


New Approach to Analytics for IBM IMS Data

Deepak Kohli



 **Analytics**

Information Management

Find and read thousands of IBM Redbooks publications

- ▶ Search, bookmark, save and organize favorites
- ▶ Get personalized notifications of new content
- ▶ Link to the latest Redbooks blogs and videos

Get the latest version of the Redbooks Mobile App



Download
Now

Android



Promote your business in an IBM Redbooks publication

Place a Sponsorship Promotion in an IBM® Redbooks® publication, featuring your business or solution with a link to your web site.

Qualified IBM Business Partners may place a full page promotion in the most popular Redbooks publications. Imagine the power of being seen by users who download millions of Redbooks publications each year!



ibm.com/Redbooks

About Redbooks → Business Partner Programs

THIS PAGE INTENTIONALLY LEFT BLANK

New Approach to Analytics for IBM IMS Data

IBM® Information Management System (IMS™) applications and data are the core of critical online transaction processing (OLTP) workloads for many of the world's major organizations. This operational data, when analyzed properly, forms the basis for making better decisions by organizations running IMS. With IBM DB2® Analytics Accelerator for z/OS®, you can exploit your IBM z Systems™ platform's IMS data where it originates so that delivering new insights to improve efficiency and drive smart outcomes is possible (Figure 1).

Critical business insights that are gained by performing analytics on IMS operational data is a valuable corporate asset and must be delivered efficiently across an organization, with high quality and proper governance, which is possible with this solution.



Figure 1 IBM z Systems platform with DB2 Analytics Accelerator for z/OS

This IBM Redbooks® Solution Guide describes DB2 Analytics Accelerator for z/OS and how it enables you to exploit the IMS data. It explains the business value of the solution, provides an overview and high-level solution architecture and includes usage scenarios.

Did you know?

Business analytics are no longer considered a *nice to have* option. Business analytics have become mission-critical applications with the same quality of service (QoS) requirements as traditional OLTP workloads. The applications must be reliable, available, and secure: attributes that are synonymous with IBM z Systems platforms and IMS.

IBM DB2 Analytics Accelerator for z/OS is a high-performance appliance that integrates the IBM z Systems infrastructure with IBM PureData® for Analytics, powered by IBM Netezza® technology. With the DB2 Analytics Accelerator for z/OS, query elapsed times can be significantly reduced, CPU cost for query processing can be moved to DB2 Analytics Accelerator for z/OS, and storage cost for stable (no longer modified) data can be reduced by moving that data onto less costly disks attached to the DB2 Analytics Accelerator for z/OS. These items are all accomplished while protecting the z Systems investment.

Business value

Getting answers quickly is vital in this fast-paced world. The DB2 Analytics Accelerator for z/OS using Netezza technology, combined with IBM DB2 for z/OS, creates a solution capable of, in many cases, increasing query processing. DB2 Analytics Accelerator for z/OS enables the completion of analytics queries in seconds and used to take hours. Most important, it allows queries that would never have been possible because of their complexity to produce an answer, in only a few minutes.

Although not its primary purpose, DB2 Analytics Accelerator for z/OS can reduce mainframe computing resources usage. For example, if a query ran for two hours on the mainframe and is now run on the DB2 Analytics Accelerator for z/OS, it both completes in a significantly shorter elapsed time, and computing resources that would have been consumed by running that query on the mainframe are available for other workloads. While the name DB2 Analytics Accelerator for z/OS might suggest it is only for DB2 data, it can be used with all z Systems data sources including IMS, DB2, and Virtual Storage Access Method (VSAM) data.

Other benefits of the DB2 Analytics Accelerator for z/OS include these:

- ▶ No application rewrites or new database designs are required.
- ▶ It uses simplified database design, resulting in no need for indexes.
- ▶ The effort to tune analytic queries is eliminated.
- ▶ The IBM DB2 for z/OS data security attributes are inherited.
- ▶ The High Performance Storage Systems (HPSS) feature improves access to and lowers the cost of storing, managing, and processing historical data.
- ▶ Latency is minimized.
- ▶ Fast deployment and quick time-to-value are inherent to the solution.

Solution overview

This solution is designed to minimize impacts to the existing z Systems platform and the business applications that rely on IMS data. The key architectural principles of this solution are as follows:

- ▶ IMS data never leaves the z Systems platform.
- ▶ IMS workload remains unaffected.
- ▶ Provides a single server for z Systems related analytics.
- ▶ Joining of IMS and DB2 data and many other types of data in the same query is possible.
- ▶ No need to extract, transform, and load (ETL) DB2 and IMS data off the z Systems platform.
- ▶ Significant IBM z Systems - Integrated Information Processor (zIIP) offload to reduce the cost of processing large amounts of IMS data.

Implementation using DB2 Analytics Accelerator for z/OS on IMS data requires these items:

- ▶ The loading of IMS data into the DB2 Analytics Accelerator for z/OS.
- ▶ IMS database metadata (segment/table metadata) is defined in the DB2 Catalog.

After setup is complete, queries are submitted directly to DB2 (and not to IMS); DB2 can then off load the query to DB2 Analytics Accelerator for z/OS (Figure 2).



Figure 2 Implementation of DB2 Analytics Accelerator for z/OS with IMS data and performing queries

Routing queries through DB2 (not IMS) and having the queries run on the DB2 Analytics Accelerator for z/OS provides several advantages:

- ▶ A single point of entry for z Systems platforms enables a wide variety of analytics and reporting queries.
- ▶ Clients have the ability to do joins between IMS and DB2 data.
- ▶ Minimal impact occurs to IMS performance.
- ▶ DB2 (and z/OS) control access to the IMS and DB2 data.
- ▶ Significant performance gains can be achieved for analytical queries that access large sets of both DB2 and IMS data.

Solution architecture

The entire process of extracting IMS data, performing data transformation, loading the data into DB2 Analytics Accelerator for z/OS, and defining the IMS metadata into the DB2 Catalog can all be done using the IBM DB2 Analytics Accelerator Loader for z/OS Version 2.1 (5639-OLE). DB2 Analytics Accelerator Loader for z/OS Version 2.1 helps simplify, increase the speed, and automate the process of loading non DB2 data, including IMS data, directly into DB2 Analytics Accelerator for z/OS. In fact, as shown in Figure 3, it can load data from many sources: IMS, VSAM, sequential files, IBM Distributed Relational Database Architecture™ (IBM DRDA®) data sources, and Oracle. The roadmap indicates many more data sources could be included in the future.

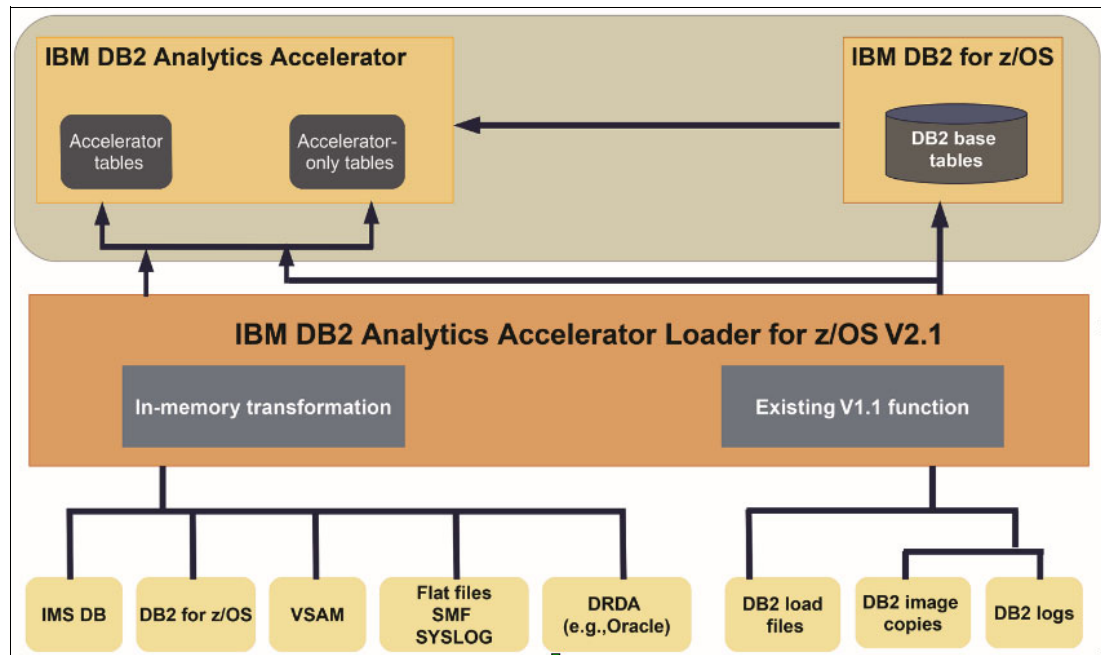


Figure 3 Data sources that can be loaded into DB2 Analytics Accelerator for z/OS using DB2 Analytics Accelerator Loader for z/OS V2.1

DB2 Analytics Accelerator Loader for z/OS V2.1 involves the following steps for loading IMS data into DB2 Analytics Accelerator for z/OS:

1. Describe your IMS databases and their data layout by creating virtual tables or views.
2. Build an SQL statement that selects from these virtual tables or views.
3. Build a job to load the result set of the SQL statement into DB2 Analytics Accelerator for z/OS.
4. Execute the batch job to perform the load.

You perform the first three steps by using IBM DB2 Analytics Accelerator Studio, a IBM Data Studio graphical user interface plug-in that is included with the product.

Let's dig into the steps a little deeper, describing the process in the following sections.

Describe your IMS databases and their data layout by creating virtual tables or views

Because IMS is a hierarchical database, and DB2 and DB2 Analytics Accelerator for z/OS are relational databases, when the data is extracted from IMS, it must be converted to a relational table record format before it can be loaded into DB2 Analytics Accelerator for z/OS.

DB2 Analytics Accelerator Loader for z/OS V2.1 defines and creates virtual tables to represent non-relational data sources, such as IMS. Specifically for IMS, a virtual table is created for each IMS segment in a database. To create a virtual table, you point DB2 Analytics Accelerator Loader for z/OS V2.1 at the COBOL or PL/I copybook that describes the data format, Program Specification Block (PSB), and database description (DBD). Each field in the segment is represented as a column in the virtual table. Figure 4 shows how to create virtual tables for IMS by importing a COBOL or PL/I copybook that describes the data layout.

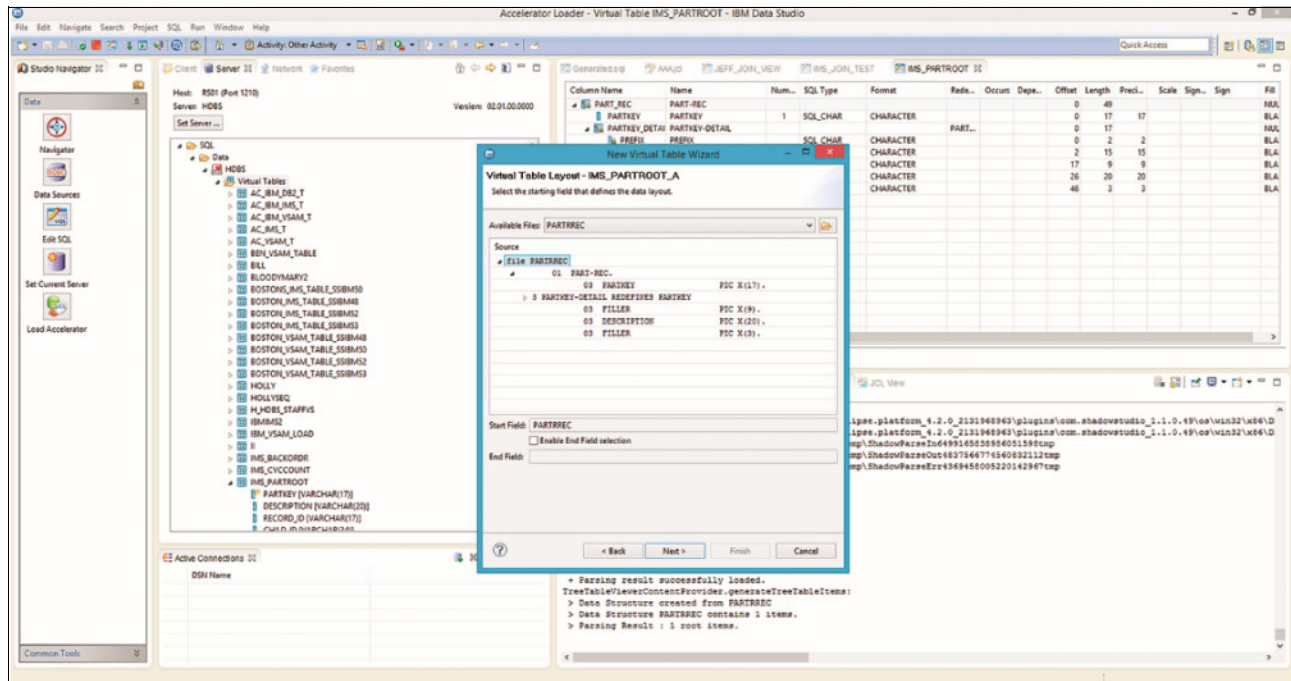


Figure 4 Creating virtual tables for IMS by importing descriptions of the data layout

The relationships between the segment tables are maintained with primary and foreign keys that are assigned automatically by DB2 Analytics Accelerator Loader for z/OS V2.1.

Creating virtual tables or views is a one-time process that needs to be updated only if the IMS data format changes.

Build an SQL statement that selects from these virtual tables or views

After you create your virtual tables or views, which describe the format of the IMS data, the next step is to create an SQL statement that identifies the IMS data to be loaded into DB2 Analytics Accelerator for z/OS. From within DB2 Analytics Accelerator Studio, you can build a *select statement* by selecting columns, which are in essence IMS data fields, from the virtual tables or views. You can add a WHERE clause to specify a subset of data to be loaded or data can be joined from multiple IMS databases

In addition, data types can be transformed from one type to another by way of an SQL CAST function, if you want. Data can even be joined between different data sources, such as IMS and VSAM, by joining an IMS virtual table or view with a VSAM virtual table with SQL. The beauty of this solution is that almost everyone is familiar with SQL and can easily write a SQL statement that selects the columns or IMS fields that need to be loaded into DB2 Analytics Accelerator for z/OS.

Build a job to load the result set of the SQL statement into DB2 Analytics Accelerator for z/OS

From within DB2 Analytics Accelerator Studio, you generate the job (job control language) that performs the entire process from end to end.

Execute the batch job

When executed, the batch job does the following steps:

1. Creates a DB2 table that matches the result set of the SQL.
2. Adds the newly created table into DB2 Analytics Accelerator for z/OS.
3. Extracts the data from the IMS databases that are referenced in the virtual tables or views in the SQL statement.
4. Converts the IMS data (in memory) to the required DB2 and DB2 Analytics Accelerator for z/OS format.
5. Loads the converted data into DB2 Analytics Accelerator for z/OS.

This final step is simple. You can place this job in a scheduler to refresh the data in DB2 Analytics Accelerator for z/OS at any time-interval you want. Also, you only need to repeat the first three steps if the IMS database changes.

Advantages of the DB2 Analytics Accelerator Loader for z/OS V2.1

Although extracting and transforming IMS data by using other tools is possible, using DB2 Analytics Accelerator Loader for z/OS V2.1 offers the following benefits:

- ▶ One tool is used for the entire process of extracting, transforming, and loading into the DB2 Analytics Accelerator.
- ▶ The data is extracted from the source, converted to the necessary format, and loaded into DB2 Analytics Accelerator for z/OS without landing or writing the data to any intermediate or staging files.
- ▶ The conversion of the data to the necessary DB2 and DB2 Analytics Accelerator for z/OS format is performed in-memory and takes advantage of zIIP processors to significantly lower costs.
- ▶ DB2 Analytics Accelerator Loader for z/OS V2.1 operates where the majority of your data resides, that is on z/OS. This process avoids potential security exposures of the traditional method of copying and propagating data off of the secure z/OS platform.
- ▶ DB2 Analytics Accelerator Loader for z/OS V2.1 provides a fast and low-cost process to load IMS and many other types of data into DB2 Analytics Accelerator for z/OS. This approach provides the ability to consolidate your entire enterprise data quickly and efficiently in DB2 Analytics Accelerator for z/OS to gain better insight into your business.
- ▶ Because the DB2 Analytics Accelerator Loader for z/OS V2.1 user interface is simply a Data Studio plug-in, customers are familiar with this and will find using it to be easy.

Usage scenarios

With IMS operational data made available to the DB2 Analytics Accelerator for z/OS, comprehensive analytics can be delivered to decision-makers with industry-leading security and availability, taking the risk out of operational business analytics. These use cases illustrate how the addition of IMS data to the DB2 Analytics Accelerator for z/OS moves organizations even closer to real-time insight.

Case in point: Responsiveness to clients

One measure of success for any organization is how responsive it is to its clients, whether they are business partners, external clients, or internal organizations. Client requests for customized reports, sourced with data from IMS, DB2 for z/OS, or both, occur frequently. If the necessary information for the report cannot be efficiently gathered, the request cannot be satisfied in a timely manner, which can introduce risk. Each business decision might require a unique set of data. The need to develop, test, and deploy applications in pursuit of the required data from multiple subsystems is costly and time-consuming. Time-to-market is negatively impacted, because the develop-test-deploy cycle can take weeks, even months. In addition, the efficiency of the reporting itself (running the application) is diminished because of the amount of data being processed. In some cases, resource limits are reached and the queries do not complete successfully. Incorporating IMS operational data into analytics powered by DB2 Analytics Accelerator for z/OS creates a single high speed solution.

Case in point: Personalized customer care

Being responsive is one cornerstone of customer retention. Being able to personalize communication and offer customized services, another important cornerstone, helps clients feel like individuals. Analytics plays a key role in identifying customized services and new opportunities. Data warehouses and repositories depend on data from operational systems, such as IMS, which have the transactional level detail but the need to keep data volumes down means this level of detail is frequently sacrificed. This situation limits the analytics to only the summarized data points instead of the full wealth of transactional data that does, in fact, exist in the operational systems themselves. Performing analytics against low-latency warehouses that include IMS transactional detail can provide insights that enable you to enrich your customer retention offerings. In addition, clients who choose to archive the deep transactional detail often maintain that data in managed files, which makes searching and filtering it costly, complex, and time-consuming. A low-latency warehouse powered by DB2 Analytics Accelerator for z/OS serves both IMS and DB2 for z/OS data and delivers efficient storage and retrieval of sensitive, historic, and valuable IBM z Systems data.

Case in point: Data integrity

IBM z System clients frequently have both IMS and DB2 data assets. A common scenario in this situation is that a correlation exists between the two versions of the data in the two subsystems to ensure data integrity. This correlation, a form of referential integrity, exists between the IMS and DB2 database; however, it is not managed at the database management system (DBMS) layer but at the application layer. In this situation, additional applications (usually batch processes) need to be run at periodic intervals to ensure that the data is synchronized between the two subsystems. This process becomes obsolete with the introduction of an efficient z Systems warehouse powered by DB2 Analytics Accelerator for z/OS containing both IMS and DB2 data. Such a warehouse also maintains the security and availability clients have come to expect from the mainframe.

Supported platforms

IBM DB2 Analytics Accelerator for z/OS runs with DB2 for z/OS (currently DB2 10 and above) environments and complies with their prerequisites in terms of z System resources. The DB2 Analytics Accelerator Loader for z/OS Version 2.1 is supported. IMS versions 12, 13, and 14 are the supported versions. For more details, review the product announcement letter:

<http://www.ibm.com/common/ssi/cgi-bin/ssialias?infotype=an&subtype=ca&appname=gpat&eam&supplier=897&letternum=ENUS215-385>

Ordering information

You can request delivery for DB2 Analytics Accelerator for z/OS, V5.1 (product number 5697-DA5) through Custom-Built Product Delivery Option (CBPDO) and ServerPac.

These customized offerings are for Internet delivery in countries where ShopzSeries product ordering is available. Internet delivery reduces software delivery time and allows you to install software without the need to handle tapes. For details about Internet delivery, see the help information at the Shopz website:

<http://www.software.ibm.com/ShopzSeries>

Related information

For more information, see the publications and web pages listed here.

IBM Redbooks publications:

- ▶ *Reliability and Performance with IBM DB2 Analytics Accelerator Version 4.1*, SG24-8213
<http://www.redbooks.ibm.com/abstracts/sg248213.html>
- ▶ *Hybrid Analytics Solution using IBM DB2 Analytics Accelerator for z/OS V3.1*, SG24-8151
<http://www.redbooks.ibm.com/abstracts/sg248151.html>

Other information:

- ▶ *Accelerate business insights with IMS Transactional data*
http://ibm.biz/accelerate_insights_ims_transactional_data
- ▶ DB2 Analytics Accelerator for z/OS product page:
<http://www.ibm.com/software/products/en/db2analacceforzos>
- ▶ IBM Offering Information page (announcement letters and sales manuals)
http://www.ibm.com/common/ssi/index.wss?request_locale=en

On this page, enter DB2 Analytics Accelerator for z/OS, select the information type, and then click **Search**. On the next page, narrow your search results by geography and language.

Authors

This solution guide was authored by **Deepak Kohli**, a Senior Software Engineer with the IMS Lab at the IBM Silicon Valley Laboratories, and working with the International Technical Support Organization, Poughkeepsie Center.

Thanks to the following people for their contributions to this project:

LindaMay Patterson
International Technical Support Organization, Rochester Center

Kyle Charlet
Distinguished Engineer IMS Architect: Cloud, Mobile, Analytics

Tim Willging
Distinguished Engineer, Database Tools Architect and Strategist

Now you can become a published author, too!

Here's 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 will 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 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 US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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, MD-NC119, Armonk, NY 10504-1785, US

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 jurisdictions 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 provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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 actual people or business enterprises 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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

| | | |
|--|----------------|---|
| DB2® | IBM PureData® | Redbooks (logo)  ® |
| Distributed Relational Database Architecture™ | IBM z Systems™ | z Systems™ |
| DRDA® | IMS™ | z/OS® |
| IBM® | PureData® | |
| | Redbooks® | |

The following terms are trademarks of other companies:

Netezza, and N logo are trademarks or registered trademarks of IBM International Group B.V., an IBM Company.

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



REDP-5313-00

ISBN 0738455229

Printed in U.S.A.

Get connected

