 cores	32 GB	6.4 TB 48 TB HDD
W2500- 384	384 cores	32 GB	6.4 TB 48 TB HDD

Table 7. IBM POWER® 7 Mini systems

Model	Compute	Memory per core	Storage
W2700- 32m	32 cores	16 GB	2.4 TB SSD 24 TB HDD
W2700- 64m	64 cores	16 GB	2.4 TB SSD 24 TB HDD
W2700- 96m	96 cores	16 GB	2.4 TB SSD 24 TB HDD
W2700- 128m	128 cores	16 GB	2.4 TB SSD 24 TB HDD

Table 8. IBM POWER 7 Enterprise systems

Model	Compute	Memory per core	Storage
W2700- 32	32 cores	32 GB	6.4 TB 48 TB HDD
W2700- 64	64 cores	32 GB	6.4 TB 48 TB HDD
W2700- 96	96 cores	32 GB	6.4 TB 48 TB HDD
W2700- 128	128 cores	32 GB	6.4 TB 48 TB HDD
W2700- 160	160 cores	32 GB	6.4 TB 48 TB HDD
W2700- 192	192 cores	32 GB	6.4 TB 48 TB HDD
W2700- 224	224 cores	32 GB	6.4 TB 48 TB HDD
W2700- 320	320 cores	32 GB	6.4 TB 48 TB HDD
W2700- 384	384 cores	32 GB	6.4 TB 48 TB HDD

Ordering information

This product is only available through IBM Passport Advantage®. For more information, contact your WebSphere sales representative or organizational Passport Advantage representative.

Ordering information is shown in Table 9.

Table 9. Ordering part numbers and feature codes

Program name	PID number	Charge unit description
PureApplication System W2500	5725-Q93	Per appliance installation
PureApplication System W2700	5725-Q94	Per appliance installation
IBM Business Process Manager Pattern V8.5	5725-L40	PVU based
IBM WebSphere MQ Hypervisor Edition V7.5.0.x for Red Hat Enterprise Linux Server	5725-C79	PVU based
IBM WebSphere MQ Hypervisor Edition V7.5.0.x for AIX®	5725-F22	PVU based

Related information

For more information, see the following documents:

- IBM PureApplication System product page
  http://www.ibm.com/ibm/puresystems/us/en/pf_pureapplication.html
- A tour of the hardware in IBM PureApplication System: The second generation
  http://www.ibm.com/developerworks/websphere/techjournal/1407_woolf2/1407_woolf2.html
- Achieving business continuity in IBM PureApplication System
  http://www.ibm.com/developerworks/websphere/techjournal/1309_brown1/1309_brown1.html
- Achieving high availability across multiple sites using IBM PureApplication System
  http://www.ibm.com/developerworks/cloud/library/cl-ps-aim1305-hadrmultisites/index.html
- IBM Offering Information page (to search on announcement letters, sales manuals, or both)
  http://www.ibm.com/common/ssi/index.wss?request_locale=en
- PureSystems® Centre:
  http://www.ibm.com/ibm/puresystems/us/en/puresystemscentre.html

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 that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right 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 10504-1785 U.S.A.

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. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products. 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.

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 measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate 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.

This document was created or updated on December 12, 2014.

Send us your comments in one of the following ways:

- Use the online **Contact us** review form found at:
  ibm.com/redbooks
- Send your comments in an e-mail to:
  redbooks@us.ibm.com
- Mail your comments to:
  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400 U.S.A.

This document is available online at http://www.ibm.com/redbooks/abstracts/tips1222.html .

Trademarks