

IMS and VSAM replication using GDPS

IBM Redbooks Solution Guide

Geographically Dispersed Parallel Sysplex™ (GDPS®)/Active-Active is an IBM® solution designed to provide near-continuous availability and disaster recovery for critical business applications. It relies on software-based replication techniques to synchronize data between sites, in unlimited distances. This IBM Redbooks® Solution Guide publication describes GDPS/Active-Active and its benefits.

For IMS™ and VSAM workloads, it uses IBM InfoSphere® Data Replication (IIDR) for IMS for z/OS® and IBM InfoSphere Data Replication (IIDR) for VSAM for z/OS products, respectively, to replicate data, while GDPS manages, controls, and monitors the environment for planned and unplanned events.

The replication technology used by IIDR products consists of capture engines, apply engines, and a transport infrastructure. Capture engines monitor the source data repository for committed transaction data, which is captured and placed in a transport infrastructure for transmission to a target location. Apply engines write the data in near-real time to an active copy of the data repository. This is shown in Figure 1.

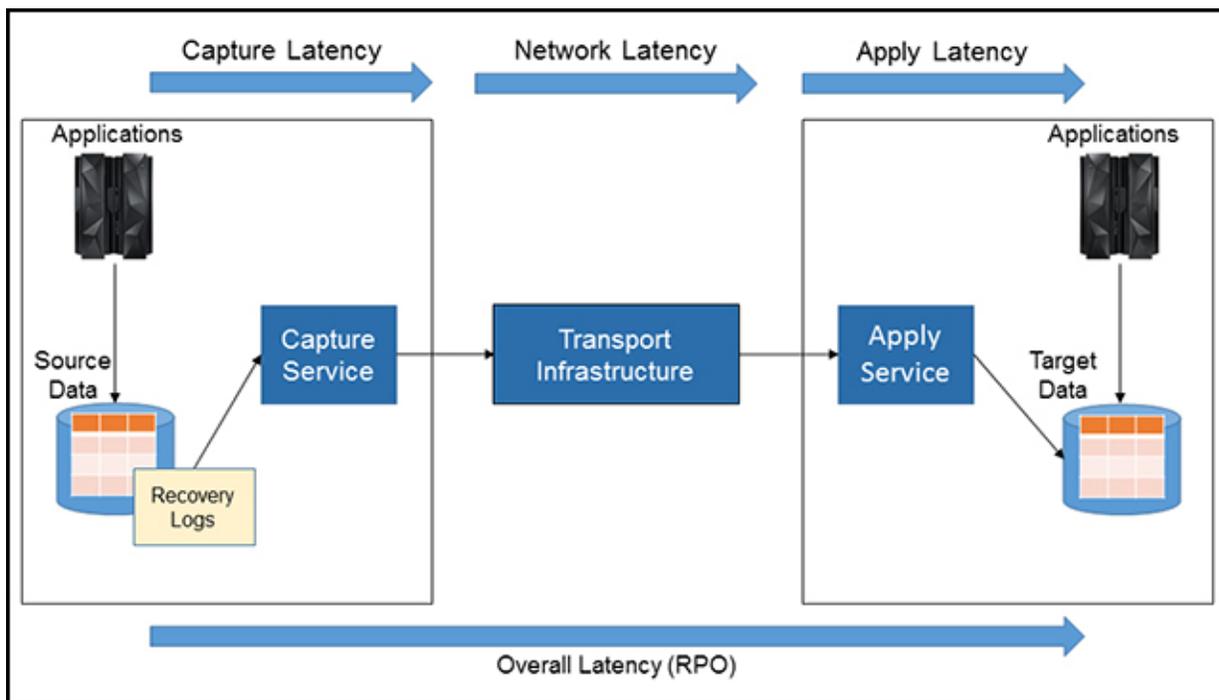


Figure 1. IBM InfoSphere Data Replication for z/OS data flow (including IMS and VSAM data)

Did you know?

GDPS/Active-Active and the IBM software-based replication products, such as IIDR for IMS and IIDR for VSAM, are not intended to be a replacement to the existing hardware-based or storage-based replication solutions, but are an option to allow customers to achieve higher IT resiliency objectives. They can also be combined with other replication techniques and GDPS offerings, such as GDPS/Peer-to-Peer Remote Copy (PPRC) and GDPS/Metro Global Mirror (MGM), to achieve even higher levels of reliability, availability, and serviceability for enterprises.

In a GDPS/Active-Active environment, qualified IMS and VSAM workloads can be automatically switched and balanced between sites depending on the solution configuration, by using the IBM Multi-site Workload Lifeline for z/OS component.

Business value

GDPS/Active-Active solution provides continuous availability for critical business applications. At the same time, the system resources and near real-time production data available on the alternate site can be used for workload balancing and data analytic workloads exploring large amount of real-time data.

To understand how IMS and VSAM replications can be valuable to clients, along with GDPS/Active-Active management and control, a few questions must be answered:

- How long can you afford to be without your critical applications?
- During planned or unplanned events, how long does it take to switch or recover your systems in an alternate location?
- What if you can run and balance your IMS and VSAM workloads in two different sites, in unlimited distances?

GDPS/Active-Active and IBM InfoSphere Data Replication for IMS and VSAM manage data replication at an application, or a workload, level of granularity, which means that they allow different workloads to have different levels of protection, performance, and design. With GDPS/Active-Active, customers can use one site as the primary site for a workload, where the data is updated, and the alternate site can be either a standby or even a query site for the same workload. This allows customers to create reports and run data analytics processing concurrently at the alternate site, while production is running and doing updates at the primary site.

IMS and VSAM replication can provide the following main benefits in a GDPS/Active-Active environment:

- Reduced planned and unplanned outages for critical applications
- Help clients meeting recovery point objectives (RPO) and recovery time objectives (RTO)
- Provide automation to allow a customer's staff to be more productive and reduce the overall support and operational skills required to manage the environment
- Simplify management of complex operating environments, monitoring, and reporting on events that can affect recovery
- Help maintain compliance with business continuity regulations
- Isolation from catastrophic failures
- Both primary and alternate sites running workloads
- Application-level granularity
- Site switch in seconds (after a failure event is detected)
- RPO of seconds/minutes (with the ability to issue reports that can be generated showing orphaned data)
- Protection against metro and regional disasters (distance between sites is unlimited)

Solution overview

GDPS/Active-Active and IBM InfoSphere Data Replication for IMS and VSAM products operate at the application level to manage the availability of the workloads and the routing of transactions between sites. GDPS acts principally as the coordination point or controller for these activities, including being a focal point for operating and monitoring the solution and readiness for recovery, and relies on IIDR products to replicate and synchronize the data repositories between the primary and alternate sites.

When the workloads are defined, the data sources that support a certain application, such as Payroll or Customers, are mapped and organized in what is called a subscription, which consists of a combination of source data sets, for VSAM replication, or source databases, for IMS replication. When replication is started, IIDR for IMS and VSAM create a mapping of the source and target data repositories and track the changes made at the source, keeping both sites synchronized in near real time.

This allows transactions to be switched to either site, depending on the workload type. Based on customized policies, GDPS, using IBM Multi-site Workload Lifeline for z/OS, directs workloads to run on a primary site or an alternate site and remains monitoring the environment for failure detections. In unplanned events, such as workload, system, or site failures, GDPS can automatically switch the failing workloads to the alternate site, in seconds. Again, because this is a workload-based configuration, each application or workload can have different goals, according to business needs.

The age of the data at the target site is monitored by GDPS and depends on the overall replication latency. This includes capture processing, transmission of the data from the source to the target site, which relies on the network transport infrastructure, and apply processing, as shown in the Figure 1 in the beginning of this guide. It will vary from customer to customer because it is business related, but the goals will typically be around a few seconds.

Solution architecture

To manage IMS and VSAM data replication, GDPS/Active-Active uses a few infrastructure components, including both hardware and software products, as shown in Figure 2.

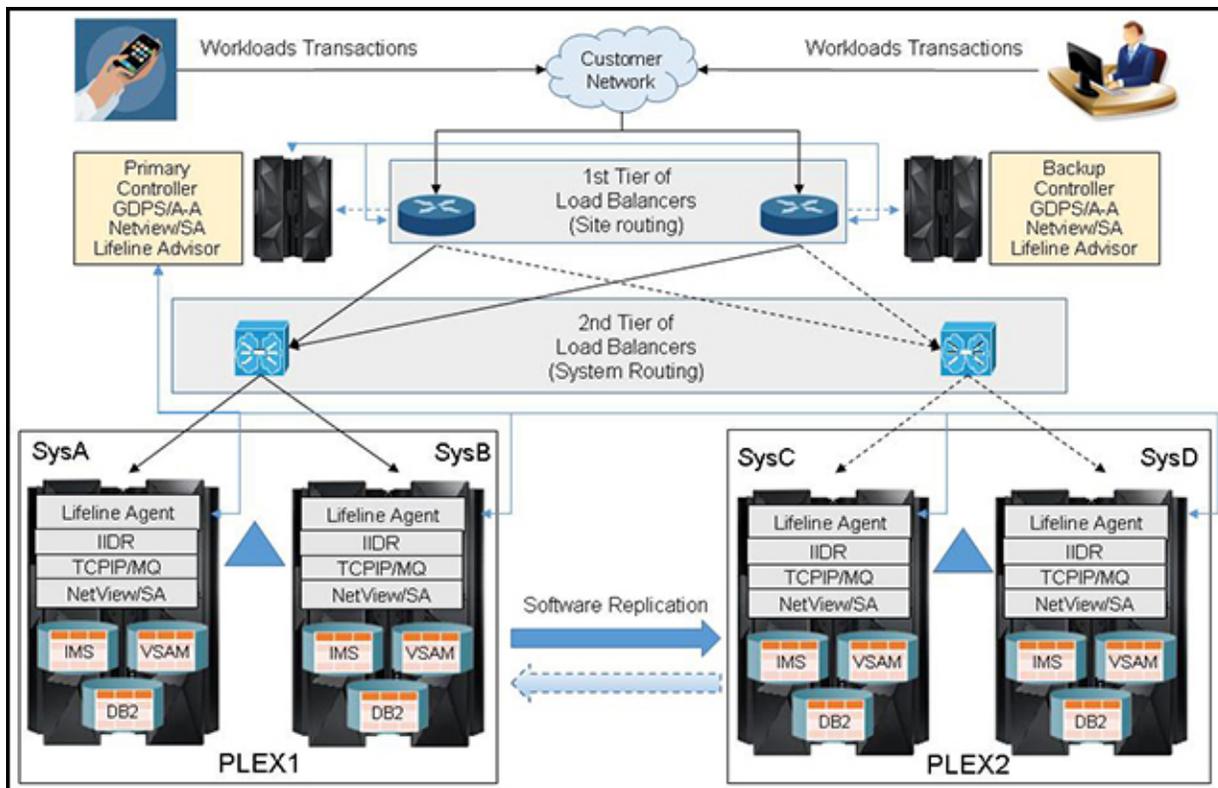


Figure 2. IBM InfoSphere Data Replication for z/OS data flow (including IMS and VSAM data)

When transactions arrive to the customer network, they are directed to a first-tier Load Balancer, which is a Server/Application State Protocol (SASP)-compliant router. It tracks the primary site for each workload based on customized policies and whether the workloads are reaching their predefined thresholds. It is monitored by the IBM Multi-site Lifeline for z/OS product, which has Agents running on every production logical partition (LPAR) and monitors the health of the server applications and the systems where the Agents run. The agents periodically send this information to the primary Lifeline Advisor, which is an address space that runs on the primary (also called Master) GDPS controlling system, having a standby address space running on the alternate GDPS controlling system, for redundancy purposes. The Advisor uses this information to calculate routing recommendations for the workloads that use these server applications, so in planned or unplanned events, it can automatically switch and route the failing workload to the alternate location.

After the workload is directed to one of the sites, it goes to a second-tier Load Balancer, which is either another hardware SASP-compliant router or a software-based router, such as a Sysplex Distributor, running on z/OS. The second-tier routing does the workload balancing within the sysplex to the application server that can effectively process the workload. An SASP-compliant second-tier load balancer is optional. If it is present, the Lifeline Advisor can manage and effect the workload distribution within the sysplex.

When transactions process and update the data at the primary site, if they are associated with a GDPS/Active-Active workload, the updates are also replicated to the alternate site through the IBM InfoSphere Data Replication products. As of today, the supported data types that can take advantage of this replication technique are DB2®, IMS, and VSAM data.

IIDR uses three main components to replicate IMS and VSAM data: the Capture server that reads updates at the primary site from the VSAM and IMS logs, the Transport Infrastructure that transmits the updates to the alternate site using TCP/IP, and the Apply Server that writes the updates to the target data repositories.

All of this data replication infrastructure and environmental health are monitored using IBM Tivoli® Monitoring and IBM Tivoli NetView® products, which are controlled by GDPS and Tivoli System Automation, and automate switch and recovery actions when needed.

Note: Although the intention of this solution guide is to cover how IMS and VSAM software replication products work when using GDPS/Active-Active, we purposely mention DB2 data replication, which is also a supported workload type of the GDPS/Active-Active solution. For more information, see the IBM Redbooks publication, *GDPS Family - An Introduction to Concepts and Capabilities*, SG24-6374.

Usage scenarios

At this time, GDPS/Active-Active solution supports software-based replication for DB2, IMS, and VSAM data types. It allows customers running any of these workloads to take advantage of the GDPS/Active-Active benefits. An example of how the solution might fit into the customer's environment is shown in Figure 3.

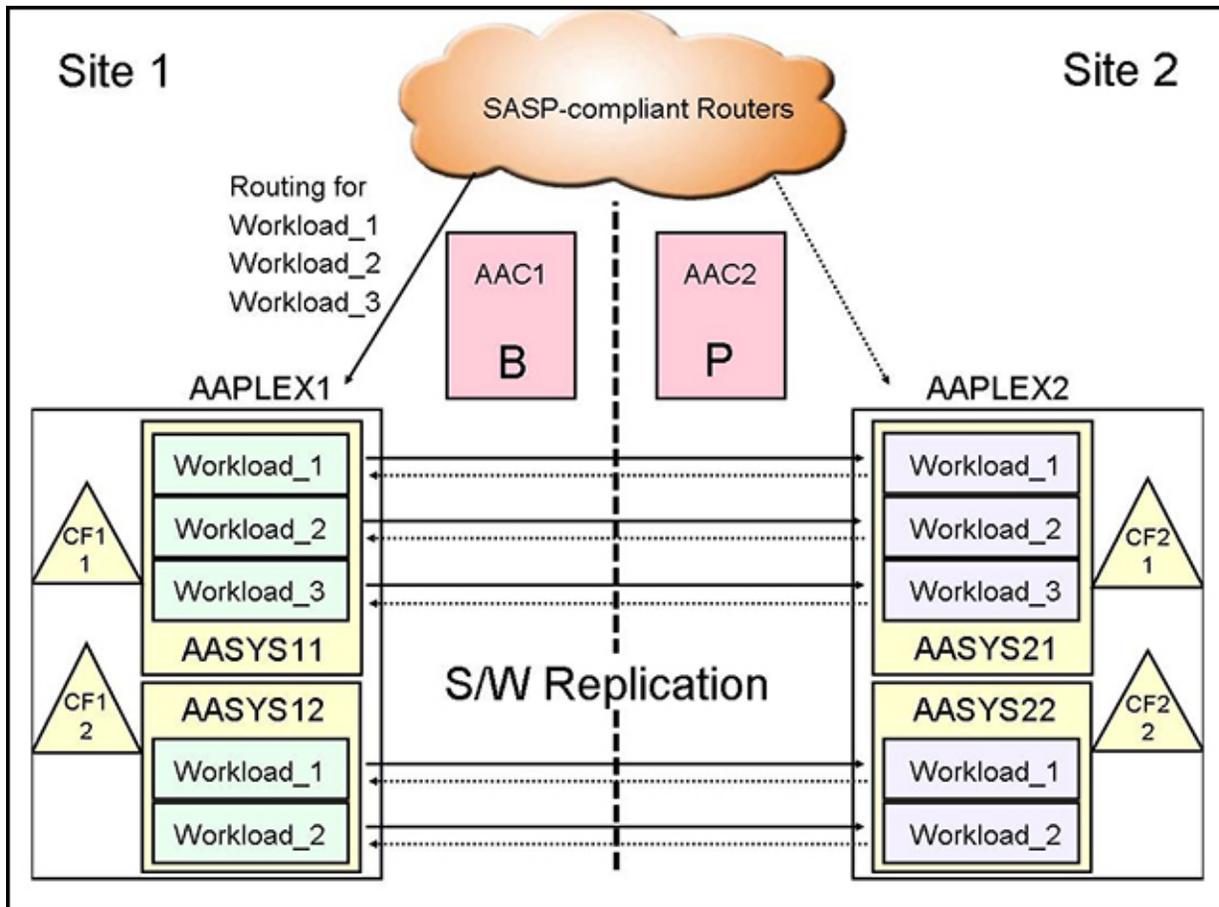


Figure 3. GDPS/Active-Active environment with different workloads active in different sites

Figure 3 demonstrates the following information:

- Two z/OS sysplexes with two z/OS partitions each. AAPLEX1 is running at Site 1 and has the z/OS partitions AASYS11 and AASYS12. AAPLEX2 is running at site 2 and has the z/OS partitions AASYS21 and AASYS22, separated geographically by virtual unlimited distance.
- The router cloud shows to which site the transactions for each of the workloads is routed. In this scenario, AAPLEX2 is the active sysplex for Workload_2. AAPLEX1 is the active sysplex for Workload_1 and Workload_3. To make it easier to differentiate which site is the active site for a workload, in the figure, it is shaded green.
- Replication of the data that belongs to Workload_2 is from AAPLEX2 to AAPLEX1. Replication of the other two workloads, Workload_1 and Workload_3, are going to the opposite direction, from AAPLEX1 to AAPLEX2. This is shown by the solid arrows, which display the active replication direction.

You might also hear the term “dual active/active” used to describe this kind of an environment where both sites/sysplexes are actively running different workloads, but each workload is active/standby. However, it is also possible to have one site with an active update workload (read/write processing) and the other site active for Query (read-only processing) of the same workload. For example, in Figure 3, Workload_1 is active for update (read/write) processing at Site 1, replicating the changes to Site 2. However, you might also have Site 2 with active Query (read-only) processing for Workload_1, as long as it does not perform updates to that data. You can associate up to two query workloads with a single update workload. Workload distribution between the sites is based on policy options, and takes into account environmental factors, such as the latency for replication that determines the age (or currency) of the data in the standby site.

In a planned or unplanned event that will cause impacts to Workload_1, for example, an outage at AAPLEX1, this is a failure that is detected and, based on your workload failure policy, can trigger an automatic switch of the failed workload to AAPLEX2, which is the standby (for update processing) for that workload.

However, if you do not want GDPS to perform automatic workload switch for failed workloads, you can also select the option of an operator prompt. The operator is prompted whether GDPS is to switch the failed workload or not. If the operator accepts switching of the workload, GDPS will perform the necessary actions to switch the workload. For this kind of switch resulting from a workload failure, whether automatic or operator confirmed, no pre-coded scripts are necessary. GDPS understands the environment and performs all the required actions to switch the workload.

This allows the switch of workloads from sites separated by unlimited distances typically in seconds.

Integration

GDPS/Active-Active is a solution that relies on other IBM software products, each providing a specific function, and acts as a single point of control, giving a high level of detail of the environment at the application level. The following IBM software products that integrate this solution are required products:

- GDPS/Active-Active V1.4
- IBM Tivoli System Automation for z/OS V3.3 (+SPE APAR) or V3.4
- IBM Tivoli NetView for z/OS V6.2
- IBM Tivoli NetView Monitoring for GDPS V6.2
- IBM Multi-site Workload Lifeline for z/OS V2.0
- IBM InfoSphere Data Replication for DB2 for z/OS V10.2.1
- IBM InfoSphere Data Replication for IMS for z/OS V11.1
- IBM InfoSphere Data Replication for VSAM for z/OS V11.1
- IBM Tivoli Monitoring V6.3
- IBM Tivoli Management Services for z/OS V6.3
- z/OS V1.13 or higher

- WebSphere® MQ V7.0.1 - required for DB2 data replication
- CICS® Transaction Server for z/OS V5.1 - required for VSAM data replication
- CICS VSAM Recovery for z/OS V5.1 - required for VSAM data replication
- DB2 for z/OS V9 or higher - workload dependent
- IMS V11 - workload dependent

GDPS/Active-Active offers monitoring integration with the IBM Tivoli OMEGAMON® XE family of monitoring products for various parts of the solution:

- IBM Tivoli OMEGAMON XE on z/OS
- IBM Tivoli OMEGAMON XE for Mainframe Networks
- IBM Tivoli OMEGAMON XE for Storage
- IBM Tivoli OMEGAMON XE for DB2 Performance Expert (or Performance Monitor) on z/OS
- IBM Tivoli OMEGAMON XE on CICS for z/OS
- IBM Tivoli OMEGAMON XE on IMS
- IBM Tivoli OMEGAMON XE for Messaging

It also provides integration support with GDPS/MGM and cooperation support with GDPS/PPRC.

Supported platforms

GDPS/Active-Active solution with IMS and VSAM replication is supported under the z/OS platform. It replicates data between two separate sysplexes, in unlimited distances. It is also necessary to have two monoplexes, one in each location, for the GDPS/Active-Active Controlling systems because they must not share components with the production workloads.

Also, to allow workload routing and switching between the sites, a Server/Application State Protocol (SASP)-compliant router is also required. This capability, described in RFC 4678, enables GDPS (through the IBM Multi-site Workload Lifeline for z/OS) to instruct the router to direct transactions to one site or the other. Check with your network provider for the best solution to fit within your existing network.

NetView Web Application, a feature that is part of the IBM Tivoli NetView for z/OS product, requires a server for hosting the application that can run under the Microsoft Windows, IBM AIX®, or Linux (including Linux on System z®) operating systems. Details of the supported server platforms, operating systems, and web browsers is in the NetView V6R2 documentation. In addition to the physical server platform considerations, also plan for redundancy in your application server infrastructure by providing at least one instance of the NetView Web Application in each of your sites to ensure that the Web Application is available in the event of a failure in either site.

In addition to the described requirements, there are other IBM System z Hardware Management Console (HMC), z/OS-provided Base Control Program Internal Interface (BCPii), and IBM System z Support Element (SE) considerations. More information about the hardware and software prerequisites is in the *GDPS/Active-Active Planning and Implementation Guide, ZG24-1767*.

Ordering information

The GDPS offerings are a combination of IBM products and services, provided by the IBM Global Technology Services® organization. The IBM installation services for GDPS are listed:

- Assists in planning, configuring, and automation code customization
- Provides onsite assistance
- Provides an automated, cross-platform disaster recovery solution (GDPS/PPRC, GDPS/PPRC HyperSwap® Manager, GDPS/XRC, GDPS/GM, GDPS/MGM, GDPS/MzGM, and GDPS/Active-Active)
- Includes onsite delivery, configuration, implementation, and testing
- Provides training for your support staff
- Provides centralized management of your data replication and recovery environment using automated technologies to help provide an end-to-end disaster recovery solution
- Project management and support throughout the engagement, and might include assistance to help you implement any prerequisite software

For questions regarding ordering, pricing, or any other GDPS-related information, contact your local IBM representative or send an email to: gdps@us.ibm.com

Related information

For more information, see the following documents:

- IBM Redbooks: *GDPS Family - An Introduction to Concepts and Capabilities*, SG24-6374
<http://www.redbooks.ibm.com/abstracts/sg246374.html>
- GDPS product page
<http://www-03.ibm.com/systems/z/advantages/gdps/resources.html>
- IBM Offering Information page (announcement letters and sales manuals). On this page, enter GDPS, select the information type, and then click **Search**. On the next page, narrow your search results by geography and language:
http://www.ibm.com/common/ssi/index.wss?request_locale=en

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.

© Copyright International Business Machines Corporation 2014. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

This document was created or updated on November 17, 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/tips1221.html> .

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at : <http://www.ibm.com/legal/copytrade.shtml>

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

AIX®
CICS®
DB2®
GDPS®
Geographically Dispersed Parallel Sysplex™
Global Technology Services®
HyperSwap®
IBM®
IMS™
InfoSphere®
NetView®
OMEGAMON®
Redbooks®
System z®
Tivoli®
WebSphere®
z/OS®

The following terms are trademarks of other companies:

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

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

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

