



Richard Tran
Sachchidanand Singh
Poonam Chitale
Tushar Patil

IBM Cognos Business Intelligence 10.2.0 Reporting on IMS

Overview

The IBM IMS™ database stores data using a hierarchical model that differs from the relational databases like IBM DB2. A typical IMS database stores the data using blocks called *segments* and each segment can contain several pieces of data called *fields*. Most organizations have their operational databases stored in the IMS databases; now they can use IBM Cognos® to make informed decisions over the IMS database.

This IBM Redpaper™ provides guidance for consuming IMS databases through the IBM Cognos Business Intelligence tool.

IBM Cognos BI integration with IMS

IMS is a database and transaction management system introduced by IBM in 1968. An IMS database contains high-volume transactional data from different industry domains like banking, finance, manufacturing, retailing, insurance, refinance, healthcare, aerospace, government, communications, and others.

The IMS data contains useful information about customer and business trends that are hidden in the huge volume of data. And in today's competitive business market prompt decision making is a key factor for any organizations BI requirements.

IBM Cognos connects to the IMS database using IMS JDBC drivers. Figure 1 shows the high-level architecture.

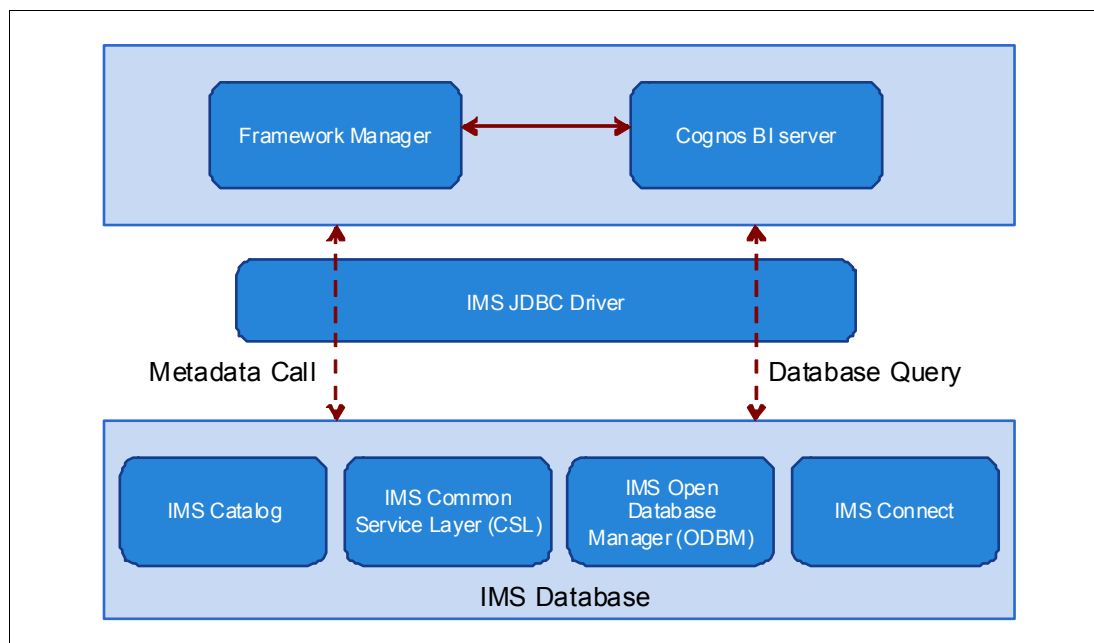


Figure 1 High-level architecture

IMS data source connection within IBM Cognos Administration

The first step, before you can do anything with IMS in Cognos, is to copy IMS JDBC driver to IBM Cognos Business Intelligence Server and then create a new data source connection in the IBM Cognos Administration.

Procedure

Complete the following steps to create an IMS data source in Cognos; these steps can be repeated as many times as needed for as many connections as required:

1. To create a new data source in the IBM Cognos Administration, select JDBC in connection type and then click **Next**.
2. In the Type field, select **IBM IMS** (Figure 2 on page 3).

Specify the JDBC connection string - New Data Source wizard

Edit the parameters to build a JDBC connection string.

Type:
 IBM IMS

JDBC URL:

Connection properties:

Driver class name:
 com.ibm.ims.jdbc.IMSDriver

Connection string:
 ;LOCAL;JDBC;URL=jdbc:ims;;DRIVER_NAME=com.ibm.ims.jdbc.IMSDriver;

Figure 2 JDBC connection string for IMS

3. In the JDBC URL field, enter the IBM IMS server URL in the following format:

//<host>:<port>/<DBName>

User ID and Password are optional fields (Figure 3) used for authorization. They can be specified either in the connection URL or in the Cognos Authorization box. The authorization parameters here correspond to the RACF permissions:

//<host>:<port>/<DBName>:user=<username>;password=<password>

Signon

Select whether a user ID and password is required in the connection string and, if so, whether to create a signon.

☐ No authentication

☐ An external namespace:

☒ Signons

☐ Password

☒ Create a signon that the Everyone group can use:

User ID:

Password:

Confirm password:

Figure 3 Signon fields

The following line is an example:

//localhost:10000/default:user=admin;password=pwd123;

4. To generate the trace file that can be used to debug issues, add details to the connection properties, described in the next section.

Connection properties

Add these connection properties:

- `traceFile`

Optional: Specifies the name that the trace should be stored in.

- `traceFileAppend`

Optional: Set this property to `true` to append the log data to the end of an existing trace; if set to `false`, the existing trace file is overwritten.

- `traceDirectory`

Optional: Specifies the directory where the trace file is located.

- `traceLevel`

Optional: Specifies the trace level to be recorded. A value of `-1` enables all trace.

- Connection string

```
;LOCAL;JDBC;URL=JDBC:ims://<host>:<portno>/<DBName>;DRIVER_NAME=com.ibm.ims.JDBC.IMSDriver;traceFile=myTrace.log;traceFileAppend=true;traceLevel=-1
```

How hierarchical IMS database is converted to relational

The hierarchical model in IMS is implemented using blocks of data known as *segments*. Each segment can contain several pieces of data, which are called *fields*.

The IMS database segment defines the fields for a set of segment instances in the same way a relational table defines columns for a set of rows in a table. Therefore, segments relate to relational tables and fields in a segment relate to columns in a relational table. The name of an IMS segment becomes the table name in an SQL query and the name of a field becomes the column name in the SQL query. For details, see the “Comparison of hierarchical and relational databases” topic:

http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.ims11.doc.apg/ims_comparehierandrel dbs.htm

Key feedback area (KFBA) contains the concatenated keys of the segments in a logical database record; sequence is from top to bottom. The KFBA is how IMS positions its cursor to the correct record that you are trying to retrieve or update. It provides the referential integrity you need to properly point to the correct resource.

An example of the three tables is shown in Figure 4 on page 5.

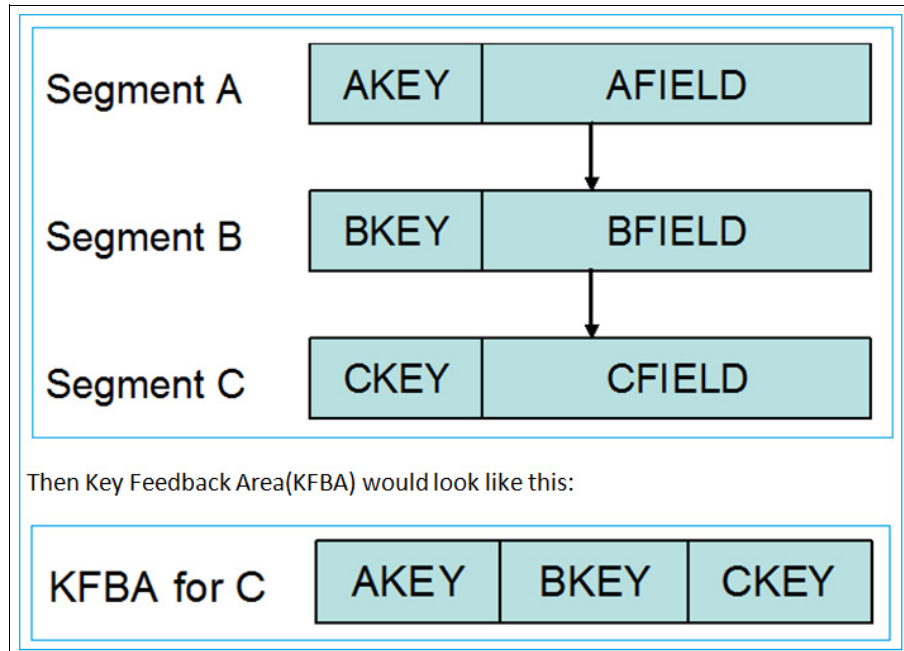


Figure 4 IMS key feedback area

To find a specific instance of C, the user must specify the primary key CKEY=CK1 and the foreign keys BKEY=BK1 and AKEY=AK1. The primary and foreign key values are used for referential integrity.

At the completion of a retrieval call, IMS identifies the location of the record by concatenating the primary and foreign keys in the key feedback area. For insert calls, IMS builds out the KFBA to position the inserted record.

IBM Cognos metadata modeling over IMS

IMS provides a Java Database Connectivity (JDBC) driver for SQL-based database connectivity to access IMS databases over TCP/IP, with the IMS Universal JDBC driver that is included in the IMS Universal drivers. The IMS Universal JDBC driver is based on the JDBC 3.0 standard.

Although IBM Cognos BI query engine is aware of the SQL syntax that is supported by the IMS Universal JDBC driver, restrictions are imposed by the IMS Universal JDBC driver related to the context of an expression or column references, which cannot be detected and can cause parsing errors. For details, see the “IBM Business Analytics Proven Practices: Best Practices for Using IMS Data in IBM Cognos BI Reports” topic:

http://www.ibm.com/developerworks/library/ba-pp-infrastructure-cognos_specific-page630/

For example, see the restriction of portable SQL keywords by the classic IMS JDBC drivers as described in the “SQL keywords supported by the IMS JDBC driver” topic:

http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/index.jsp?topic=%2Fcom.ibm.ims13.doc.apg%2Fims_sqlkeywordsjdbc.htm

In that topic, and to read more about metadata modeling, see the “Configuring IBM Cognos BI for IMS data sources” subtopic.

Performance considerations for IMS reporting

Cognos can access all of the various IMS data types through the JDBC driver, with the exception of MSDB, which has been functionally stabilized.

IMS is a popular online transaction processing (OLTP) systems in the world, with more than 50 billion transactions run through IMS databases each day.

IMS databases have hierarchical structure with optimized storage and retrieval, and with a high level of data integrity and recovery. The hierarchical data structure relationships are predefined, having access paths distinct and ready to use by the application. IMS offers ultimate data access speed because of its existing mapped-out roadway to navigate data structures.

IMS 13 can deliver unprecedented performance gains:

- ▶ 117,000 transactions per second on a single IMS Fast Path system.
- ▶ CPU time reductions of up to 62% in some Java applications, leveraging IMS 13 Universal Database Type-2 access.
- ▶ CPU reductions of up to 15% for certain zIIP-enabled general-purpose processing workloads. For details, see *Developing a blueprint for IBM IMS application modernization*:
<http://public.dhe.ibm.com/common/ssi/ecm/en/imw14726usen/IMW14726USEN.PDF>

Sample reports created using IMS database

The following figures are Cognos sample reports, created using the IMS database. Figure 5 shows a sales commissions report for the central Europe branches in the year 2005.




Figure 5 Sales commission report

The report in Figure 6 shows how satisfied employees are with their jobs. The report is built using data present in IMS database.




Figure 6 Employee survey score

The report in Figure 7 uses colored symbols to show an aggregation of how happy employees are by department for different metrics.




Figure 7 Department survey score

Troubleshooting considerations

Cognos can be installed on multiple Java EE servers.

WebSphere Application Server configuration

For IBM WebSphere® Application Server, certain settings are required.

First, set the JVM environment variable (`javax.xml.bind.JAXBContext`) in the IBM Websphere administrative console, as shown in Figure 8.

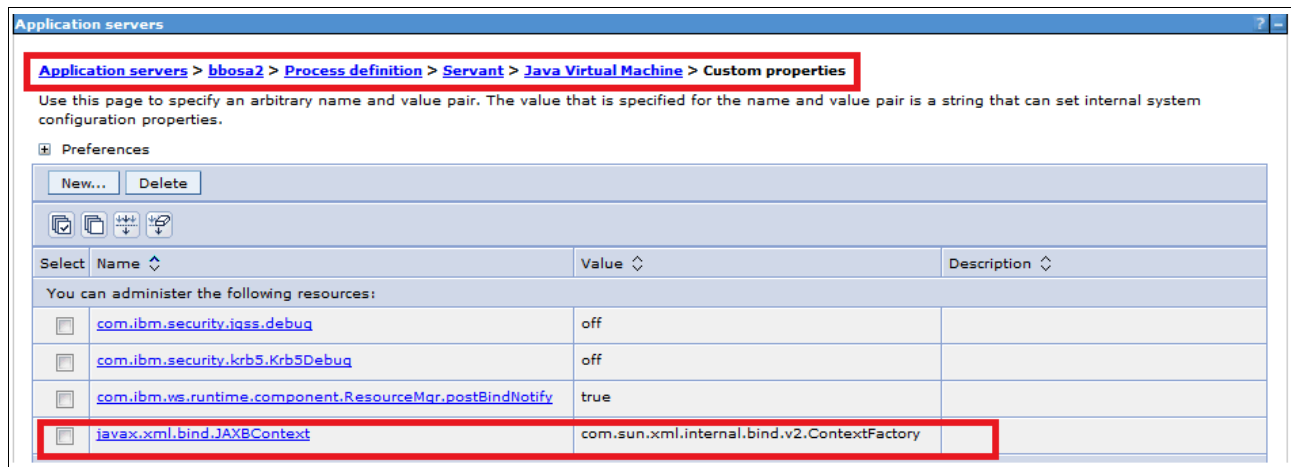


Figure 8 IBM WebSphere custom properties

For IBM Cognos® z/OS® deploying on WebSphere Application Server for z/OS, Cognos z/OS V10.2.1 is required in order to properly retrieve metadata from the IMS catalog.

Two methods enable tracing the IMS database:

- ▶ Using JRE logging.properties
- ▶ Configuring J2EE tracing

Method 1: Using JRE logging.properties

The trace can be enabled by setting the trace level for the IMS Universal drivers loggers in logging.properties file in the JRE lib folder. For example, in a WebSphere Application Server environment, the logging.properties is in the following locations.

- ▶ For 32 Java: <WAS base> /java/Java version/lib
- ▶ For 64 Java: <WAS base> /java64/java version_64/lib

Figure 9 on page 9 shows the setting of the log detail level for tracing.

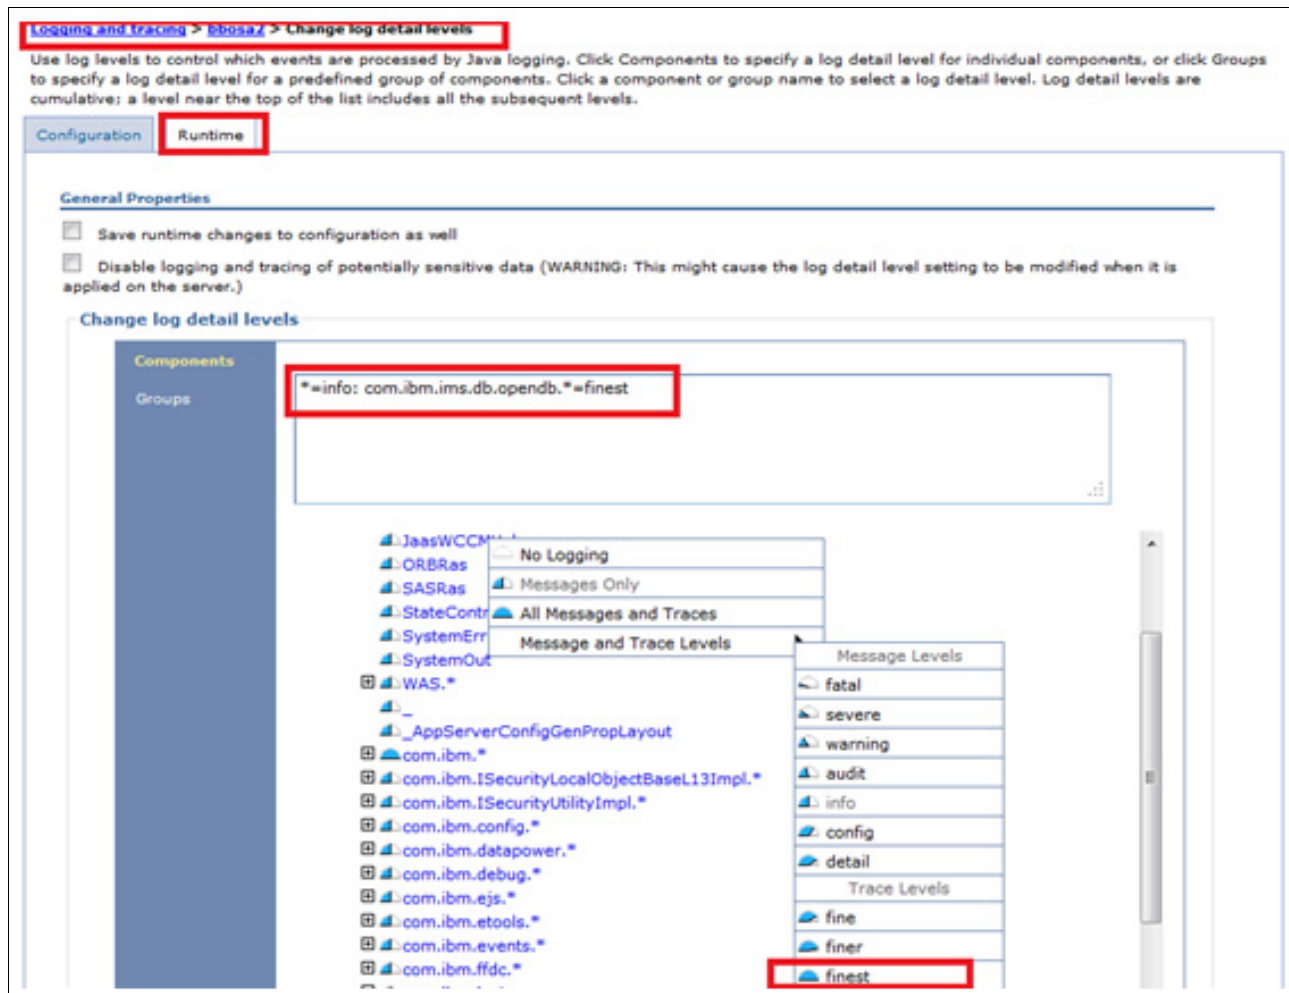


Figure 9 Log detail level for tracing

The recommended trace level is FINEST.

To set the trace for all IMS Universal drivers loggers, add the following line to the logging.properties file:

```
com.ibm.ims.db.opendb.level = FINEST
```

To send the trace output to a file, add the following lines to your logging.properties file:

```
java.util.logging.FileHandler.level = FINEST
java.util.logging.FileHandler.pattern = <C10>/logs/UniversalDriverTrace.txt
java.util.logging.FileHandler.formatter = java.util.logging.SimpleFormatter
```

Method 2: Configuring J2EE tracing

Tracing can be enabled in the WebSphere Application Server from the administrative console. The IMS Universal DB resource adapter must be deployed on the WebSphere Application Server before tracing can be configured.

To get the most detailed trace from the IMS Universal DB resource adapter, use the following steps:

1. Start the WebSphere Application Server administration console.
2. Select **Troubleshooting**.

3. Select **Logs and Trace**.
4. Select your application server from the table.
5. Select **Change log detail levels**.
6. Select the **Runtime** tab. Complete the following steps:
 - a. Make sure that the **Save runtime changes to configuration** check box is selected.
 - b. Under the Change Log Details Levels section, select the **com.ibm.ims.db.opendb.*** component. This Message and Trace levels menu opens. Do the following steps:
 - i. Select the message level FINEST.
 - ii. Click **Apply**.
 - c. To save these changes for the next time the application server is started, click **Save** at the top of the page. IBM WebSphere Application Server does not need to be restarted.

Enabling trace for the IMS Universal JDBC driver

The IMS Universal JDBC driver runs as a separate component from Cognos, and additional configurations beyond the standard Cognos logging settings are needed to generate JDBC trace.

Conclusions

This paper highlights the importance of IBM Cognos reporting on the IMS database. The IMS database is built on IBM System z® and delivers the lowest cost transaction and hierarchical database management system for mission-critical OLTP. Over 90% of the fortune 500 companies are using IMS to run their daily business operations. For business intelligence (BI) reporting using IBM Cognos, the data present in an operational database should be warehoused first. The BI reporting over IMS database can be a key enabler and crucial factor in the organizational success.

References

For additional information, see the following publications:

- *IMS Integration and Connectivity Across the Enterprise*, SG24-8174
- *IMS Cognos Dynamic Query*, SG24-8121
- Comparison of hierarchical and relational databases:
http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/topic/com.ibm.ims11.doc.apg/ims_comparehierandrel dbs.htm
- IBM Business Analytics Proven Practices: Best Practices for Using IMS Data in IBM Cognos BI Reports:
http://www.ibm.com/developerworks/library/ba-pp-infrastructure-cognos_specific-page630/
- SQL keywords supported by the IMS JDBC drivers:
http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/index.jsp?topic=%2Fcom.ibm.ims13.doc.apg%2Fims_sqlkeywordsjdbc.htm
- *Developing a blueprint for IBM IMS application modernization*:
<http://public.dhe.ibm.com/common/ssi/ecm/en/imw14726usen/IMW14726USEN.PDF>

Authors

This paper was produced by a team of specialists working at the Silicon Valley Laboratory, San Jose.

Richard Tran is an Advisory Software Engineer working as the IMS Open Database team lead. He is involved in integrating IMS with other business solutions such as Business Analytics and Reporting, Content Discovery, and Web Enablement.

Sachchidanand Singh is a Cognos Consultant at IBM India, working on enablement and implementation of the IBM Cognos BI product suite for customers. In the past, he was involved in the mainframe application development.

Poonam Chitale is a Software Engineer working on the IMS Open Database team. She is primarily involved in function test and web technologies for IMS productivity tools. In the past, she worked in the IBM Rational® Quality and Requirements Management area.

Tushar Patil is a staff Software Engineer at IBM India, working on enablement and implementation of Cognos BI product suite for customers.

Thanks to the following contributor to this paper:

Paolo Bruni
ITSO, Silicon Valley Lab, San Jose

Now you can become a published author, too!

Here is an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:
<http://www.facebook.com/IBMRedbooks>
- ▶ Follow us on Twitter:
<http://twitter.com/ibmredbooks>
- ▶ Look for us on LinkedIn:
<http://www.linkedin.com/groups?home=&gid=2130806>
- ▶ Explore new IBM Redbooks® publications, residencies, and workshops with the IBM Redbooks weekly newsletter:
<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:
<http://www.redbooks.ibm.com/rss.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 grant 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites 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.

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.

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.

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 REDP-5091-00 was created or updated on March 10, 2014.



Send us your comments in one of the following ways:

- ▶ Use the online **Contact us** review Redbooks form found at:
ibm.com/redbooks
- ▶ Send your comments in an email 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.




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:

Cognos®
IBM®
IMS™
Rational®

Redbooks®
Redbooks (logo) ®
Redpaper™
System z®

WebSphere®
z/OS®

The following terms are trademarks of other companies:

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