IBM

Getting Started with the IBM Smart Analytics System 9600

- Getting the most from the IBM Smart Analytics System 9600
- Understanding System z and the IBM Smart Analytics System 9600
- Managing the components

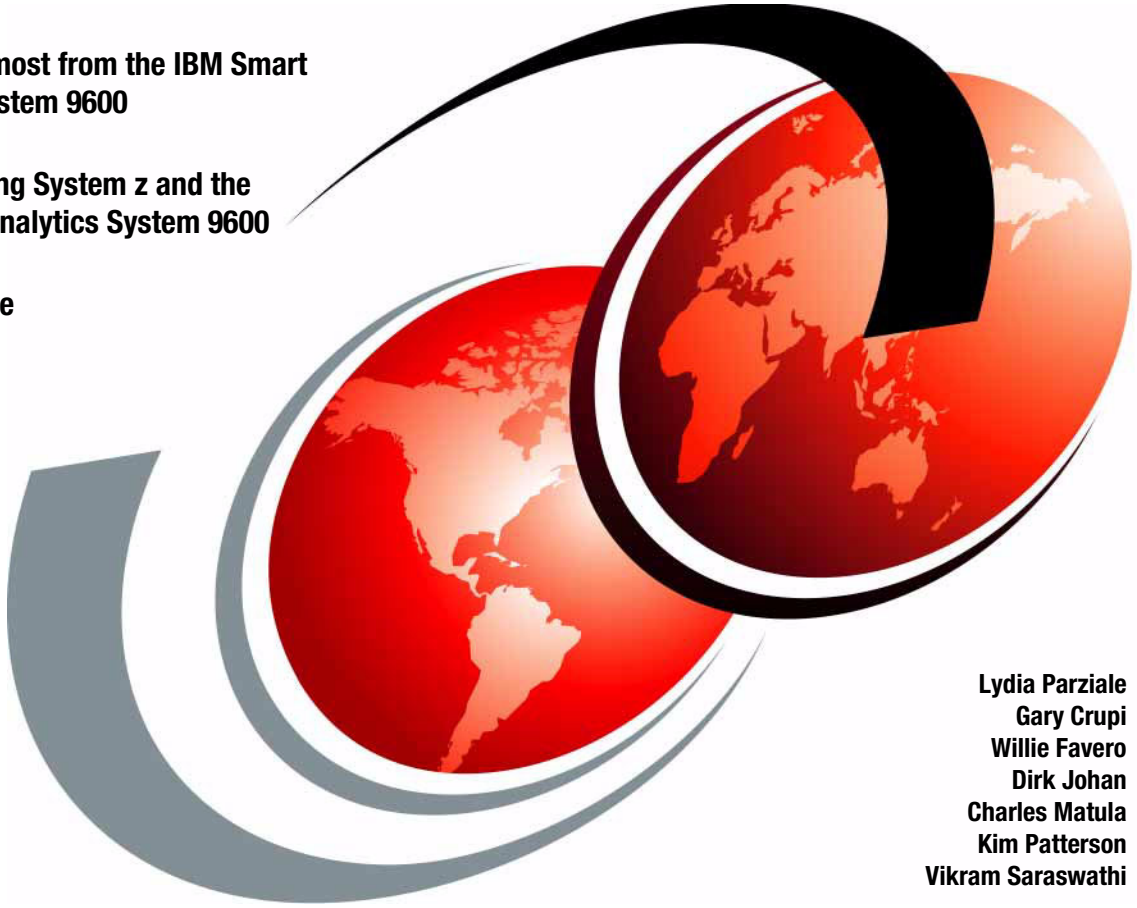

Lydia Parziale
Gary Crupi
Willie Favero
Dirk Johan
Charles Matula
Kim Patterson
Vikram Saraswathi

Redbooks

ibm.com/redbooks

International Technical Support Organization

Getting Started with the IBM Smart Analytics System 9600

April 2011

SG24-7902-00

Note: Before using this information and the product it supports, read the information in “Notices” on page vii.

First Edition (April 2011)

This edition applies to DB2 9 for z/OS Value Unit Edition (VUE) or MIPS Based License Charging (MLC), InfoSphere Warehouse for Linux on System z Version 9.5.2, Cognos Business Intelligence for Linux on System z Version 8.4, z/OS Version 1.11 Operating System Stack, z/VM Version 6.2, and Novell SUSE Linux Enterprise Server 10, SP2, WebSphere Application Sever Version 7 Fix Pack 9.

Contents

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 about 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.

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.

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:

Blox®
CICS®
Cognos®
DB2 Connect™
DB2®
DirMaint™
DRDA®
GDPS®
Global Business Services®
IBM®
IMS™
InfoSphere™
MVS™
OMEGAMON®
Parallel Sysplex®
pureXML®
QMF™
RACF®
Redbooks®
Redbooks (logo) ®
Resource Measurement Facility™
RMF™
System Storage®
System z10®
System z®
Tivoli®
VTAM®
WebSphere®
z/OS®
z/VM®
z10™
z9®

The following terms are trademarks of other companies:

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

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

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

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

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

Preface

The IBM® Smart Analytics System 9600 is a single, end-to-end business analytics solution to accelerate data warehousing and business intelligence initiatives. It provides integrated hardware, software, and services that enable enterprise customers to quickly and cost-effectively deploy business-changing analytics across their organizations.

As a workload-optimized system for business analytics, it leverages the strengths of the System z® platform to drive:

- Significant savings in hardware, software, operating, and people costs to deliver a complete range of Data Warehouse and Business Intelligence (BI) capabilities
- Faster time to value with a reduction in the time and speed associated with deploying the foundation for Business Intelligence applications
- Industry-leading scalability, reliability, availability, and security
- Simplified and faster access to the data on System z

Using the IBM Smart Analytics System 9600 helps ensure that a solution is quickly up and running and remains as relevant and powerful in the future as it is today. At the core of the IBM Smart Analytics System is DB2® for z/OS® and the powerful warehouse capabilities from IBM InfoSphere™ Warehouse. This foundation not only manages the data store, but also is essential for speeding system deployment and enabling advanced analytics. The analytic information is then made available to the users where and when it is needed using the breadth of reporting, analysis, and dashboarding capabilities available with IBM Cognos® 8 Business Intelligence.

Each configuration can be augmented at any time to meet new requirements by adding new analytic capability or data and user capacity building block components. Because all of these components use the same foundation, the system is easy to maintain, preserves existing investments, and delivers results quickly.

This flexibility and scalability enable customers to select the best combination to meet their requirements today and retain that investment for future growth.

The IBM Smart Analytics System 9600 takes existing IBM hardware, maintenance, and software and packages that with IBM LAB Services to create a fast and easy-to-deploy, end-to-end business intelligence environment. The IBM Smart Analytics System 9600 is shipped directly to the customer floor, where

IBM Lab services will come on site to install and prepare the system for turnover to the customer, ready for them to define and load their database.

This reduces the time necessary to install the system and software from months to weeks. The set of IBM products in this offering has been tested together, removing many of the risks associated with integrating many different pieces of a solution, ensuring that the customer will have a functional working system.

This IBM Redbooks® publication will assist customers in getting started with their IBM Smart Analytics System 9600. In addition to identifying first tasks, this book provides overviews of key concepts and an introduction to systems management information.

This book is intended for system administrators, data warehouse administrators, database administrators, and other technical personnel who will be managing the IBM Smart Analytics System.

The team who wrote this book

This book was produced by a team of specialists from around the world working at the International Technical Support Organization, Poughkeepsie Center.

Lydia Parziale is a Project Leader for the ITSO-GCS team in Poughkeepsie, New York, with domestic and international experience in technology management, including software development, project management, and strategic planning. Her areas of expertise include e-business development and database management technologies. Lydia is a certified PMP and an IBM Certified IT Specialist with an MBA in Technology Management and has been employed by IBM for over 24 years in various technology areas.

Gary Crupi is a Senior Certified Executive IT Specialist who joined IBM in 2001. Building on his prior experience as a Senior Systems Analyst with Northwestern Mutual and the United States Air Force, he helps customers leverage IBM Information Management Solutions across platforms. In addition, he encourages System z platform customers to maximize their investments through software currency and modernization. Throughout his career, Gary has helped customers and IBMers position for success by selecting the right platform for DB2 based on requirements. Most recently, Gary defined and led the creation of the IBM Smart Analytics System 9600. This solution is the backbone of the IBM Business Analytics on System z initiative. Today, Gary continues leading multiple organizations in his hybrid role of lead Technical Architect for the 9600 and Senior Technical Sales Leader for System z Data Warehouse and Business Intelligence solutions.

Willie Favero is an IBM Senior Certified IT Software Specialist and DB2 SME for the IBM Silicon Valley Lab Data Warehouse on System z Swat Team. He has over 35 years of experience working with databases with more than 25 years of that working with DB2. He is a sought-after international speaker for conferences, user groups, and seminars, and he publishes articles, white papers, and IBM Redbooks publications, and has one of the most-read technical blogs on the Internet.

Dirk Johan is an IT Architect at the IBM Boeblingen Lab in the Center of Excellence for Data Warehouse on System z. His team conducts POCs for large and complex Data Warehousing implementations and supports customers in all areas of Data Warehouse topics on System z. With more than 20 years of hands-on experience in the mainframe world, Dirk has been working on systems operations, databases, and application programming. He has presented at several IDUG, GSE, and IOD conferences.

Charles Matula is an IT Architect/Specialist on the IBM Global Account in Poughkeepsie, New York. He has been in the Cognos COC for the past two years, developing data-mart reporting solutions. He has over 20 years of experience developing database application and warehouse solutions, using a variety of DBMS, including DB2 LUW, DB2 z/OS, Sybase, Oracle, and others. Charles received his BS in Electrical/Computer Engineering at the State University of New York. Charles presented his design “The Account Data Model” at the Conceptual Modeling - ER 2002, 21st International Conference on Conceptual Modeling, Tampere, Finland, which is a Star-Schema hybrid.

Kim Patterson is a Managing Consultant in the United States. She works with DB2 for System z. She holds a master’s degree in information systems from Rutgers University. Her areas of expertise include DB2 installation, configuration, performance monitoring, and SQL tuning. She works with ISV systems, which run on DB2, including SAP and Siebel. She has taught as an instructor for IBM and has co-authored an IBM Redbooks publication on APPC protected conversations and WebSphere® for z/OS to CICS® and IMS™ Connectivity Performance. She has also presented at Data Management conferences and zNTP.

Vikram Saraswathi is a an IT Specialist with IBM Global Business Services®. He is a certified DB2 Database Administrator for z/OS, a solution designer for Business Intelligence Solutions using DB2. He lives in Bangalore, India, and has four years of experience working with Mainframes and DB2 as a DBA, Systems Programmer, Data Modeler, WebSphere MQ Solution Designer and System Administrator. He holds a bachelor’s degree in electrical and electronics engineering from Jawaharlal Nehru Technological University, India.

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

Roy P Costa, Bob Haimowitz
International Technical Support Organization, Poughkeepsie Center

Mei Hing (Ann) Jackson
IBM USA

Andrew Perkins, IT Architect
IBM USA

Jonathan Sloan, IT Architect
IBM USA

Dino Tonelli, Software Performance Analyst: System z
IBM USA

Mark Nover, Systems Programmer: MVS™
IBM USA

James Jasper, Management Consultant
IBM USA

Thanks to the authors of the following IBM Redbooks publications:

- Chuck Ballard, Nicole Harris, Andrew Lawrence, Meridee Lowry, Andy Perkins, and Sundari Voruganti, authors of *InfoSphere Warehouse: A Robust Infrastructure for Business Intelligence*, SG24-7813
- Mike Ebbers, Dino Tonelli, Jason Arnold, Patric Becker, Yuan-chi Chang, Willie Favero, Shantan Kethireddy, Nin Lei, Shirley Lin, Ron Lounsbury, Susan Widing Lynch, Cristian Molaro, Deepak Rangarao, and Michael Schapira, authors of *Co-locating Transactional and Data Warehouse Workloads on System z*, SG24-7726
- Paolo Bruni, Gaurav Bhagat, Lothar Goeggelmann, Sreenivasa Janaki, Andrew Keenan, Cristian Molaro, and Frank Neumann, authors of *Enterprise Data Warehousing with DB2 9 for z/OS*, SG24-7637
- Bertrand Dufrasne, Werner Bauer, Brenda Careaga, Jukka Myyrrylainen, Antonio Rainero, and Paulus Usong, authors of *IBM System Storage DS8700 Architecture and Implementation*, SG24-8786, and *IBM System Storage DS8700 Easy Tier*, REDP-4667

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

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications 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

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

1

Overview of the IBM Smart Analytics System 9600

This chapter provides an overview of the IBM Smart Analytics System 9600. We describe the architecture and the special requirements and conditions of the IBM Smart Analytics System 9600. In the architectural overview you will see how the components of the hardware and software solution fit together. In the latter parts of this chapter we describe the content and functionality of the IBM Smart Analytics System 9600 solution.

1.1 Architectural overview

The IBM Smart Analytics System 9600 is part of the Smart Analytics System family. The IBM Smart Analytics System 9600 is based on a transparent, modular architecture that allows you to choose the way that your data warehouse solution develops. You start with a base configuration and add capacity in granular, balanced increments as required.

The IBM Smart Analytics System 9600 is a System z-based solution. Its implementation is based on System z hardware and contains an integrated stack of software, operating system, and hardware that solves the needs of an Enterprise Data Warehouse (EDW) and Business Intelligence (BI) environment. It starts with a balanced stack, and then grows incrementally, in a balanced fashion, as your data and query volumes dictate. The bundling and pricing is built on the Solution Edition (SE) foundation:

- The IBM System z Solution Edition for Data Warehousing creates a single LPAR that supports the data warehouse data store. This LPAR executes all queries that are submitted to the DB2 for z/OS database within the LPAR.
- The Solution Edition for Enterprise Linux® provides the tooling to deliver a collocated environment for business intelligence workload.

The IBM Smart Analytics System 9600 creates an end-to-end environment for Business Intelligence (BI) that includes the DB2 for z/OS LPAR with Linux LPARs for BI tools, such as InfoSphere Warehouse and Cognos 8.4 BI.

However, the IBM Smart Analytics System 9600 is much more than a bundle of discounted products.

The IBM Smart Analytics System 9600 includes everything required to serve as a foundation for your analytics and business intelligence solutions. It delivers a system of software, server and storage hardware, and services to eliminate the time and cost of integrating and optimizing analytics solutions for business use, while preserving the flexibility not offered by single use appliances. Different from the other members of the Smart Analytics System family, the IBM Smart Analytics System 9600 provides the highest qualities of service and can be delivered in two ways:

- As a stand-alone system.
- As an upgrade to an existing environment which then manifests itself as a virtual appliance. This option is unique to the IBM Smart Analytics System 9600 and related to the flexible resource allocation and configuration abilities of System z.

Regardless of which deployment option is selected, the IBM Smart Analytics System 9600 contains the same software packages and hardware resources. The difference is only the physical implementation. For the stand-alone system, IBM Smart Analytics System 9600 runs on a separate System z, while the virtual upgrade shares an existing System z. Figure 1-1 shows the general layers of the configuration.

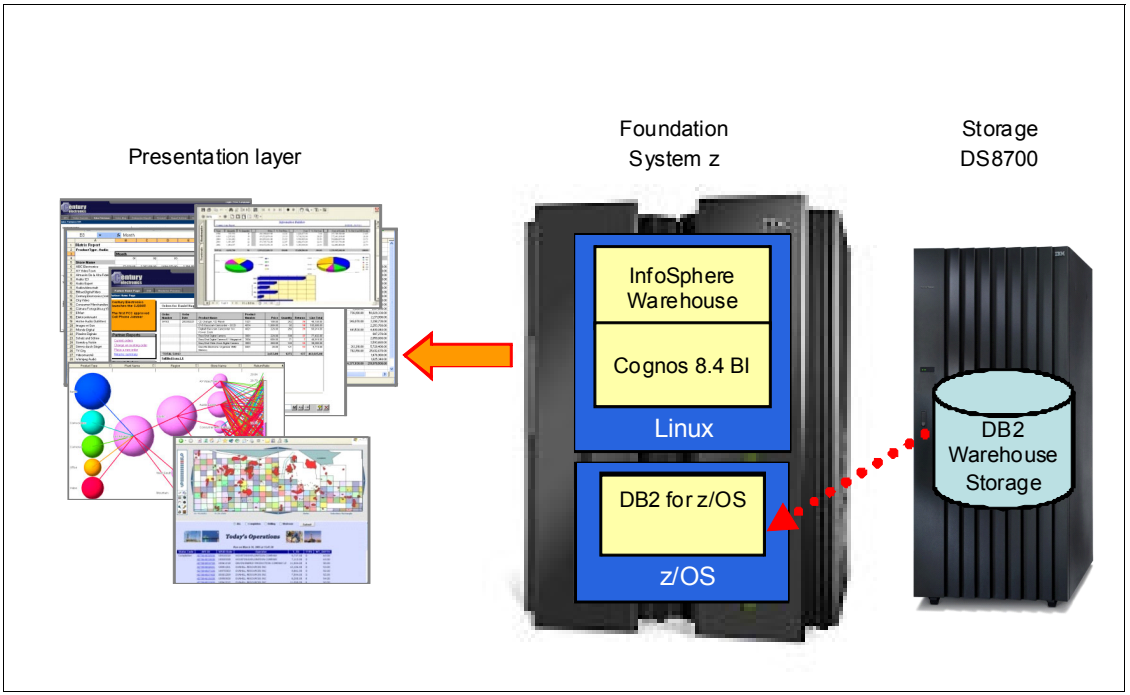

Figure 1-1 General layers of the IBM Smart Analytics System 9600 solution

At the presentation layer:

- Data can be accessed from anywhere in the enterprise.
- IBM Cognos 8 BI runs on an application server.

At the foundation layer is:

- A Linux LPAR for enterprise business intelligence with InfoSphere Warehouse and Cognos 8.4 Business Intelligence for reporting and business analytics
- A z/OS operating system
- A DB2 for z/OS for the enterprise data warehouse

At the warehouse storage layer is System Storage®, a high-performance, high-capacity storage subsystem. It offers balanced performance and storage capacity that scales linearly up to hundreds of terabytes.

1.2 Hardware specification

The IBM Smart Analytics System 9600 uses either the IBM System z10® or System z196 hardware in either a logical partition (LPAR) or a stand-alone system.

The System z196 hardware includes:

- Up to 80 processing units (some can be reserved for ICFs) consisting of general-purpose processors (CP) and specialty engines (zIIPs and IFLs)
- 5.2 GHz cores
- Approximately 12 -14 GB memory per CP/zIIP
- 16 GB memory per IFL
- Network connections
- Concurrent Hardware Management Console (HMC) and Support Element

The System z10 hardware includes:

- Up to 64 processing units (some may be reserved for ICFs etc) consisting of general-purpose processors (CP) and specialty engines (zIIPs and IFLs)
- 4.4 GHz cores
- 8 GB memory per GP/zIIP
- 16 GB memory per IFL
- Network connections
- Concurrent Hardware Management Console (HMC) and Support Element

The Smart Analytics System 9600 can be integrated into an existing environment via an additional member of a DB2 data sharing group. Using the same environment for both operational data and the data warehouse optimizes the extract, transform, and load (ETL) processes as well as simplifying applications that leverage both operational and data warehouse content. IBM Services can assist in integrating the new system into your environment.

The hardware components feature a call-home capability in the z10 and z196 servers, as well as the DS8700 storage subsystem, should there be any hardware issues. Additionally, owners of an IBM Smart Analytics System 9600 will have access to support by calling 1-800-IBMSERV.

Storage overview

The DS8700 Storage Subsystem is included in the IBM Smart Analytics System 9600 configuration.

The DS8700 Enterprise Class Storage provides:

- Ten to 278 TB usable storage in pre-configured solutions
- RAID 5
- Leveraging:
  - 300 GB or 450 GB DDMs running at 15 K RPM
  - 64 GB to 512 GB cache
  - Two to 76 drive sets (16 DDMs each)
  - Two to 48 host adapters
  - Four to 96 host channels to DS8700
  - Two to 32 disk adapters
  - HyperPAV
  - MIDAW
  - zHPF

As part of the IBM Smart Analytics System 9600, storage will come configured for your environment according to best practices guidelines.

The following is the configuration for all of the pre-configured sizes (customized sizes are available) that are offered by this solution:

4 TB	Two LCUs, each with 114 3390B (real 3390s) and forty-eight 3390A (aliases) (total of 162 UCBs/LCU). The first 16 addresses on each LCU are 3390-9 and the rest are 3390-27.
12 TB	Eight LCUs, each with 147 3390B (real 3390s) and sixty-four 3390A (aliases) (total of 211 UCBs/LCU).The first eight addresses on each LCU are 3390-9 and the rest are 3390-27.
25 TB	Sixteen LCUs, each with 110 3390B (real 3390s) and sixty-four 3390A (aliases) (total of 174 UCBs/LCU). The first four addresses on each LCU are 3390-9 and the rest are 3390-54.
50 TB	Sixteen LCUs, each with 163 3390B (real 3390s) and sixty-four 3390A (aliases) (total of 227 UCBs/LCU). The first four addresses on each LCU are 3390-9 and the rest are 3390-54.
100 TB	Eighteen LCUs, each with 192 3390B (real 3390s) and sixty-four 3390A (aliases) (total of 256 UCBs/LCU).

For the first 16 LCUs, the first four addresses are 3390-9 and the rest are 3390-54.

For the next 2 LCUs, all 192 addresses are 3390-54. Twelve additional LCUs, each with 143 3390B (real 3390s) and sixty-four 3390A (aliases) (total of 207 UCBs/LCU). All 143 addresses are 3390-54.

The DS8700 configuration supports two LPARs:

- LPAR 1: z/OS, DB2 for z/OS VUE/MLC (optional alternative)
- LPAR 2: Multiple Linux on System z guests running InfoSphere Warehouse, and Cognos 8.4 BI with 5 - 10,000 Users (optional de-selection)

1.3 Software overview

IBM has assembled a base set of software that is pre-optimized for out-of-the-box performance to enable you to build a comprehensive data warehouse. At the time of this writing, the IBM Smart Analytics System 9600 uses the following software:

- DB2 for z/OS Value Unit Edition (primary) v9 with an option for Monthly Licence Charge (MLC)
- DB2 Utilities Suite v9
- InfoSphere Warehouse for Linux on System z v9.5.2
- IBM Cognos 8.4 BI for Linux on System z
  - IBM Cognos 8 BI reporting
  - IBM Cognos 8 BI analysis
  - IBM Cognos 8 BI dashboard
- z/OS operating system stack v1.11
- z/VM® v6.1 with a Linux guest pre-installed and configured (It should be noted here that the customer must provide a supported Linux license. SUSE 10 SP2 was validated in this IBM Redbooks publication.)

For the current validated stack of software for the IBM Smart Analytics System 9600, see:

http://www-01.ibm.com/support/docview.wss?uid=swg21450964

DB2 for z/OS provides the software backbone of the solution. Advanced query prioritization capabilities allow identification of critical, specific user queries within a large query workload, and have them executed without delay. With Workload

Manager (WLM) provided by z/OS and DB2 for z/OS, you can prioritize individual users to ensure that the application of processing resources is based on business requirements.

The system, as deployed, has two LPARs included. One contains DB2 for z/OS in a native z/OS LPAR and the other one has z/VM installed. Multiple z/VM Linux on System z guests are configured to support InfoSphere Warehouse on System z and Cognos BI.

The InfoSphere Warehouse Cubing Services capability, the WebSphere Application Server, and the Administration Console run on Linux on System z. The target DB2 warehouse database runs on the z/OS operating system. The source data can come from any mainframe or distributed system.

Because z/VM is part of the solution, it is easy to build up a customized system. For example, you can clone the existing Linux partition as many times as you like to create a test environment for each of your developers.

1.4 Network specifications

The IBM Smart Analytics System 9600 uses the following networks:

- z/OS network: Uses the z/OS Communications Server (TCP/IP), which is connected to the z/VM LPAR via a Hipersocket.
- z/VM: TCP/IP has been configured according to specifications provided by the customer. For example, in our case, the z/VM user ID of the z/VM TCP/IP stack virtual machine is TCPIP. The hostname, domain name, domain IP address, device number, and IP address have already been pre-configured according to the installation specifications. Path MTU discovery is enabled and QDIO (layer 3) has been selected. The network type will be Ethernet and the maximum transmission unit (MTU) size is set to 1500.
- Linux on System z network: TCP/IP connectivity for Linux guests. Virtual network interfaces allow the real connections to be shared. The virtual network connection used here is via VSWITCH. In addition to providing a network of virtual adapters, the switch is connected directly to an OSA-Express QDIO adapter.

Figure 1-2 shows an overview of the IBM Smart Analytics System 9600 Network.

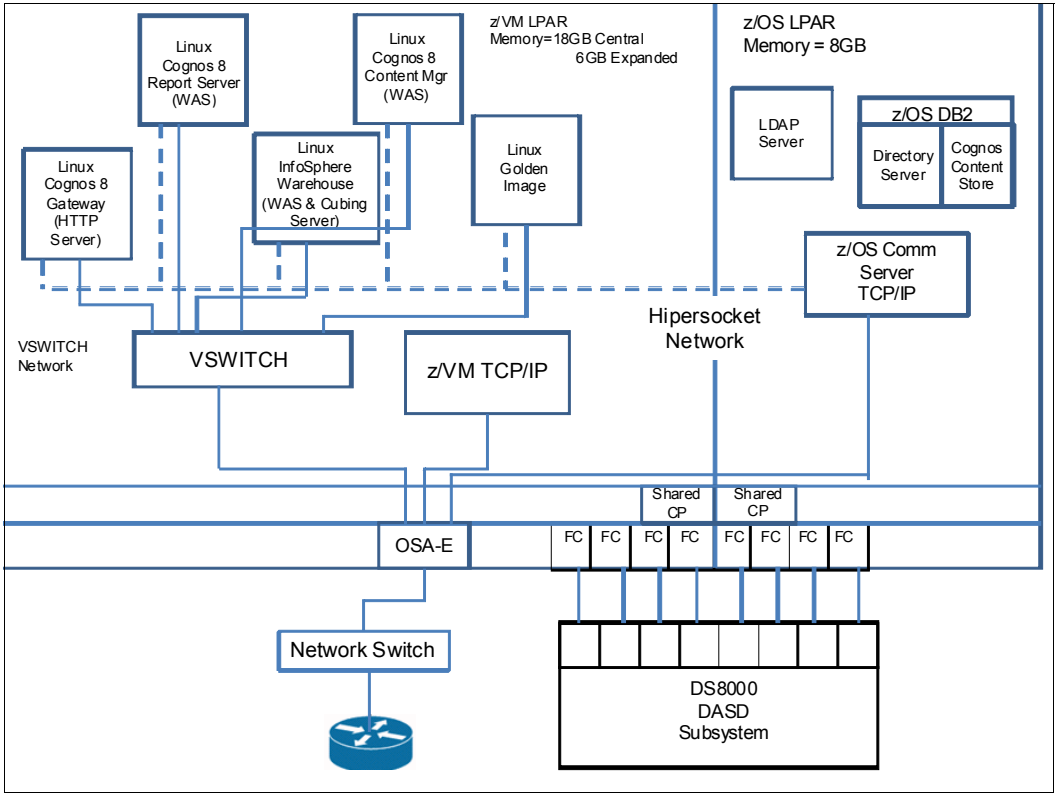

Figure 1-2 Network overview

1.5 Optional software components overview

To further enhance the IBM Smart analytics System 9600, optional compatible components are available, such as:

- InfoSphere Master Data Management Server
- InfoSphere Information Server
- InfoSphere Replication Server (Q-Rep, CDC and Event Publisher eligible)
- InfoSphere Federation Server plus Classic Federation on System z
- SPSS
- Cognos Now! for Linux on System z

- IBM Smart Analytics Optimizer for DB2 for z/OS V1.1

  This intended to speed up Data Warehouse and Business Intelligence workloads. For more information about this, see *Co-locating Transactional and Data Warehouse Workloads on System z*, SG24-7726.
- Tivoli® OMEGAMON® for DB2 Performance Expert
- Tivoli Directory Server

2

Getting started

This chapter gives an overview of the next steps to get started using the IBM Smart Analytics System 9600.

First, the procedures necessary to get started are outlined, then the roles that will be involved in using the IBM Smart Analytics System 9600 will be discussed. Here we provide some guidance on what needs to be done in order to set up a data warehousing environment using the InfoSphere Warehouse for System z before a business intelligence specialist can start to create their first report.

This chapter discusses:

- Procedure overview
- Roles
- Enterprise Data Warehouse
- InfoSphere Warehouse for System z overview
- Preparing Cognos BI to create reports

2.1 Procedure overview

The following list is an overview of things to consider before getting started. Each step is discussed in a section that will give you more details on the step.

1. Identify roles (2.2, “Identifying the roles” on page 12).
2. Define users and authorities (Chapter 8, “Managing users of the IBM Smart Analytics System 9600” on page 99).
3. Implement startup procedures (4.1, “Startup procedure for IBM Smart Analytics System 9600 components” on page 52).
4. Implement WLM policies for a business intelligence environment (see 7.1, “IBM Smart Analytics System 9600 WLM Policies” on page 79).

2.2 Identifying the roles

The following roles will need to be involved in various capacities:

- Systems programmer

  The Systems Programmer is responsible for all tasks related to the operating system (z/OS or Linux on System z). Setting up the basic environment, which includes workload management (WLM) and security (RACF®, LDAP), is also included in this role.

- Database administrator (DBA)

  The database administrator is responsible for all tasks related to the database system (DB2 for z/OS). This includes tasks for maintaining the database (utilities, backup, and recovery), as well as managing ongoing performance. The physical implementation of the database objects is also managed by the DBA who will work closely with the warehouse administrator.

- Warehouse administrator

  The warehouse administrator performs tasks such as creating the tables and extract, transform, and load (ETL) or data movement processes or flows to populate the data structures. This person uses the SQW, SQW run time, Admin Console, and cubing services to perform their work.

- Data modeler

  The data modeler provides the definition and format of the data. The person in this role has to do an analysis to first understand all the data in the organization and then decide what should be sent to a warehouse.

- Warehouse/Business Intelligence (BI) architect

  The warehouse/BI Architect models the overall BI system for the BI users. The functional architecture and designs of the main data flows for the reports are managed by the person in this role.

- Business Intelligence specialist

  The BI specialist is also referred to as the BI developer or OLAP developer. All specific requests for the BI system are handled by this role.

Generally, the systems programmer and DBA will define users and authorities according to the security guidelines of their enterprise. The systems programmer would also define the WLM policies for the BI environment according to company security guidelines. Implementing startup procedures would be performed by the systems programmer, the warehouse administrator, and the BI architect.

2.3 InfoSphere Warehouse for System z

InfoSphere Warehouse on System z can be used to build your data warehouse leveraging existing data sources. The data warehouse, or subsequent data marts, are then used to perform multidimensional analysis and reporting of data. Cubing Services can also be implemented to provide exceptional performance. You can also use the in-database data movement and manipulation capabilities of the SQL Warehouse Tool (SQW) to transform and load your data. Your InfoSphere Warehouse Server product is on a Linux on System z partition connecting to your remote DB2 for z/OS database server. The steps to consider to complete the InfoSphere Warehouse setup are:

1. Implement connections to the EDW and the online transaction processing OLTP sources. The BI specialist would work with the DB2 DBA on this.
2. Implement startup procedures for InfoSphere Warehouse (systems programmer).
3. Model data sources in InfoSphere Warehouse (data modeler, warehouse/BI architect).
4. Define stages for extract, load, and transform (ELT) processes in InfoSphere Warehouse (data modeler, warehouse/ BI architect).
5. Define Cube model in InfoSphere Warehouse.
6. Implement Cube model in InfoSphere Warehouse.
7. Make Cube model available for Cognos.

See Chapter 5, “InfoSphere Warehouse administrative tasks” on page 55, for more details on each task.

For more information about InfoSphere Warehouse for System z tasks, see Chapter 5, “InfoSphere Warehouse administrative tasks” on page 55.

2.4 The Enterprise Data Warehouse

As outlined in 3.1, “Database design” on page 16, the IBM Smart Analytics System 9600 comes with two databases:

- SQWCTRL
- DWESAMP

From there, the data modeler and the Warehouse/BI Architect design the enterprise data warehouse (EDW) structures. This includes the DB2 DBA and systems programmer defining high-availability procedures for DB2.

The DB2 DBA is then responsible for the physical implementation of the EDW. For more information about setting up an EDW, see *Enterprise Data Warehousing with DB2 9 for z/OS*, SG24-7637, which can be found at:

`http://www.redbooks.ibm.com/abstracts/sg247637.html?Open`

Additionally, the DB2 DBA will define backup and recovery strategies for the EDW, as well as maintenance procedures.

2.5 Preparing Cognos BI to create reports

Cognos 8 BI is installed with the IBM Smart Analytics System 9600 and will be on the same LPAR as InfoSphere Warehouse. It also runs in a Linux on System z guest. Cognos BI connects to DB2 on z/OS to access and retrieve the data. This connection is primarily through Java™ Database Connectivity (JDBC). The system installed at your installation will have Java available as part of the system pack.

Details of the following steps can be found at Chapter 6, “Cognos 8 Business Intelligence” on page 69:

1. Define the scope for BI reports (DWH/BI architect).
2. Make EDW structures available for Cognos Reporting (DWH/BI specialist).
3. Make the Cube model available for Cognos (BI specialist). (Also see Chapter 5, “InfoSphere Warehouse administrative tasks” on page 55, for more details.)
4. Create Cognos reports (BI specialist).

3

DB2 design for the Enterprise Data Warehouse

This chapter provides an overview of the design implications and configuration changes for DB2 for z/OS in support of data warehousing on the IBM Smart Analytics System 9600. The IBM Smart Analytics System 9600 makes DB2 for z/OS a powerful platform for the InfoSphere data warehouse and infrastructure for business analytics.

In this chapter, we discuss:

- Database design
- DB2 for z/OS settings and configuration
- DB2 for z/OS special features for data warehousing
- Database and enterprise data warehouse (EDW) design considerations
- XML and the data warehouse
- DB2 tuning and optimization considerations

3.1 Database design

Two DB2 for z/OS databases were created in the DB2 subsystem using DB2I SPUFI:

- SQWCTRL
- DWESAMP

These databases have been created using the ISWZADM user ID.

The SQWCTRL database is the runtime metadata database for InfoSphere Warehouse. When a user deploys SQW applications or Cubes to the Admin Console, all of the metadata is inserted into SQWCTRL. It also has all of the information that InfoSphere Warehouse needs to operate on a daily basis. When you interact with the Admin Console, all of the information is stored in SQWCTRL. SQWCTRL is created the first time that the Admin Console tries to connect. If the tables do not exist, the Admin Console creates them. SQWCTRL is created and gets cataloged during pre-installation activities.

DWESAMP is a sample DB2 database that includes a set of tables that contain data about a fictitious retail company that sells various types of products through a number of different channels and stores. A set of metadata objects that describe the sample data tables is also included. This is the same sample database that is provided for the Cubing Services tutorial, and you must install the sample and set up the data before you can use it to create your first cubes.

DWESAMP is not installed by the InfoSphere Warehouse installation routine. In fact, it does not even ship with the product. However, when the IBM Smart Analytics System 9600 is set up prior to customer turnover, DWESAMP is created and used in the verification process. So, it is available when you purchase the IBM Smart Analytics System 9600. For a tutorial on how to design and deploy a data warehousing solution that expands the capabilities of the DB2 data warehouse for a fictional company that uses this database, go to:

http://publib.boulder.ibm.com/infocenter/db2luw/v9r5/index.jsp?topic=/com.ibm.dwe.samples.doc/tutgenintrodetails.html

On the chance that you do not have the DWESAMP database and want to create the database and tables, run the installation program and install only the Documentation and Samples choice. To do this, run the SetOlapAndMining script:

1. Open a DB2 command-line interface.
2. On the command line, change the directory to `InfoSphereWarehouseHome\samples\data`.
3. Run the appropriate SetOlapAndMining script.

3.1.1 Buffer pool design

The database manager uses buffer pools to cache data in database memory. For every different table space page size specified, there must be at least one buffer pool with that same page size. Table space page sizes can be 4 KB, 8 KB, 16 KB, or 32 KB. From a number of buffer pools standpoint, a data warehouse environment is not all that different from online transaction processing (OLTP). The correct number of buffer pools is whatever number is necessary to satisfy a warehouse's caching requirement, keeping in mind that the total storage used by the combination of all buffer pools must be something less than the amount of virtual storage available to the warehouse DB2. If the environment is supporting a large number of buffer pools, you must decide which buffer pool must have more pages, and ensure that each buffer pool is performing to its maximum efficiency. This task is best performed using some type of monitoring tool. Make sure that you examine all possible buffer pool page sizes when creating your warehouse table spaces. In some cases, better buffer pool utilization and/or buffer pool efficiency can be achieved by using a large 8 K or 16 K buffer pool.

The size of a buffer pool is determined by the type and amount of warehouse data that will utilize that pool. Which table spaces or indexes will use which buffer pool is often determined by the data characteristics. For example, smaller indexes and dimension tables might be placed in separate buffer pools in order to "pin" them in memory. In these cases, the buffer pool would have to be large enough to contain the objects intended to stay in memory. On the other hand, some table spaces might be so large that they would never be contained entirely in a buffer pool. If these are read-only tables, with lots of prefetch activity and minimal random read, less space might work better than more space. Also, in most cases, least recently used (LRU) is the appropriate setting for the page steal algorithm. However, if the buffer pool is large enough to allow the entire table space or indexspace to be pinned in the buffer pool, first-in-first-out (FIFO) could be considered as a performance enhancement.

Table 3-1 lists suggestions for defining different buffer pools with different characteristics.

Table 3-1 Bufferpool suggestions

	Sequential Steal Threshold	Parallel Sequential Threshold	Assisting Parallel Sequential Threshold	Deferred Write Threshold	Vertical Deferred Write Threshold		Page Fix	Page-Stealing Algorithm	AUTOSIZE
	VPSEQT	VPPSEQT	VPXPSEQT	DWQT	VDWQT		PGFIX	PGSTEAL	AUTOSIZE
Catalog - BP0	50	0	0	50	10		?	LRU	NO
Catalog - BP8K0	50	0	0	50	10		?	LRU	NO
Catalog - BP16K0	50	0	0	50	10		?	LRU	NO
Buffer pool without parallelism	80	0	0	50	5,	0	YES	LRU	NO
Buffer pool w/parallelism	80	100	0	50	5,	0	YES	LRU	NO
Buffer pool all compressed indexes	80	100	0	50	5,	0	NO	LRU	NO
4K sort buffer pool	95	50	0	30	5,	0	YES	LRU	NO
32K sort buffer pool Make larger than 4K	95	50	0	30	5,	0	YES	LRU	NO
	0-100%	0-100%	0-100%	0-100%	0-90	0-9999	YES/NO	LRU/FIFO	YES/NO
	Default = 80%	Default = 50%	Default = 0%	Default = 30	0	0	Default = NO	Default = LRU	Default = NO

To alter a buffer pool's characteristics, including its initial size if it already does not exist, use the following command (all on one line):

```
-ALTER BUFFERPOOL (bpname) VPSIZE(integer) VPSEQT(integer)
VPPSEQT(integer) VPXPSEQT(integer) DWQT(integer)
VDWQT(integer1,integer2) PGSTEAL( LRU/FIFO/NONE ) PGFIX( YES/NO )
AUTOSIZE( YES/NO )
```

All keywords listed in Table 3-1 can be altered (changed) using this -ALTER buffer pool command.

The PGFIX keyword on the ALTER buffer pool command can be key to reducing the CPU used by DB2 when processing a buffer pool. It is discussed in more detail in the following subsection.

DB2 for z/OS manages multiple buffer pools very well. Do not be afraid to separate table spaces and indexes into multiple buffer pools. Between 4 kb, 8 kb, 16 kb, and 32 kb buffer pools, more than 30 different buffer pools are defined when validating the IBM Smart Analytics System 9600. You will find it necessary to define 4 kb (BP0), 8 kb (BP8K0), and 16 kb (BP16K0) buffer pools for use by the DB2 catalog. It is also necessary to define a large number of 32 kb pool pages for use by DB2's RDS sort component, in addition to the 4 kb sort pages. We used buffer pools BP7 for 4 kb and BP32K7 for the 32 kb buffer pools. You can name the sort buffer pools anything you like. In fact, if you are using data

sharing, you will have to use different sort buffer pools names on each data sharing member.

Buffer pool page fixing

To help reduce CPU overhead, DB2 V8 conversion mode (CM) introduced a buffer pool feature that can have a significant effect on CPU usage by long-term page fixing selected DB2 buffer pools.

When a page is brought into storage, DB2 will fix and release the page for I/O processing as required by the channel. The CPU cost for this operation using the 64-bit instruction can be as high as 10%. To avoid this CPU cost for every page being touched by DB2, the ALTER BUF-FERPOOL command has an option to page fix one or more entire buffer pools. After monitoring buffer pool activity and storage usage, we suggest first page fixing the pools with the highest I/O rates. If you need more granularity, start with the pools with the poorest hit ratios.

Tip: Use relative *I/O intensity* to determine which buffer pools are best candidates for PGFIX(YES). The higher the I/O intensity is, the better. The following formula can be used to calculate I/O intensity.

```
((Sync Reads + Async Pages Read) by (SPF + Async Pages Read) by (LPF
+ Async Pages Read) by (DPF + Sync Writes + Async Pages Written)) /
VPSIZE
```

In either case, always start with the pools most critical to the overall performance of your DB2 subsystem. We also suggest just page fixing a few pools at a time, then monitoring your storage usage before page fixing additional pools. There usually is not a need to page fix all of the buffer pools, although the more buffer pools that are long-term page fixed, the greater the possible CPU savings. However, nothing is free. You must ensure that the storage page fixed does not exceed your real storage. DB2 needs to remain 100% backed by real storage and the buffer pools, in most cases, can account for the greatest percentage of real storage use.

To change the PGFIX option, which is NO by default, use the ALTER BUFFERPOOL command:

```
ALTER BUFFERPOOL (bpname) PGFIX ( NO | YES )
```

When altering the PGFIX option, the buffer pool does not get long-term page fixed in real storage until that buffer pool's next allocation. To have the page fixed pool take affect sooner rather than later, some actions needs to take place that will force the buffer pools to be reallocated. For any buffer pool other than the three pools used by the catalog, perform the following three commands:

```
ALTER BUFFERPOOL (bpname) PGFIX(YES)
ALTER BUFFERPOOL (bpname) VPSIZE(0)
ALTER BUFFERPOOL (bpname) VPSIZE(integer value)
```

3.1.2 Stored procedures

DB2 for z/OS provides stored procedures that you can call in your application programs. The following stored procedures reside on the database server:

- DSNUTILU
- DSNWZP
- ADMIN_JOB

DSNUTILU

The DSNUTILU stored procedure enables you to provide control statements in Unicode UTF-8 characters instead of EBCDIC characters to execute DB2 utilities from a DB2 application program. For more information, see:

http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/index.jsp?topic=/com.ibm.db29.doc.ugref/db2z_sp_dsnutilu.htm

DSNWZP

Use DSNWZP to retrieve the DSNZPARMs of the connected subsystem. DSNWZP has only one OUT parameter, which is registered and then called to execute the stored procedure. Because DSNWZP does not issue a return code, we can simply retrieve the OUT parameter and tokenize it using the split() method, which is new since Java 1.4. The split() method splits the string around matches of the given regular expression and returns a string array, which we print to the terminal (Example 3-1).

Example 3-1 Calling DSNWZP and handling the output

```
//Query ZPARM
cs = con.prepareCall("CALL SYSPROC.DSNWZP(?)");
cs.registerOutParmater(1,Types.LONGVARCHAR); //ZPARMs
cs.execute();

String[] zparms = cs.getString(1).trim().split("[/\n]");

System.out.println(“ ----------------------------------------");
```

```
for (int i = 0; (i + 7) < zparms.length; i += 7)
{
   System.out.println("Internal field name = " + zparms[i]);
   System.out.println("Macro name = " + zparms[i + 1]);
   System.out.println("Parameter name = " + zparms[i + 2]);
   System.out.println("Install panel name = " + zparms[i + 3]);
   System.out.println("Install panel field number = " + zparms[i + 4]);
   System.out.println("Install panel field name = " + zparms[i + 5]);
   System.out.println("Value = " + zparms[i + 6]);
}
cs.close();
```

ADMIN_JOB

ADMIN_JOB uses ADMIN_JOB_SUBMIT to submit JCL to compress an existing partitioned data set (PDS). Then it uses ADMIN_JOB_QUERY to poll the job status until the job is in the OUT queue. It then uses ADMIN_JOB_FETCH to fetch the job output and print it. Finally, it calls ADMIN_JOB_CANCEL to purge the job output.

- The ADMIN_JOB_SUBMIT stored procedure load module name is DSNADMJS and its package name is DSNADMJS. ADMIN_JOB_SUBMIT runs in a WLM-established stored procedures address space. For more information about this stored procedure, see:

  http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/index.jsp?topic=/com.ibm.db29.doc.admin/db2z_sp_adminjobsubmit.htm

- The ADMIN_JOB_QUERY stored procedure load module name is DSNADMJQ and resides in an APF authorized library. ADMIN_JOB_QUERY runs in a WLM-established stored procedure's address space, and all libraries in this WLM procedure STEPLIB DD concatenation must be APF-authorized. For more information about this stored procedure, see:

  http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/index.jsp?topic=/com.ibm.db29.doc.admin/db2z_sp_adminjobquery.htm

- The load module for ADMIN_JOB_FETCH is DSNADMJF and also must reside in an APF authorized library. For further details about the ADMIN_JOB_FETCH stored procedure, see:

  http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/index.jsp?topic=/com.ibm.db29.doc.admin/db2z_sp_adminjobfetch.htm

- ADMIN_JOB_CANCEL's load module name is DSNADMJP. This procedure runs in a WLM-established stored procedures address space, and all libraries in this WLM procedure STEPLIB DD concatenation must be APF-authorized. For more information, see:

  http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/index.jsp?topic=/com.ibm.db29.doc.admin/db2z_sp_adminjobfetch.htm

- ADMIN_DS_BROWSE returns either text or binary records from certain data sets or their members. You can browse a physical sequential (PS) data set, a generation data set, a partitioned data set (PDS) member, or a partitioned data set extended (PDSE) member. This stored procedure supports only data sets with LRECL=80 and RECFM=FB. The load module for ADMIN_DS_BROWSE, DSNADMDB, must reside in an APF-authorized library. ADMIN_DS_BROWSE runs in a WLM-established stored procedures address space, and all libraries in this WLM procedure STEPLIB DD concatenation must be APF-authorized. For more information, see:

  http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/index.jsp?topic=/com.ibm.db29.doc.admin/db2z_sp_adminjobfetch.htm

3.1.3 Database partition group design

Range partition data is used in the IBM Smart Analytics System 9600, which allows you to put data into different "buckets" depending on the ranges that were specified. So, for example, all of the January 2010 data could be in one tablespace partition with the data for all of June 2010 in another tablespace partition, and so on. The data is spread across multiple tablespace partitions for better query performance, improved parallelism, and easier tablespace management.

The large volume of data stored in data warehouse environments can introduce challenges to database management and query performance. The tablespace partitioning feature of DB2 for z/OS currently has the following characteristics to aid in addressing those challenges:

- Maximize availability or minimize run time for specific queries by allowing queries and utilities to work at the partition level.
- Grow to 4096 partitions, with each partition as a separate physical data set.
- Allow loading and refreshing activities, including the extraction, cleansing, and transformation of data, in a fixed operational window.
- Increase parallelism for queries and utilities. Parallelism can be maximized by running parallel work across multiple partitions. The number of CPs (general purpose processors) and the number or partitions are the greatest influencers of the degree of parallelism that can be obtained by a query running in DB2.

- Accommodate data growth. A partition-by-growth universal table space can grow automatically up to 128 TB and has the functionality of segmented table spaces while retaining the size and partition independence allowed by a partitioned table space.
- Perform data recovery or restoration at the partition level if data should become damaged or otherwise unavailable, improving availability and reducing elapsed time.

3.2 DB2 for z/OS settings and configuration

Enterprises across the world are increasing focus on their data warehouse and BI initiatives. It continues to be a strong focus area in the overall strategic plan of most enterprises across the world.

As data warehouses and BI environments continue to grow rapidly, and BI insights become critical components to operational workloads, customers, as well as IBM organizations, have expressed a strong interest in how this growing data can be efficiently stored, processed, and managed.

Customers have expressed a strong interest in understanding how data warehousing solutions built on the System z platform running with DB2 9 or DB2 10 for z/OS can be the answer to their growing requirements. The configuration, scalability, and management of data warehousing solutions on System z to create a balanced data warehouse is where enterprises would like to see themselves.

It is critical to explore the scaling and management of very large data warehouses with DB2 for z/OS and System z.

DB2 for z/OS has been supporting data warehousing for more than 25 years. It has continually delivered features and functions in direct or indirect support of data warehousing and the associated BI applications. The following list details the more significant DB2 features that can enhance your data warehousing experience:

- Resource Limit Facility (RLF)

  Introduced in DB2 V2.1, this allows for the control of the amount of CPU resource that a task (in this case a query) can actually use. RLF affects dynamic SQL, which can comprise a significant portion of the data warehouse SQL workload. As examples, it can be critical in controlling system resources, and can help you control the degree of parallelism obtained by a query.

- Hardware assisted data compression

  Delivered with DB2 V3, this still has a major and immediate effect on data warehousing. Enabling compression for table spaces can yield significant disk savings. In testing, numbers as high as 80% have been observed. I/O parallelism, CP parallelism, and Sysplex query parallelism

  These features became available in DB2 Version 3, Version 4, and Version 5 respectively. With the first iteration, multiple I/O could be started in parallel to satisfy a read request. Next, a query could run across two or more CPs. For example, a query could be segmented into multiple parts, and each part could run against its own Service Request Block (SRB), performing its own I/O. With the delivery of data sharing, a query can run across multiple CPs on multiple Central Electronic Complexes (CECs) in the Parallel Sysplex®. There is additional CPU used for setup when DB2 first decides to run a query parallel. There is a correlation between the degree of parallelism achieved and the elapsed time reduction. There are also DSNZPARMs and bind parameters that need to be set before parallelism can be used.

  Because DB2 compression is specified at the tablespace level and is assisted by the System z hardware, compressed data is also carried through into the buffer pools. This means that compression could have a positive effect on reducing the amount of logging that you do because the compressed information is carried into the logs. This reduces your active log size and the amount of archive log space needed.

  Compression also can improve your buffer pool hit ratios. With more rows in a single page after compression, fewer pages need to be brought into the buffer pool to satisfy a query get page request. An additional advantage of DB2 hardware compression is the hard speed. As hardware processor speeds increase, so does the speed of the compression built into the hardware chipset.

  When implementing a data warehouse, the growth in size can become problematic, regardless of the platform. DB2 hardware compression can help

address that issue by reducing the amount of disk needed to fulfill your data warehouse storage requirements. For further details see Chapter 3, “DB2 design for the Enterprise Data Warehouse” on page 15.

- Index compression

  Introduced in DB2 9, index compression is a mechanism used to reduce the amount of storage used by indexes. Index compression averages around 50% for most indexes. Indexes can be a significant performance tool in a data warehouse environment. Reducing the space used by indexes can help to provide more storage.

  Index compression has the ability to compress an index without the use of a dictionary. Without the need for a dictionary, compression starts immediately when the first key is added to the index without any additional processing needed.

  There are no performance gains when using index compression. It is available as a device that helps to save disk space only.

- Parallelism

  One method of reducing the elapsed time of a long-running query is to segment that query across multiple processors. This is exactly what DB2 parallelism does. Parallelism allows a query to run across multiple CPs. A query is segmented into multiple parts, with each part running under its own SRB and performing its own I/O. Although there is additional CPU used when DB2 decides to take advantage of query parallelism for its setup, there is a close correlation between the degree of parallelism achieved and the time reduction for the query. There also are DSNZPARMs and bind parameters that need to be set before parallelism can be used.

  There are three different types of parallelism available with DB2:

  – I/O parallelism
  – CP parallelism
  – Sysplex query parallelism

  I/O parallelism became available with DB2 Version 3. With I/O parallelism, multiple I/O could be started in parallel to satisfy a read request.

  Next, DB2 Version 4 introduced CP parallelism. CP parallelism allowed a query to take advantage of multiple (two or more) CPs. CP parallelism enables true multitasking within a query. A large query can be broken into multiple smaller queries. These smaller queries run simultaneously on multiple processors accessing data in parallel, reducing the elapsed time for each query. Starting with DB2 V8, the parallel queries exploit zIIPs when they are available on the system, thus reducing the costs.

In DB2 Version 5, taking advantage of the recently delivered data sharing feature, Sysplex query parallelism allowed a query to run across multiple CPs on multiple Central Electronic Complexes (CECs) in a Parallel Sysplex.

Additional CPU resources are used to manage running a query in parallel, so parallelism does come at a slight CPU cost. In addition, there is a correlation between the degree of parallelism achieved and the elapsed time reduction.

Parallelism must be enabled through DSNZPARMs and bind parameters that need to be set before parallelism can be used.

- Star schema

  Star schema is a specialized case of the use of parallelism, representing multi-dimensional data, which is often a requirement for data warehousing applications. A star schema usually consists of a large fact table with a number of smaller dimension tables.

- Snowflake schema

  A snowflake schema is similar to a star schema. However, with a snowflake, the dimension tables can have additional dimensions.

- DB2 data sharing

  Data sharing was delivered along with CP parallelism in DB2 Version 4. High availability for data warehousing has now become the norm rather than the exception, and data sharing is capable of giving data warehousing that kind of high availability. DB2 data sharing allows access to the operational data by the data warehouse and analytics, yet still lets you separate those applications into their own DB2, reducing the chances of the data warehouse activity impacting operational transactions.

3.2.1 DSNZPARM

This section discusses the suggested DSNZPARM options that should be considered when implementing a data warehouse with the IBM Smart Analytics System 9600. This list is fairly extensive and should only be used for initial guidance. Although these are the DSNZPARMs that were modified during our testing and used during installation, how we chose to use them should not be considered *global best practices*. Some could be set differently and still allow for positive results.

We have included an explanation as to why we set them to the values that we used to aid in your understanding of the DSNZPARM and to assist you should you decide to modify to a different value in the future. The keyword explanations are grouped together by subsystem parameter macros. Additional information about these keywords is available in the *DB2 Installation Guide* or, in some cases, on an APAR's cover.

Table 3-10 on page 36 lists the subsystem parameters and the values that we changed, along with a brief explanation of why we changed them. All DSNZPARM keyword descriptions and default values are based on DB2 9 for z/OS.

DSN6SPRM

DSN6SPRM is a macro found on the DSNTIPO panel. Table 3-2 lists the subsystem parameters and the values that we changed, along with a brief explanation of why we changed them.

Table 3-2 Subsystem parameters values for macro DSN6SPRM

Subsystem parameter	Allowed values	Default value	Set value	Explanation
CACHEDYN	YES,NO	YES (Enable dynamic SQL caching.)	YES	Caching dynamic SQL allows it to be considered for reuse, thus reducing the necessity for PREPAREs and resulting in a potential reduction in CPU. The initial installation sets this keyword to its default, YES.
CDSSRDEF	1,ANY	1 (Disable use of parallelism.)	ANY	CDSSRDEF defines the default value for the parallelism special register. If the special register is not set prior to using parallelism in an SQL statement, DSNZPARM is used to determine whether parallelism is enabled. The default is 1. However, because this is a warehouse system and we want to minimize the run time (elapsed time) of any SQL statements, we suggest setting this keyword to ANY, enabling parallelism by default for those situations where the special register is not explicitly set. Bind and CURRENT DEGREE special register also need to be set to ANY.

Subsystem parameter	Allowed values	Default value	Set value	Explanation
CONSTOR	NO,YES	NO (YES in DB2 v10)	NO	During installation, this keyword is set to NO. CONTSTOR saw a lot of use when DB2 was having storage issues back in Version 7. Each new version of DB2 has reduced the occasions for that to happen. However, some might still set the keyword to YES just out of habit or "because that was how it was always done." If you are not storage constrained, do not set this to YES. If you *are* storage constrained for the DBM1 address space, then turn this option on by setting it to YES.
DSMAX	1 to 100,000	Default calculated		DSMAX determines the maximum number of data sets that can be open in DB2 at one time. At first installation, the initial value for this system parameter is calculated, although in many cases it still needs to be resized based on the actual database configuration implemented.
IRLMRWT	1 to 3600	60	15	IRLMRWT represents the amount of time in seconds that DB2 will wait before timing out. Fifteen seconds is the suggested value only because the default is too large. If a query cannot get to a resource, it is better to find out as quickly as possible.
LRDRTHLD	1 to 1439	0 (DB2 v10 reduces this default to 10.)		This subsystem parameter describes how long a claim has been held in minutes. It is a good indicator of what read operations (queries) run for an excessive amount of time. We suggest that this value be set fairly high. Our sample is set to 20 minutes. Any query that holds a claim longer than 20 minutes will have DB2 cut a trace record to report the offending query.

Subsystem parameter	Allowed values	Default value	Set value	Explanation
MAXRBLK	0, 128k to 10,000,000	8000 k (DB2 v10 increases this to 400,000 k.)	100,000	MAXRBLK is the size of this DB2 subsystem's RID pool. The value used at installation was 100,000. However, this number should be adjusted based on your query workload's use of the RID pool.
MAX_OPT_CPU	0 to 1000 seconds	100 seconds	based on installa-tion	Thresholds for the enhanced internal monitoring of how much CPU can be consumed by the optimization process to avoid excessive resource consumption. The value used for this opaque subsystem parameter should be determined by the amount of PCU available and query processing time.
MINSTOR	NO,YES	YES (DB2 v10 is NO.)	NO	The MINSTOR subsystem parameter controls whether DB2 is to use storage management algorithms that minimize the amount of working storage that is consumed by individual threads. If set to yes, the reduction in storage does come at a CPU cost. Because the default was changed to YES in DB2 9, it is mentioned here to make sure that it is set back to NO. It should be set to YES only if this subsystem is having an available storage issue.
MXDTCACH	0 to 512	20	128	MXDTCACH specifies the maximum size of memory for data caching of each thread. Increasing MXDTCACH to 128 MB could minimize the amount of random activity that might spill over to the sort work buffer pool.

Subsystem parameter	Allowed values	Default value	Set value	Explanation
MXQBCE			1023	MXQBCE limits how many different join sequences DB2's optimizer will consider. The lower the number, the fewer the number of join sequences considered. By reducing the number of join sequences considered, you could reduce the time that DB2 spends in bind processing. Of course, spending less time could also mean that a less appropriate access path could be chosen.
OPTHINTS	NO/YES	NO	YES	Set to YES to enable the use of optimization hangs. In general, this is the preferred setting, and using YES has no negative effects.
OPTHYBCST	NO/YES	NO	NO	When set to YES, this DSNZPARM enables a cost model improvement for hybrid joins with SORTN_JOIN = N. This parameter is deprecated in DB2 9 and removed from the product in DB2 10. The DB2 10 behavior is the same as the DB2 9 behavior if this parameter was set to YES.
OPTIOWGT	ENABLE DISABLE	ENABLE	ENABLE	When this parameter is set to ENABLE, DB2 uses a new formula that better balances the cost estimates of I/O response time and CPU usage when selecting access path. We left this parameter set to its default. This parameter is also deprecated in DB2 10.
OPTJBPL			ON	
OBTJBPR			OFF	
OPTOFNRE			ENABLE	
OPTOIRCPF			ENABLE	

Subsystem parameter	Allowed values	Default value	Set value	Explanation
PARAMDEG	0 to 254	0		PARAMDEG controls the maximum degree of parallelism allowed by the DB2 subsystem. The value suggested for this keyword is based on the formula #Processors <= X <= 2*#Processors. For a new data warehouse, where query performance metrics have not yet been established, we suggest that the default of 0 not be used. Specifying 0 or taking the default could allow the initial over allocation of processors, potentially causing performance issues. The actual degree of parallelism chosen by DB2 is highly influenced by the number of processors available and the number of partitions in the table space being accessed. The higher either or both of these values are, the higher the degree of parallelism used. Parallel processing is one of the functions within DB2 that is eligible to run on the IBM System z Integrated Information Processor (zIIP). Once an internal threshold is reached for the amount of processing performed by a parallel query, the parallel child task will be zIIP eligible with no predetermine zIIP limit. Up to 90% of all parallel child task work is eligible to run on a zIIP after the initial threshold is reached.
PLANMGMT	OFF, BASIC, EXTENDED	OFF (DB2 v10 default is extended)	OFF	Setting PLANMGMT to a value other than OFF allows DB2 to use a previous copy of a packages should a REBIND introduce an issue through the currently bound copy. The assumption is made here that most analytic SQL is dynamic. PLANMGMT currently only applies to static SQL. If static SQL should be used by the analytical process, then the use of PLANMGMT should be considered.

Subsystem parameter	Allowed values	Default value	Set value	Explanation
PREDPRUNE			YES	
RRULOCK			YES	
SEQCACH			SEQ	
SJMISSKY			ON	
SJTABLES			10	Based on installation.
SRTPOOL			LARGE	Based on installation. 8000 = 8 MB sort pool.
STARJOIN			DISABLE	
STATCLUS			ENHANCED	
STATROLL			YES	
UNION_COLNAME_7			YES	
WFDBSEP			YES	

Using star joins in DB2 requires enabling the feature through a DSNZPARM keyword. You also should check a few other DSNZPARMs before using star joins because they can affect a star join's performance.

There is a specialized case of parallelism called a star schema—a relational database's way of representing multi-dimensional data—that is often popular with data warehousing applications. A star schema is usually a large fact table with lots of smaller dimension tables. For example, you might have a fact table for sales information. This sales table would hold most of your data. The dimension tables could represent products that were sold, the stores where those products were sold, the date the sale occurred, any promotional data associated with the sale, and the employee responsible for the sale. Using star joins in DB2 requires enabling the feature through the DSNZPARM keyword (Table 3-2 on page 27).

DSN6ARVP

DSN6ARVP is a macro found on the DSNTIPA panel. Table 3-4 lists the subsystem parameters and the values that we changed with a brief explanation of why we changed them.

Table 3-3 Subsystem parameters values for macro DSN6ARVP

Subsystem parameter	Allowed values	Default value	Set value	Explanation
PRIQTY			1000	
UNIT				Based on installation
UNIT2				Based on installation

DSN6LOGP

This is the third group of are keywords on the DB2 system parameter (DSNZPARM) macro. DSN6LOGP is found on the DSNTIPO panel. Table 3-4 lists the subsystem parameters and the values that we changed with a brief explanation of why we changed them.

Table 3-4 Subsystem parameters values for macro DSN6LOGP

Subsystem parameter	Allowed values	Default value	Set value	Explanation
OFFLOAD			NO	Based on installation

DSN6SYSP

Table 3-5 lists the subsystem parameters and the values that we changed with a brief explanation of why we changed them.

Table 3-5 Subsystem parameters values for macro DSN6SYSP

Subsystem parameter	Allowed values	Default value	Set value	Explanation
ACCUMACC			NO	
ACCUMUID			0	
CHKFREQ			3	
CONDBAT				This is based on the size of the installation.

Subsystem parameter	Allowed values	Default value	Set value	Explanation
CTHREAD				This is based on the size of the installation.
DSVCI	NO,YES	YES	YES	Although this is the default, ensure that it has not been changed to NO.
IDBACK			>100	Based on installation.
IDFORE				Lower than IDBACK and based on installation.
MAXDBAT				This is based on the size of the installation.
MGEXTSZ	NO,YES	YES	YES	Although this is the default, ensure that it has not been changed to NO.
PCLOSEN			5	
PCLOSET			10	
PTASKROL			YES	
RLF				Consider based on installation standards.
SMFACCT			(1,2,3,7,8)	
SMFSTAT			(1,3,4,5,6)	
STATIME			3	
SYNCVAL			0	
OTC_LICENSE			NOT_USED	Turn on for DB2 VUE

DSN6FAC

Table 3-6 lists the subsystem parameters and the values that we changed with a brief explanation of why we changed them.

Table 3-6 Subsystem parameters values for macro DSN6FAC

Subsystem parameter	Allowed values	Default value	Set value	Explanation
CMTSTAT			INACTIVE	
IDTHTOIN			120	
TCPKPALV			120	Validate existing TCP/IP settings.
PRIVATE_PROTOCOL			YES	

DSN6SPRC

Table 3-7 lists the subsystem parameters and the values that we changed with a brief explanation of why we changed them.

Table 3-7 Subsystem parameters values for macro DSN6SPRC

Subsystem parameter	Allowed values	Default value	Set value	Explanation
SPRMPTH			200	

DB2 customization

This section describes the processes, steps, and related information that is utilized in the delivery of the DB2 subsystem. The delivered system included the DB2 SMP/E target and distribution library data sets and was used as the base in generating the system.

As a reference, the default values shown in Table 3-8 and Table 3-9 on page 36 were used during the z/OS and DB2 installation process when defining your DB2 subsystem.

Table 3-8 SMS configuration values

Storage class	Default VOLSER	Usage
DB2SYSTM	PJDSC1	DB2 cat/dir tables
	PJDSC2	DB2 cat/dir indexes

Storage class	Default VOLSER	Usage
DB2LOGS	PJDLG1,2,3,4	Logs, BSDS's
	PJDLG5,6,7,8	Logs, BSDS's
DB2WORK	PJDW01-10	DSNDB07
NONSMS	PJDDLB	Libraries, ZFS file
DB2DATA	PJD000-10	Application data

Table 3-9 D B2 DSNZPARM values

Parm	Default
SSID	DB2I
IRLMID	I2BD
CRC	-DB2I
EBCDIC CCSID	37
ASCII CCSID	437
CHKFREQ	15 MIN
LOCATION	DB2I
VTAM® LU	DB2ILU
DRDA® PORT	446
RESYNC PORT	5020

Table 3-10 contains the remaining DB2 DSNZPARM values as installed with the IBM Smart Analytics System 9600.

Table 3-10 DSNZPARMS

DSNZPARM macro name	ISAS 9600 Config
DSN6SPRM	CACHEDYN=YES
DSN6SPRM	CDSSRDEF=ANY
DSN6SPRM	CONTSTOR=NO
DSN6SPRM	DBACRVW=YES

DSN6SPRM	DSMAX= needs adjusting based in size
DSN6SPRM	INLISTP=50
DSN6SPRM	IRLMRWT=15
DSN6SPRM	LRDRTHLD=20
DSN6SPRM	MAXRBLK >= 100000
DSN6SPRM	MAX_OPT_CPU=based on installation
DSN6SPRM	MINSTOR=NO
DSN6SPRM	MXDTCACH 128
DSN6SPRM	MXQBCE 1023
DSN6SPRM	NUMLKTS=1000
DSN6SPRM	NUMLKUS=10000
DSN6SPRM	OPTHINTS=YES
DSN6SPRM	OPTHYBCST=OFF
DSN6SPRM	OPTIOWGT=ENABLE
DSN6SPRM	OPTJBPL=ON
DSN6SPRM	OPTJBPR=ON
DSN6SPRM	OPTOFNRE=ENABLE
DSN6SPRM	OPTOIRCPF=ENABLE
DSN6SPRM	PARAMDEG=#Processors <= X <= 2*#Processors
DSN6SPRM	PLANMGMT=OFF
DSN6SPRM	PREDPRUNE=YES
DSN6SPRM	RRULOCK=YES
DSN6SPRM	SEQCACH=SEQ
DSN6SPRM	SJMISSKY=ON
DSN6SPRM	SJTABLES=10 (based on installation)
DSN6SPRM	SRTPOOL= large (based on installation)
DSN6SPRM	STARJOIN=DISABLE
DSN6SPRM	STATCLUS=ENHANCED

DSN6SPRM	STATROLL=YES
DSN6SPRM	UNION_COLNAME_7=YES
DSN6SPRM	WFDBSEP=YES
DSN6ARVP	PRIQTY=1000
DSN6ARVP	UNIT=based on installation
DSN6ARVP	UNIT2=based on installation
DSN6LOGP	OFFLOAD=NO based on installation
DSN6SYSP	ACCUMACC=NO
DSN6SYSP	ACCUMUID=0
DSN6SYSP	CHKFREQ=3
DSN6SYSP	CONDBAT=based on 9600 Size
DSN6SYSP	CTHREAD=based on 9600 Size
DSN6SYSP	DSVCI=YES
DSN6SYSP	IDBACK= >100 (based on installation)
DSN6SYSP	IDFORE= lower than IDBACK based on installation
DSN6SYSP	MAXDBAT=based on 9600 Size
DSN6SYSP	MGEXTSZ=YES
DSN6SYSP	PCLOSEN=5
DSN6SYSP	PCLOSET=10
DSN6SYSP	PTASKROL=YES
DSN6SYSP	RLF=consider (based on installation standards)
DSN6SYSP	SMFACCT=(1,2,3,7,8)
DSN6SYSP	SMFSTAT=(1,3,4,5,6)
DSN6SYSP	STATIME=3
DSN6SYSP	SYNCVAL=0
DSN6SYSP	OTC_LICENSE=NOT_USED (turn on for DB2 VUE)
DSN6FAC	CMTSTAT=INACTIVE
DSN6FAC	IDTHTOIN=120

DSN6FAC	TCPKPALV=120 (Validate existing TCP/IP settings)
DSN6FAC	PRIVATE_PROTOCOL=YES
DSN6SPRC	SPRMPTH=2000
DSNTIPD	
DSNTIP7	

3.2.2 Logging and backup considerations

DB2 for z/OS logging considerations

3.3 DB2 9 for z/OS enhancements and features for data warehousing

DB2 9 for z/OS delivers changes that will directly impact your data warehouse and application analytics, including the following:

- New row internal structure for faster VARCHAR processing
- Fast delete of all the rows in a partition (TRUNCATE)
- Deleting first n rows
- Skipping uncommitted inserted/updated qualifying rows
- Index on expression
- Dynamic index ANDing
- Reducing temporary tables materialization
- Generalizing sparse index/in-memory data caching
- Clustering decoupled from partitioning
- Indexes created as deferred are ignored by DB2 optimizer
- Fast cached SQL invalidation
- Statements IDs of cached statements as input to EXPLAIN
- Universal tablespaces
- Partition-by-growth to remove non-partitioned tablespace size limit
- Implicating objects creation
- Cloning tables
- MERGE statement
- Identifying unused indexes
- Simulating indexes in EXPLAIN (Optimization Service Center)
- More autonomic buffer pools tuning for WLM synergy
- Resource Limit Facility (RLF) support for end-user correlation
- RANK, DENSE_RANK, and ROW_NUMBER

- EXCEPT and INTERSECT
- pureXML®

DB2 10 for z/OS enhancements and features for data warehousing

DB2 10 for z/OS delivers scale, complexity, and productivity changes that will directly impact your data warehouse and application analytics, including the following:

- Enhanced query parallelism (restrictions removed)
- On-the-fly data compression
- Temporal (versioned) data support
- More online schema changes (data definition on demand)
- More SQL compatibility
- Moving SUM, moving AVG
- Improved pureXML performance and usability
- Hash access
- Index include columns
- Inline large objects
- Parallel index updates
- Work file in memory
- Member clustering of universal table spaces
- Efficient caching of dynamic SQL statements with literals
- Security enhanced for better granularity for admin privileges
- IBM Smart Analytics Optimizer

  The IBM System z10 EC is a general-purpose server for computation-intensive workloads (such as business intelligence) and I/O-intensive workloads (such as transaction and batch processing). It continues to offer all the specialty engines available with its predecessor z9®, such as:

  – ICF: Internal Coupling Facility used for z/OS clustering. ICFs are dedicated for this purpose and exclusively run Coupling Facility Control Code (CFCC).

  – IFL: Integrated Facility for Linux Exploited by Linux and for z/VM processing in support of Linux. z/VM is often used to host multiple Linux virtual machines (called guests.)

  – SAP: System Assist Processor offloads and manages I/O operations. Several are standard with the z10 EC. More can be configured if additional I/O processing capacity is needed.

- zAAP: System z10 Application Assist Processor is exploited under z/OS for designated workloads, which include the IBM JVM and some XML System Services functions.
- zIIP: System z10 Integrated Information Processor is exploited under z/OS for designated workloads, which include some XML System Services, IPSec off-load, part of DB2 DRDA, complex parallel queries, utilities, global mirroring (XRC), and some third-party vendor (ISV) work.

For more details on the features and enhancements of DB2 10 for z/OS see the following websites:

http://www.ibm.com/software/data/db2/zos/db2-10

http://wwwlibm.com/common/ssi/rep_ca/5/877/ENUSZP10-0015/ENUSZP10-0015.pdf

http://www.ibm.com/support/docvierw.wss?uid=swg27017960

3.4 Database and enterprise data warehouse design considerations

DB2 for z/OS can contain an enormous amount of data. DB2 for z/OS can support up to 64,000 databases. With each database containing up to 32,000 objects, it can easily cater to the growing need of a data warehouse environment.

DB2 9 for z/OS utilizes universal tablespaces, table partitioning, indexing, data and index compression, stored procedures, materialized query tables (MQTs), work files, cubes, fact tables, dimension tables, and multi-level security.

3.4.1 Tablespaces, tables, indexes, compression, stored procedures

The large volume of data stored in data warehousing environments can introduce challenges to database management and query performance.

Universal tablespace

This is a key DB2 enhancement in support of data warehousing. Consider the sometimes unpredictable but expected growth of a data warehouse and the high possibility that many tables could be frequently refreshed. A universal tablespace is a cross between a partitioned tablespace and a segmented tablespace, giving you many of the best features of both. When using a universal tablespace, you get the size and growth of partitioning while retaining the space management, mass delete performance, and insert performance of a segmented tablespace. It is similar to having a segmented tablespace that can grow to a 128 TB of data,

assuming that the correct DSSIZE and correct number of partitions are specified, and that also gives you partition independence.

For more details about the partitioning feature of DB2 9 for z/OS, see *Enterprise Data Warehousing with DB2 9 for z/OS*, SG24-7637

Data compression

Compression has been around for a long time, and every customer is aware of the significant storage benefits that compression offers to their data warehouse. What is not so well known is the impact of compressed data to query performance. Compressed tables use fewer pages and could lead to performance improvement for certain queries. Index compression is a new feature in DB2 9, and little performance data is available. Similar to data compression, there is an interest in understanding the effect of storage savings, query performance, and CPU overhead with index compression.

Data compression offers significant benefits:

- Reduction in the storage space used by the data warehouse
- Reduced elapsed time for most data warehouse type queries
- Reduced I/O time
- More effective use of buffer pool space
- Higher buffer pool hit ratio under certain conditions

Compression can be implemented using either hardware or software. DB2 9 on z/OS uses hardware compression for data, but uses software compression for indexes. The difference between the two is described in the following sections.

Hardware compression

Hardware compression has the compression algorithms built into the hardware. Thus, minimal CPU overhead is required to compress and decompress data. A key point is that hardware compression keeps getting faster as chip speeds increase, although software compression speed increases at the same time.

Other advantages of hardware compression on System z are:

- It reduces CPU overhead, saving valuable CPU bandwidth.
- Higher data throughput.
- Faster than software compression.
- Less costly than software compression.
- Runs as a black box, performing compression and decompression.

DB2 for z/OS compresses rows within a page, so that each data page consists of compressed rows. It uses the hardware instruction along with a data dictionary to give the most efficient compression available. The compressed data can also be encrypted, thereby saving space and implementing security requirements at the

same time. The encryption tool was recently changed to be able to compress and encrypt efficiently. It will compress the data, then encrypt it. With a 50% compression rate, a compressed page contains twice the rows that an uncompressed page would contain. This means that each I/O retrieves twice as much compressed data as it would if the data was uncompressed. The data remains compressed in the buffer pool, which means that DB2 for z/OS can cache twice as much compressed data in its buffer pool as it would if the data were uncompressed. Finally, when data is modified in a row that is compressed, the information logged about that data change is also compressed, thus reducing log volume for both the active logs and archive logs.

Not all data on a compressed page is decompressed, just the rows needed by the application. Combined with the use of the hardware instruction to perform the decompression, this serves to limit the amount of additional CPU needed to access compressed data. The larger amount of data retrieved in each I/O gets compounded with the DB2 9 for z/OS increased prefetch quantities. This provides significant elapsed time reductions for all types of sequential processes, including the typical BI queries that make use of table scans and index range scans. This also includes sequential processes for utility access, providing benefits in terms of faster reorganizations, faster unloads, and faster recovery.

Building the compression dictionary is one of the critical components of data compression. The better that the dictionary reflects the data, the higher the compression rate achieved. The dictionary is built through use of either the LOAD or the REORG utilities. These are the only two utilities and the only two ways that can be used to create a dictionary. Creating or rebuilding the dictionary can be a CPU-intensive task. If the existing dictionary results in an acceptable compression rate, we do not recommend rebuilding the dictionary. It is important to always remember this when running a LOAD or a REORG, as the dictionary will be rebuilt on invocation of the utility. If the existing dictionary is good, then the KEEPDICTIONARY keyword can be used with the LOAD/REORG utilities to keep the existing dictionary and not create a new one.

Details on whether compression is active for an index or table space and metrics describing how effective compression is can be found in the DB2 catalog.

3.4.2 MQTs, views, cubes, and fact table dimension tables

In this section we discuss MQTs, views, cubes, and fact tables dimension tables.

Table space partitioning

The large volume of data stored in data warehousing environments can introduce challenges to database management and query performance. The

table space partitioning feature of DB2 9 for z/OS has the following characteristics:

- Maximizes the availability or minimizes the run time for specific queries
- Can have 4096 partitions

A partition is a separate physical data set.

It allows the loading and refreshing of activities, including the extraction, cleansing, and transformation of data, in a fixed operational window and:

- Increases parallelism for queries and utilities
- Accommodates the growth of data

A universal table space can grow automatically up to 128 TB and adds the functionality of segmented table spaces.

Very large database

DB2 for z/OS can contain an enormous amount of data. DB2 for z/OS can support up to 64,000 databases. With each database containing up to 32,000 objects, it can easily cater to the growing need of a data warehouse environment.

Star schema enhancements

A common data model that is used in data warehouse environments is the star schema, in which a large central fact table is surrounded by numerous dimension tables. Queries generally provide filtering on the independent dimensions, which must be consolidated for efficient access to the fact table.

DB2 for z/OS Version 8 contains the following enhancements, among others, to improve the performance of star schema queries:

- In-memory work files for efficient access to materialized dimensions or snowflakes
- Improved cost formula for join sequence determination
- Predicate localization when OR predicates cross tables

DB2 9 for z/OS further enhances star schema query performance with a new access method. That is Dynamic Index ANDing for simpler index design, more consistent performance, disaster avoidance, and improved parallelism.

Query parallelism

You can significantly reduce the response time for data or processor-intensive queries by taking advantage of the ability of DB2 to initiate multiple parallel operations when it accesses data in a data warehouse environment.

Materialized query tables

MQTs can simplify query processing, greatly improve the performance of dynamic SQL queries, and be particularly effective in data warehousing applications, where you can use them to avoid costly aggregations and joins against large fact tables. The DB2 optimizer uses partial or entire MQTs to accelerate queries. Its operation and access path are also kept for an easy refresh of the MQT content without specifying the source query again.

OLAP functions

New SQL enhancements are made in DB2 9 for z/OS for improving online analytical processing (OLAP) functionalities in a data warehouse. The following OLAP expressions were introduced in DB2 9 for z/OS:

- RANK and DENSE_RANK
- ROW_NUMBER

Table space and index compression

DB2 for z/OS uses the hardware-assisted compression instructions of the System z server for compressing table spaces. DB2 9 for z/OS can also compress index spaces by using software techniques. Table space and index space compression saves a large amount of disk space (in certain cases CPU saving) when implemented in a data warehouse environment, considering the amount of data and the number of indexes that are created for query performance on the large tables.

Index on expression

DB2 9 for z/OS supports the creation of indexes on an expression. The DB2 optimizer can then use such an index to support index matching on an expression. In certain scenarios, it can enhance the query performance. In contrast to simple indexes, where index keys are composed by concatenating one or more table columns specified, the index key values are not exactly the same as the values in the table columns. The values are transformed by the expressions that are specified.

CLONE tables

To overcome the availability problem when running certain utilities, such as LOAD REPLACE in a DB2 for z/OS environment, a cloning feature was introduced in DB2 9 for z/OS. A clone of a table can be created by using the ALTER TABLE SQL statement with the ADD CLONE clause. The clone can then be used by applications, SQL, or utilities, and therefore provide high availability.

For further details about special DB2 for z/OS table considerations, see *Enterprise Data Warehousing with DB2 9 for z/OS*, SG24-7637.

3.4.3 DB2 multi-level security

Multilevel security is a security policy that allows the classification of data and users based on a system of hierarchical security levels combined with a system of non-hierarchical security categories. You can improve the security of your DB2 applications when you add RACF security labels to DB2 objects or row-level security on a multilevel-secure system. Implementing multilevel security is a system-wide endeavor. See *z/OS Planning for Multilevel Security and the Common Criteria*, GA22-7509, for more details.

The IBM Smart Analytics System 9600 DB2 security details should be similar to your existing user security, which limits and controls access to installation data (that is, DB2 Security, RACF, and so on).

The basic idea of MLS with row-level granularity is that any user reading or updating data in a DB2 table needs to be allowed to handle only the rows that his or her security label allows. Each row in a table is assigned a security label, and a user can read the row only if his label dominates the label of the row. Similar rules apply for updating rows in a table with row-level security. Only where updating within an MLS environment is concerned, other principles concerning write-down (that is, the declassification of data) influence the result of the update. Multi-level security was introduced in DB2 v8 for z/OS and improved in DB2 9 for z/OS.

A security label enables an installation to classify subjects and objects according to a data classification policy, identify objects to audit based on their classification, and protect objects such that only appropriately classified subjects can access them.

3.4.4 Subjects and objects

A subject is an entity that requires access to system resources. Examples of subjects are human users, started tasks, batch jobs, and z/OS UNIX® daemons. Examples of objects are data sets, a row within a DB2 table, commands, terminals, printers, and DASD volumes.

Subjects are defined to RACF. For example, a user or started task will have a RACF user ID. Objects (other than rows in a DB2 table) are also defined to RACF as either a resource profile or a data set profile. The terms *subject* and *user ID* have the same meaning and can be used interchangeably. In a multilevel secure system, subjects and objects have a security label associated with them.

The security label is defined to RACF in the resource class SECLABEL. Rows in a DB2 table have a security label associated with them by means of a special

column in the table that contains only the eight-character security label that defines the security classification of each row in that table. A subject's security label determines whether the subject is allowed to access a particular object. An object's security label indicates the sensitivity of that object's data.

A subject is authorized to use a security label by having been permitted READ access to the resource profile in the SECLABEL class in RACF, which defines the particular security label. A TSO user can have a default security label defined in RACF if desired.

3.4.5 Network-trusted context

A powerful security enhancement in DB2 V9 for z/OS is the introduction of the network-trusted context. In itself, it supplies the ability to establish a connection as trusted when connecting to DB2 for z/OS from a certain location. Having established the connection, it provides the possibility of switching to another user ID, thus giving the opportunity of taking on the identity of this other user ID only within the trusted context. In addition, it is possible to assign a role to a user of a trusted context. The role can be granted privileges and can therefore represent a role within the organization in the sense that it can hold the sum of privileges needed to perform a certain job or role. These two constructs together supply security enhancements for a variety of different scenarios, ranging from any three-tier layered application such as SAP to the daily duties of a DBA maintaining the DB2 subsystem. The possibilities are many and varied.

A role can be used as a single database authid that can be used to simplify administration of dynamic SQL privileges. The user's authid can be used to run database transactions, so that the DB2 audit is able to identify the users individually (an important capability for meeting some regulatory compliance requirements). The trusted context retains many of the performance benefits of connection pooling.

The trusted context and role support can be used to implement DBA privileges that can easily be disconnected and reconnected to individual employees. This provides function similar to shared SYSADM or DBADM user IDs, but avoids the audit compliance problems associated with shared user IDs.

A multilevel security system is a security environment that allows the protection of data based on both traditional discretionary access controls and controls that check the sensitivity of the data itself through mandatory access controls. These mandatory access controls are at the heart of a multilevel security environment, which prevents unauthorized users from accessing information at a classification to which they are not authorized or changing the classification of information to which they do have access. These mandatory access controls provide a way to

segregate users and their data from other users and their data regardless of the discretionary access that they are given though access lists and so on.

Creating a multilevel security environment requires a combination of several software and hardware components that enforce the security requirements needed for such a system. The security-relevant portion of software and hardware components that make up this system are also known as the trusted computing base.

For more details on defining security categories, levels, and tables, see the *DB2 v9.1 for z/OS Administration Guide*, SC18-9840, and *Securing DB2 and Implementing MLS on z/OS*, SG24-6480.

3.5 XML and the data warehouse

For IT leaders building data warehouses that meet the evolving demands of their business environments, integration of XML data into their infrastructures is critical. XML has become the preferred data exchange format across many industries. As a result, organizations must find ways to efficiently manage and manipulate XML within their data warehouses.

IBM DB2 pureXML makes it possible for organizations to manage XML data and relational data. This increases database efficiency, improves the user experience, and increases their competitive advantage by fully exploiting data interchange standards.

One of the primary goals of data warehousing is to make it as easy as possible for users to get the information that they need when they need it. Presenting this information to decision makers is a challenge in this environment. Programmers and database analysts must determine which attributes to expose to which decision maker.

Investigating warehouse data in this way requires an intimate knowledge of how the warehouse schemas are constructed. There is often no easy, efficient, or effective way in today's table-based warehouses for developers to create a search function that works like a web search.

Extending a relational data warehouse schema with one or more XML columns avoids those problems. Commonly used attributes can be stored in relational columns, while additional details can be maintained in an XML column, which readily accommodates variable structures and is easily accessible for queries and reports.

As more and more critical business data is captured and exchanged in XML, firms are recognizing the need to manage, share, query, and report on XML data. The increased use of XML standards for data interchange creates storage and management challenges. The highly variable, nested structures are difficult to accommodate using traditional relational database techniques.

The database management system (DBMS) cannot provide optimized access to specific XML elements or attributes contained within a message or document.

Some firms “shred” or decompose XML data into multiple columns of one or more tables. These complex, labor-intensive mappings are difficult to adjust as XML messaging formats change over time.

Many firms are storing XML in its native hierarchical format alongside relational data so that both types of data can be managed in an optimal manner.

DB2 and XML considerations

IBM DB2 provides firms with a common application programming interface and database management platform for data modeled in tables and XML hierarchies. This hybrid database management architecture (Figure 3-1) helps to extend traditional relational database environments to directly manage XML messages and documents without the need to shred data into columns of various tables. Applications can retrieve relevant portions of the XML data easily and efficiently, as well as integrate XML and relational data with little effort.

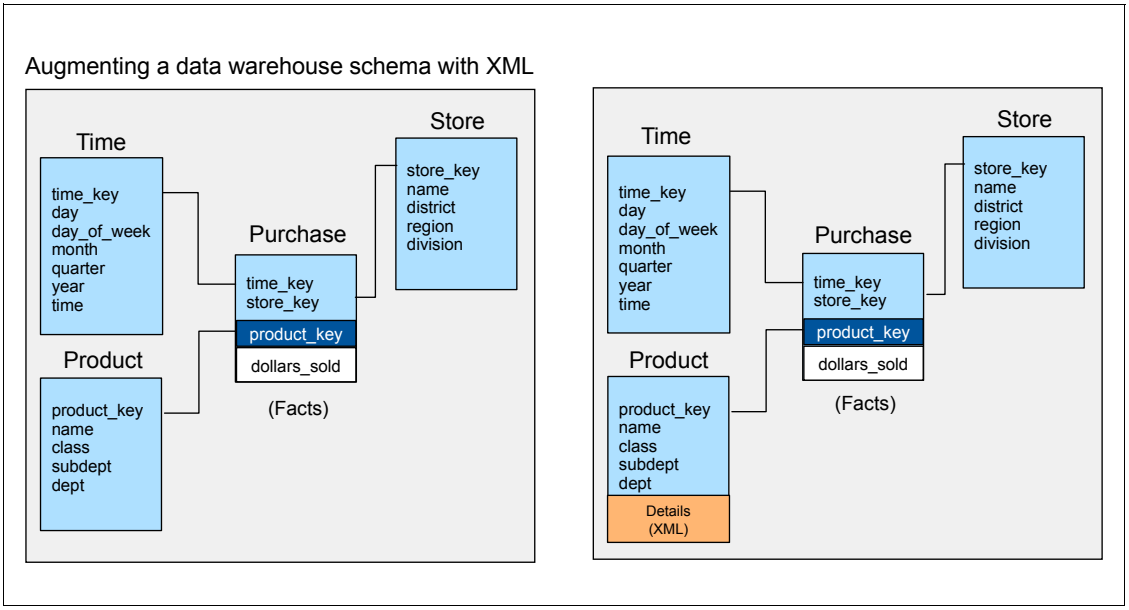

Figure 3-1 Augmenting a data warehouse schema with XML

'DB2 9 architecture with build-in support for relational and XML data helps extend traditional relational database environments.

DB2 9 for z/OS provides pureXML, which is a native XML storage technology that provides hybrid relational and XML storage capabilities. pureXML provides a huge performance improvement for XML applications while eliminating the need to shred XML into traditional relational tables or to store XML as character large objects (CLOBs), which are methods that other vendors use. DB2 9 pureXML exploits z/OS XML System Services for high-performance parsing with improved price performance by using zAAPs and zIIPs.

DB2 pureXML includes these features:

- Cost-based query optimization helps enable DB2 to select an efficient path for accessing requested data.
- Specialized XML indexing speeds retrieval of queries over XML data as well as queries over relational views of XML data.
- Hash-based partitioning provides significant scalability gains.
- Range-based partitioning helps firms “roll in” and “roll out” data over time (a common requirement in data warehouses).
- Multi-dimensional clustering often improves performance of analytic queries.
- Compression of XML data and indexes reduces storage costs, improves storage efficiency, and speeds runtime performance for many common workloads.

For more information about DB2 pureXML, see:

http://www.ibm.com/software/data/db2/xml

4

Managing the IBM Smart Analytics System 9600 components

In this chapter we discuss the startup and shutdown procedures of the IBM Smart Analytic System 9600 components. Administrative tasks are discussed in Chapter 5, “InfoSphere Warehouse administrative tasks” on page 55, and Chapter 6, “Cognos 8 Business Intelligence” on page 69.

Verify with your system administrators that they have already IPLed z/OS, brought up VTAM, started the Linux on System z guests, and brought up DB2 on System z. These tasks must be completed prior to the tasks in this chapter.

4.1 Startup procedure for IBM Smart Analytics System 9600 components

The order in which the IBM Smart Analytics System 9600 components should be started is:

1. InfoSphere components
   a. Log into the Linux on System z guest where the WebSphere Application Server resides.
   b. Start the InfoSphere Warehouse server with the following command:

      ```
      /opt/IBM/ISWarehouse/appServer/profiles/AppSrv01/bin/startServer.
      sh  server1
      ```

   c. Using the InfoSphere Warehouse Administration console from a web browser, start the cube server, start the cube, and start the XML for Analysis (XMLA) interface.
2. Cognos components
   a. Log into the Linux on System z guest where the Cognos components server reside.
   b. Log in to the content manager Linux on System z guest and start the Content Manager. Switch to the Cognos user ID (su - cognos). The command to start the content manager is:

      ```
      /opt/IBM/WebSphere/AppServer/bin/startServer.sh server1
      ```

   c. Log in to the report server Linux on System z guest and start the report servers. Switch to the Cognos user ID (su - cognos). The commands to start the report server are:

      ```
      /opt/IBM/WebSphere/AppServer/bin/startServer.sh server1
      -profileName AppSrv01
      /opt/IBM/WebSphere/AppServer/bin/startServer.sh server1
      -profileName AppSrv02
      ```

   d. Log in to the gateway guest (web server) and start the HTTP server. Switch to the Cognos user ID (su - cognos). The commands to start the HTTP server are:

      ```
      /opt/IBM/HTTPServer/bin/adminctl start
      /opt/IBM/HTTPServer/bin/apachectl -k start
      ```

4.2 Shutdown procedure for IBM Smart Analytics System 9600 components

The order in which the IBM Smart Analytics System 9600 components should be stopped is:

1. Shut down the Cognos components:
   a. Log in to the gateway guest (web server) and stop the HTTP server. Switch to the cognos user ID (su - cognos). The commands to stop the HTTP server are:

      ```
      /opt/IBM/HTTPServer/bin/apachectl -k stop
      /opt/IBM/HTTPServer/bin/adminctl stop
      ```

   b. Log in to the report server guest and stop both report servers. Switch to the Cognos user ID (su - cognos). The commands to stop the report servers are:

      ```
      /opt/IBM/WebSphere/AppServer/bin/stopServer.sh server1
      -profileName AppSrv01 -username wasadmin -password xxxxxxxxx
      /opt/IBM/WebSphere/AppServer/bin/stopServer.sh server1
      -profileName AppSrv02 -username wasadmin -password  xxxxxxxxx
      ```

      The first report server is AppSrv01 and the second is AppSrv02.

   c. Log in to the content manager guest and stop the content manager. Switch to the Cognos user ID (su - cognos). The command to stop the content manager is:

      ```
      /opt/IBM/WebSphere/AppServer/bin/stopServer.sh server1 -username
      wasadmin -password  xxxxxxxxx
      ```

2. InfoSphere components
   – Using the InfoSphere Warehouse Administration console from a web browser:
      i. Stop the XMLA interface.
      ii. Stop the cube.
      iii. Stop the cube server.
      iv. Ensure that there are no control flows running.
   – Stop the InfoSphere Warehouse server. The command to do this is:

      ```
      /opt/IBM/ISWarehouse/appServer/profiles/AppSrv01/bin/stopServer.s
      h  server1
      ```

4.3 Other administration tasks

In this section we provide basic remedies and steps to assist in managing your IBM Smart Analytics System 9600.

4.3.1 Stopping Cognos application when content store is unavailable

If the content store database becomes inaccessible during normal operations, the Cognos 8 BI Server application cannot submit new workloads for processing. In this situation, the Cognos application can be restarted by restarting the application server. If the content store becomes available within a short period of time, the Cognos application reconnects to the database and resumes workload processing. In the event that an active Content Manager is not designated after this short outage of the content store, it is necessary to stop and start the Cognos nodes to designate a new active Content Manager. See 4.2, “Shutdown procedure for IBM Smart Analytics System 9600 components” on page 53, and 4.1, “Startup procedure for IBM Smart Analytics System 9600 components” on page 52, for detailed steps to stop and start the Cognos nodes.

If the content store is inaccessible for longer than 10 minutes, stop the Cognos BI Server application and the application server. After a sustained outage during which the Cognos BI Server application cannot connect to the content store database, the Cognos application sometimes does not terminate when you issue the **`stopServer.sh server1`** command. If this scenario occurs, you might need to terminate the process manually. To terminate the process, determine the process ID of the application server by issuing the **`ps -eaf | grep server1`** command and then the **`kill`** command.

4.3.2 Backup and restore tasks

See 7.5, “Backup and restore tasks” on page 85, for information about backing up the Cognos module and components.

5

InfoSphere Warehouse administrative tasks

Administrators can use the Web-based InfoSphere Warehouse Administration Console for administrative tasks such as SQL Warehousing or OLAP in InfoSphere Warehouse.

You can use the InfoSphere Warehouse Administration Console to deploy, run, or monitor data-warehouse applications. These data-warehouse applications contain specific executable processes. You can also use the InfoSphere Warehouse Administration Console to study deployment histories, execution statistics, or log files.

In this chapter we provide an overview of InfoSphere Warehouse and its relationship to the other components in the IBM Smart Analytics System 9600, its architecture, and an overview of tasked involved in designing a warehouse using Data Studio.

For information about starting and stopping InfoSphere Warehouse or any of its components, see Chapter 4, “Managing the IBM Smart Analytics System 9600 components” on page 51.

Because the InfoSphere Warehouse Administration Console is browser based, you will use it to manage the data warehouse applications that you have deployed. During the life cycle of an application, you might need to update

application properties and eventually remove the application. For more information about this, see:

http://publib.boulder.ibm.com/infocenter/db2luw/v9r5/index.jsp?topic=/com.ibm.help.etl.doc/administering/tadmmanageapps.html

5.1 InfoSphere Warehouse and the IBM Smart Analytics System 9600

The following InfoSphere Warehouse software components have been installed on multiple z/VM Linux on System z guests:

- *The InfoSphere Warehouse Administration Console* consumes few resources beyond the normal resource requirements for WebSphere Application Server. The primary resource required is the processor. The processor utilization for the administration console directly relates to the number of concurrent administration console users.
- *The Design Studio* provides a common design environment for creating physical data models, OLAP cubes, data mining models, SQL data flows and control flows, and Blox® Builder analytic applications. The Design Studio is built on the Eclipse workbench, which is a development environment that you can customize.
- *The SQL Warehousing Tool (SQW)* executes SQL Warehouse process flows that are deployed as part of SQL Warehouse applications. Process flows can be run ad hoc or using the WebSphere-based scheduler. In general, the application server merely manages the process flows and sends jobs to the database using the administration Linux on System z guest. For this reason, most of the resource requirements are pushed into the target database rather than residing on the application server Linux on System z guest. However, some operators do impact the application server in terms of the processor, memory, and disk. When these types of operators are used, the resources that they consume depend on how they are used. One group of operators are the unstructured text analysis operators from the unstructured information management architecture (UIMA). The operators for this component consume very little memory but can have significant processor utilization on the application server guest.

  You can use the InfoSphere Warehouse Administration Console to deploy, run, and monitor data warehouse applications, which contain specific executable processes. You can also use the console to study deployment histories, execution statistics, and log files. To access the InfoSphere Warehouse Administratin Console to to:

  `http://hostname:9080/ibm/warehouse/`

  Where *hostname* is the name of your WebSphere Administration Server.

  Most users do not need to access the WebSphere Administration Console directly to manage their warehousing applications. The WebSphere Application Server provides certain functions for console processes, and the console itself is a J2EE application that WebSphere runs. However, you can

manage your deployed data warehouse applications entirely through the InfoSphere Warehouse Administration Console.

For more information about using the InfoSphere Warehouse Administration Console to manage your warehousing applications, see:

http://publib.boulder.ibm.com/infocenter/db2luw/v9r5/index.jsp?topic=/com.ibm.help.etl.doc/administering/tadmmanageapps.html

- *Cubing Services* is designed to provide a multidimensional view of data stored in a relational database. With Cubing Services, you can create, edit, import, export, and deploy cube models over the relational warehouse schema. Cubing Services also provides optimization techniques to dramatically improve the performance of OLAP queries, a core component of data warehousing and analytics. You can install one or more cube servers on one or more application server guests, thus reducing resource requirements for memory and processors. Every cube server primarily consumes memory and processor resources rather than storage and I/O resources. The factors that will affect the resources consumed by a cube server include:
  - The number of concurrent users or clients
  - The complexity of the Multidimensional Expressions (MDX) queries
  - The size and number of cubes

Figure 5-1 shows the relationships between the components of the IBM Smart Analytics System 9600. As shown in Figure 5-1, InfoSphere Warehouse is used for two purposes. First, it is used to move data between the operational system and the data warehouse system. Second, it is used to cache cubes. These cubes are created from the data warehouse data and used by the Cognos application.

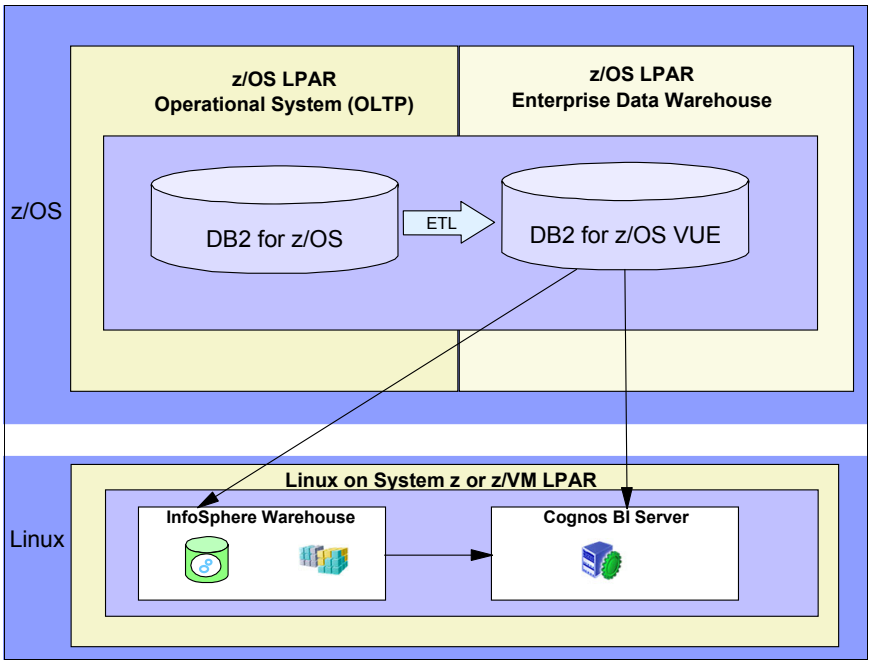

Figure 5-1 IBM Smart Analytics System 9600 component relationships

For data movement, InfoSphere Warehouse extracts, transforms, and loads from DB2 for z/OS in the OLTP LPAR to the DB2 for z/OS data warehouse LPAR. Cognos 8 BI can query the DB2 data warehouse directly or through the InfoSphere Warehouse cubes.

5.2 Architecture of InfoSphere Warehouse

InfoSphere Warehouse has a component-based architecture with client and server components. Figure 5-2 gives an overview of the architecture of InfoSphere Warehouse.

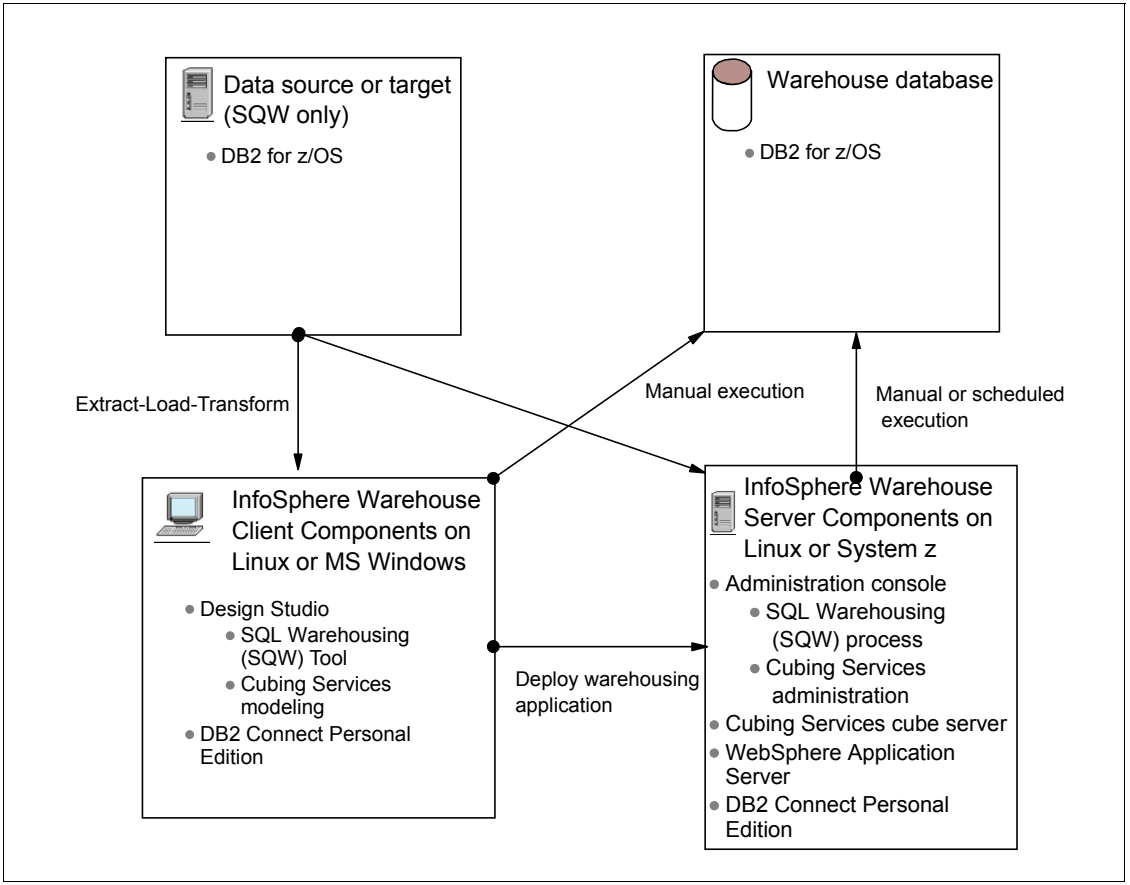

Figure 5-2 Logical component groups in the IBM InfoSphere Warehouse

IBM InfoSphere Warehouse runs on a Linux on System z guest.

- InfoSphere Warehouse server

  This includes:

  - InfoSphere Warehouse Cubing Services

    Cubing services provide OLAP access to data directly from InfoSphere Warehouse to the business intelligence tools like Cognos. There is also a cubing services tool for multidimensional modeling that runs on the InfoSphere Warehouse client. This is used to design OLAP metadata (cubes). The cube server is installed on the Linux partition. It executes as a stand-alone server and does not require WAS. The data, however, sits in DB2 tables.

    The cube server processes multidimensional queries expressed in the MultiDimensional eXpression (MDX) query language and produces multidimensional results. The cube server fetches data from DB2 through SQL queries as needed to respond to the MDX queries (if the data queried is readily available in the cubing services cache it fetches it from there, otherwise it gets the data from DB2).

  - Application server

    The application server is the WebSphere Application Server. The WebSphere Application Server is a Java-based web application server. It provides the access to runtime management capabilities from the InfoSphere Warehouse administration console, which allows the warehouse administrators to manage the runtime environment over the web using a web browser.

  - InfoSphere Warehouse Administration Console

    The Administration Console is a web application for warehouse administrators to deploy and manage warehouse applications, control flows, database resources, and system resources. It has a SQL warehousing runtime component to run and monitor data warehousing applications and view deployment histories and execution statistics. It has a cubing services component to manage cube servers, import and export cube models, explore cubes and cube models, and run the OLAP Metadata Optimization Advisor.

  - DB2 Connect™ Personal Edition

    DB2 Connect allows you to access and administer DB2 databases from remote workstations.

- Client

  In addition to the IBM Data Server Client, the InfoSphere Warehouse client runs on either an MS Windows® 32-bit or a Linux 32-bit operating system, and it has the following components:

  - Design Studio

    Design studio is an Eclipse-based integrated development environment (IDE) that facilitates design and development of data models, OLAP models, data flows, and control flows. Eclipse provides an extensible architecture based on the concept of plug-ins and extension points. This allows InfoSphere Warehouse Design Studio to take advantage of code reuse, integration, and many other development functions provided by Eclipse. Design studio has two components:

    - SQL Warehousing tool (SQW)

      The SQL Warehousing Tool is a graphical tool that generates SQL for warehouse maintenance and administration. The SQL Warehousing Tool automatically generates SQL that is based on visual operator flows that you model in Design Studio. The library of SQL operators covers the in-database data operations that are typically needed to move data between database tables and to populate analytical structures, such as multidimensional cubes.

      The basic function of SQW is to manage and move data into and around the data warehouse while transforming it for various purposes. SQW provides these services by making use of the power of the DB2 relational database engine and the SQL language, which classifies SQW in the Extract-Load-Transform (ELT) category of data movement and transformation tools. SQW also provides sequencing and flow control functions and functions to integrate non-database processing.

    - Cubing services modeling

      This component is integrated in the Design Studio's IDE. Cubing services modeling is an Eclipse-based feature for designing multidimensional models that can be consumed by the Cubing Services engine.

  - DB2 Connect Personal Edition

    DB2 Connect allows you to access and administer DB2 databases from remote workstations.

5.3 Designing Warehouse applications using Design Studio

In this section we provide an overview of designing a Data Warehouse using Design Studio. For more information about designing InfoSphere Warehouse applications using Design Studio, see *InfoSphere Warehouse: A Robust Infrastructure for Business Intelligence*, SG24-7813, which can be found at:

http://www.redbooks.ibm.com/abstracts/sg247813.html?Open

5.3.1 Data Warehouse/Business Intelligence solution design overview

There are four major steps involved in designing a Data Warehouse/Business Intelligence (DW/BI) solution:

1. Acquire the data.
2. Load the data into the warehouse.
3. Transform the data into information by building cubes for Online Analytical Processing (OLAP).
4. Present the information using BI tools.

Figure 5-3 shows this process.

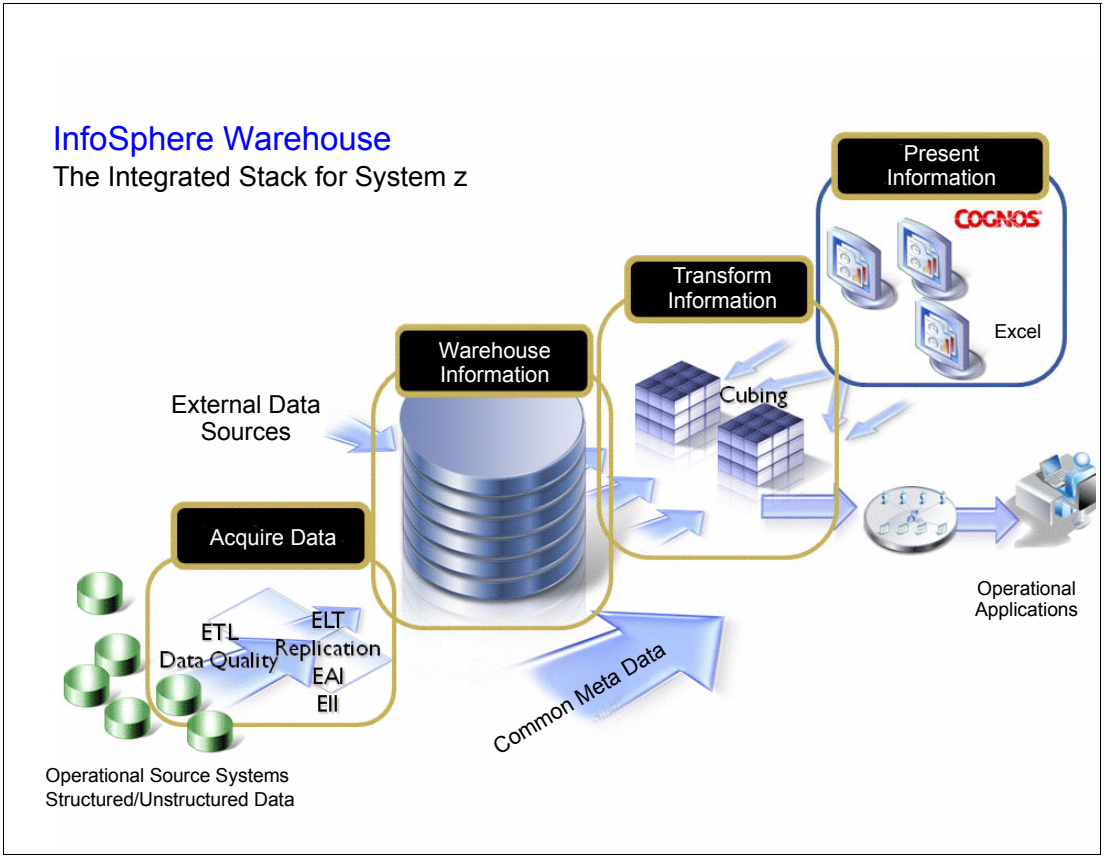

Figure 5-3 DW/BI solution design steps

The IBM Smart Analytics System 9600 uses the InfoSphere Warehouse Design Studio and SQL Warehousing tool (SQW) to acquire the data. DB2 for z/OS is used as the warehouse to store the data and then transform that data into information by using InfoSphere Warehouse Cubing Services and modeling tools and finally present the data using Cognos BI.

Many roles are involved in the DW/BI solution design, and there will be one or more people responsible for each task mentioned. Here is a list of the roles and responsibilities involved in a DW/BI solution design using the IBM Smart Analytics System 9600 system:

- Data architect

  The data architect models the database schemas that are needed to support the analytical solution. This person works with the business users and gets involved in the solution design at the conceptual design phase and will deliver the physical design to the warehouse administrator.

- Warehouse administrator

  The warehouse administrator performs tasks such as creating the tables and ETL or data movement processes or flows to populate the data structures. This person uses the SQW, SQW runtime, Admin Console, and cubing services to get his work done.

- OLAP developer

  The OLAP developer models and creates the OLAP metadata. This person uses the cubing services modeling tools to develop the OLAP models.

- BI developer

  The BI developer builds and delivers front-end analytical applications to the business owners in the corporation using Cognos Report Studio.

- BI administrator

  The BI administrator takes care of the availability of all the components of Cognos and their connections to the other components in the IBM Smart Analytics System 9600 system.

For more information about these roles, see 2.2, “Identifying the roles” on page 12.

5.3.2 The Design Studio workspace

The InfoSphere Warehouse Design Studio assists in the process of creating the physical data model.

In discussing the workspace, we mean both the Design Studio graphical user interface (GUI) and the actual directory that is used to store the Design Studio projects.

When you launch Design Studio, you are prompted to select a workspace (Figure 5-4). This is the default location where Design Studio stores projects. You can keep projects in different workspaces. After you select the workspace, you will see a welcome page. You can turn this welcome page off by selecting the “Use this as the default and do not ask again” check box. After closing the welcome page, the Design Studio workspace displays.

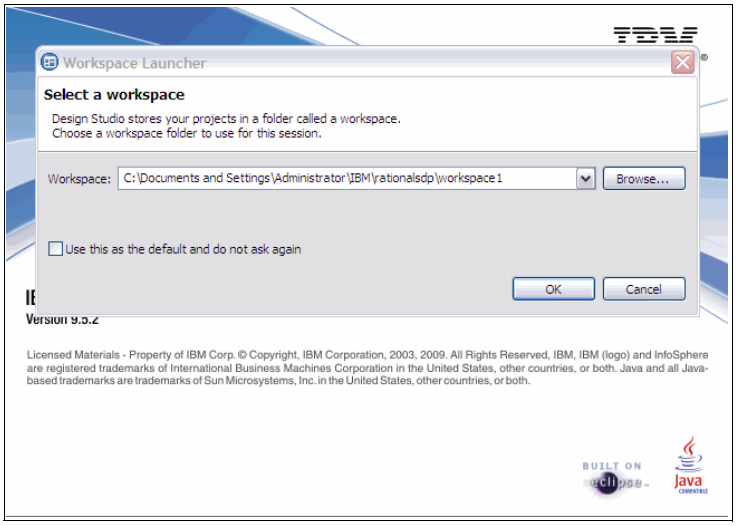

Figure 5-4 Selecting the Design Studio workspace

Figure 5-5 shows the Design Studio workspace.

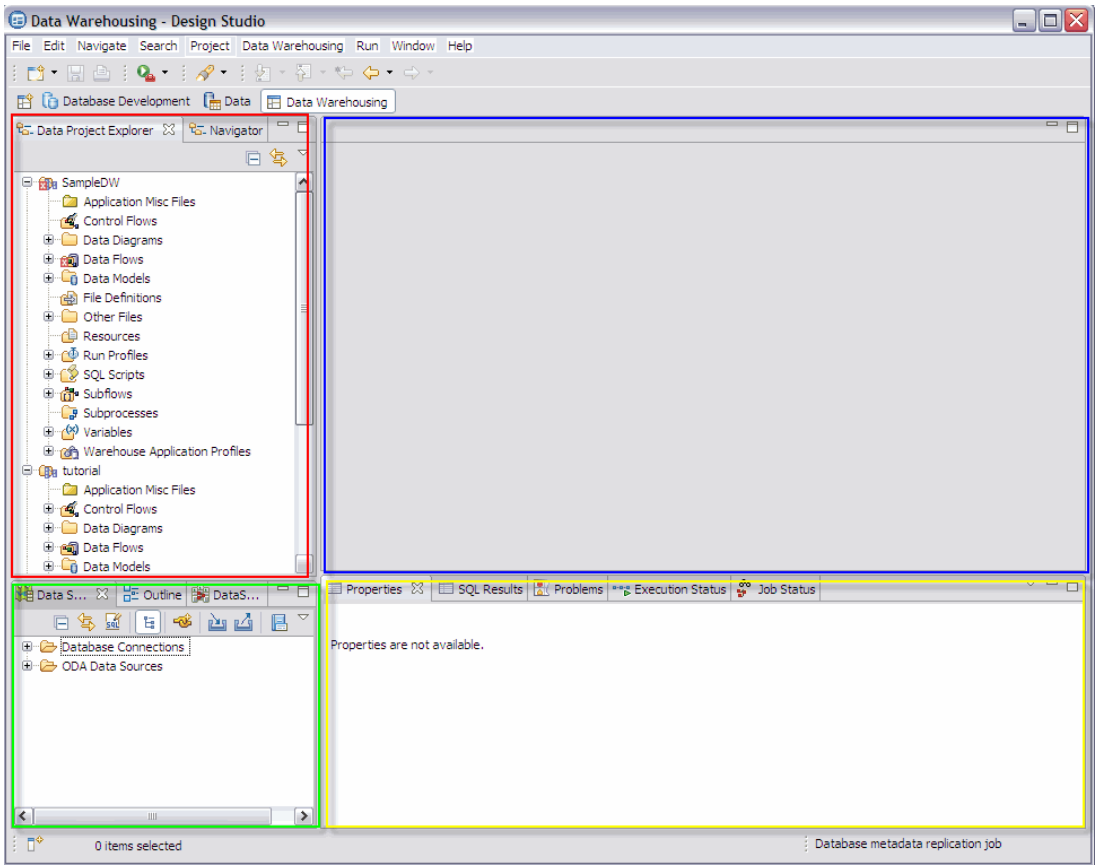

Figure 5-5 Design studio workspace

Figure 5-5 shows the Menu bar and the Toolbar across the top of the workspace. You will also see that the window is divided into four sections. The top left section is the Data Project Explorer. This is where the work is organized into projects. In Design Studio you will work with two types of projects:

- Data project

  A data project is used to develop the physical data models and the OLAP cube models.

- Data Warehouse project

  A Data Warehouse project is used to develop the data flows and control flows. It might reference one or more data projects to use the physical data model.

The bottom left section is the Data Source Explorer. This is where connections to the databases are defined. Database connections are required for a number of functions during the development process.

The top right section is the Editor pane. The type of editor used is dependent on the type of object being edited (for example, a text editor for text files and a graphical editor for data flows, physical data models and control flows). Most of the time graphical editor will be used.

The bottom right section contains a number of tabs that will be used throughout the development process to define properties of objects, show the results of SQL queries, view problems in the workspace, and view execution status and results.

5.3.3 Next steps

The next steps are:

- Develop the physical data model.
- Create the data model.
- Deploy the data model.
- Maintain the accuracy of the model.

For more information about these steps, see *InfoSphere Warehouse: A Robust Infrastructure for Business Intelligence*, SG24-7813, which can be found at:

http://www.redbooks.ibm.com/abstracts/sg247813.html?Open

6

Cognos 8 Business Intelligence

IBM Cognos 8 Business Intelligence (BI) can leverage your existing entity relationship (ER) database or cubing services investment by providing the ability to access existing data structures using either ER or warehouse data as a Cognos 8 data source.

The Cognos 8 BI system configured for the IBM Smart Analytics System 9600 was designed and built with reliability and scalability as key considerations.

In this section, we first discuss the Cognos 8 BI architecture. We next provide an overview of the Cognos 8 module and specific details about Cognos 8 BI as it pertains to the IBM Smart Analytics 9600.

6.1 Cognos architecture

Cognos 8 uses a multi-tiered architecture that allows various components to be applied within a single application framework. In this section, we focus on the components involved in reporting using the business intelligence components of the Cognos 8 platform. The base structure (Figure 6-1) consists of a tiered architecture. The individual services of the Cognos 8 server run within an application server and can be distributed across multiple application server instances.

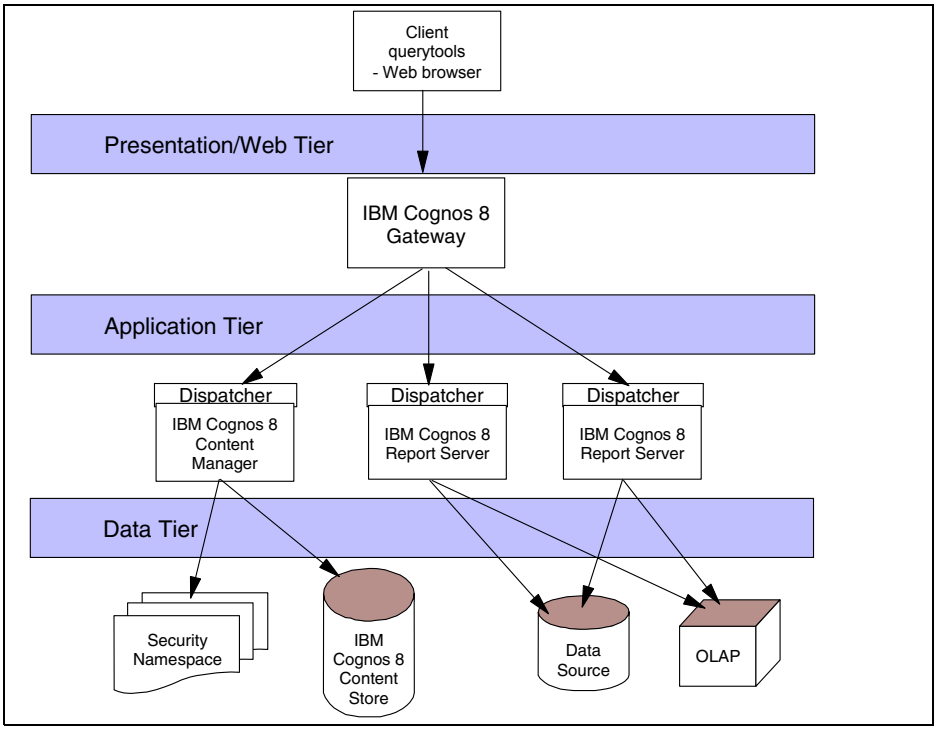

Figure 6-1 IBM Cognos 8 Business Intelligence architectural components

A browser interface at the presentation/web tier provides users with the ability to create reports and access published content from the IBM Cognos 8 Content Store database repository. This portal also allows for administration and configuration of the Cognos 8 server properties.

The IBM Cognos 8 Gateway component manages all Web communications in the IBM Cognos 8 Platform. The workload on the IBM Cognos 8 Gateway server is comparatively lightweight, but you can deploy multiple redundant gateways

along with an external HTTP load-balancing router to meet availability or scalability requirements.

The IBM Cognos 8 dispatcher performs the load balancing of requests at the application tier. The dispatcher component is a lightweight Java servlet that manages and provides communication between the application services. At startup, each IBM Cognos 8 Dispatcher registers locally available services with the IBM Cognos 8 Content Manager. During the normal operation of IBM Cognos 8 BI services, requests are load balanced across all available services by using a configurable, weighted round-robin algorithm to distribute requests. You can tune the performance of the IBM Cognos 8 Platform by defining how Dispatcher handles requests and manages services. The IBM Smart Analytics System 9600 will come with the normal configuration of two report server processes per allocated processor and 8 - 10 threads per process (either three low plus one high affinity thread or four low plus one high). Threads within the IBM Cognos 8 Platform are managed by the type of traffic that they handle (referred to as high and low affinity, where affinity relates to the report service process that handled the original user request when multiple interactions need to occur to satisfy the request_. High-affinity connections are used to process absolute and high-affinity requests from the report services, whereas low-affinity connections are used to process low-affinity requests. A high-affinity request is a transaction that can benefit from a previously processed request. It can be processed on any service, but resource consumption is minimized if the request is routed back to the report service process that was used to execute the original process. A low-affinity request will operate just as efficiently on any service.You can manage the number of threads per IBM Cognos 8 BI reporting service process through the IBM Cognos 8 Platform Administration Console by setting the number of high-affinity and low-affinity connections.

For more details, see:

http://publib.boulder.ibm.com/infocenter/c8bi/v8r4m0/index.jsp?topic=/com.ibm.swg.im.cognos.crn_arch.8.4.0.doc/crn_arch.html

The IBM Cognos 8 Report server report service (also known as the query service) is responsible for application-tier processing. These services are often referred to as the BIBus processes, as they are the services of the BI Business server.

The request flow for report execution is:

1. The user clicks a report to run it, and the request goes through the gateway and the dispatcher to the presentation service.
2. The presentation service sends the request to the report service.
3. The report service requests the report and metadata from the content manager.

4. The Content Manager sends the report XML specifications and metadata to the report service. Content manager refetches metadata only when IBM Cognos is stopped and restarted or the model is updated and republished.
5. The report service returns one of the following results to the presentation service:
   - An error page
   - A not ready page
   - A page of an HTML report
6. The presentation service sends one of the following results through the dispatcher and gateway to the browser:
   - An error page
   - A wait or cancel page
   - A page of a completed HTML report in the report viewer interface

The IBM Cognos 8 Content Manager is the service that manages the storage of customer application data, including security, configuration data, models, metrics, report specifications, and report output. It is needed to publish packages, retrieve or store report specifications, manage scheduling information, and manage the Cognos namespace. The Content Manager maintains information in a relational database that is referred to as the content store database. The IBM Smart Analytics System 9600 will come with a minimum of one Content Manager service (required for each IBM Cognos 8 Platform implementation). Content Manager performance can benefit from the availability of high-speed RAM resources and will have one processor for every four processors allocated for report server processing.

6.2 Adding authentication credentials to a data source

When DB2 databases or Cubing Services cube servers are defined as reporting data sources in the Cognos module, authentication credentials are not supplied for these data sources when the system is initially configured. When reports are executed using these reporting data sources, valid user credentials need to be provided before the report is executed. However, authentication credentials can be added to a reporting data source in the Cognos BI Server application to suppress the request for user credentials instead. An authentication credential is added to a reporting data source by creating a Cognos signon and then associating it with the data source defined in Cognos. You can add multiple authentication credentials to a reporting data source by creating multiple signons and associating each one with the Cognos data source. More information about user IDs that are predefined for the IBM Smart Analytics System 9600 can be

found in Chapter 8, “Managing users of the IBM Smart Analytics System 9600” on page 99.

6.3 Accessing Cognos 8 BI components

To use Cognos 8 BI reporting and query functions, and to manage the Cognos 8 BI application, you will need to access the following hosted components:

- Cognos Connection: The web portal used to manage Cognos 8 resources and content. It provides the user interface to the Cognos 8 BI server application and is a single point of access to other Cognos components, such as the Cognos Administration portlet and the Cognos Viewer portlet that displays report output. It is hosted on the application server provided by the WebSphere Application Server and can be accessed by navigating to:

  `http://`*`cogsip`*`/cognos8`

  Where *cogsip* is the gateway service IP address.

- Cognos Content Manager status page: A web page hosted on the application server provided by WebSphere Application Server that shows the status of the Cognos Content Manager. The status page for the Content Manager hosted on a particular server can be accessed by navigating to:

  `http://`*`cognos001`*`:9081/p2pd/servlet`

  Where *cognos001* represents the host name of the Cognos server.

  The Content Manager status page displays the following information:

  – Cognos build number
  – Start time
  – Current time
  – Content Manager state

  An active Content Manager is displayed as “Running” and a standby Content manager displays a “Running as standby” state.

- Cognos 8 BI Administration console: Hosted on the WebSphere Application Server, the administration console is deployed as an Enterprise Archive (EAR) file and can be accessed by navigating to:

  `http://`*`cognos001`*`:9061/ibm/console`

  Where *cognos001* represents the host name of the application server.

  When the application server is started or stopped on the application server, it automatically starts or stops the Cognos BI Server application.

Note: When the Cognos application is stopped with the application server, all running workloads on that server are disrupted and need to be resubmitted for processing.

6.4 Cognos 8 BI performance configuration settings

The IBM Smart Analytics System 9600 Cognos 8 BI comes preconfigured for performance. These settings are based on the following:

- IBM Cognos recommends setting the following for the maximum number of processes for the report service for peak period to two times the number of cores or CPUs. For example, if your environment had two CPUs, the equation would be 2 * 2 = 4 processes.
- IBM Cognos recommends setting the maximum number of processes for the batch service for peak period to two times the number of cores or CPUs. For example, if your environment had two CPUs, the equation would be 2 * 2 = 4 processes.

The `CQEConfig.xml` file has also been modified for performance on the report server. In this file, timeout has been changed to 300 from 900 and PoolSize has been changed to 75 from 20.

The IBM Cognos Configuration tool comes with generic configuration entries. These entries have been changed accordingly to effectively communicate with other servers. For example, the correct hostnames have been configured for the Content Manager, gateway, application, DB2, security, and so on.

6.5 Accessing IBM Cognos 8 BI Metadata

Cognos 8 supports a direct connection to the data sources for Cube Services and can store data structures for traditional database access. In the case of cube services, this means that the metadata for the published package can be obtained directly from the Cube Server at run time instead of requiring full metadata import into Framework Manager. In the case of traditional database access, the meta model is stored and accessed in the Cognos Server. In either case, you will still need to publish a package from Framework Manager to enable access to the different data sources, but in the case of Cube Services, there will not be any required changes to the cube properties.

In either a traditional data source or a Cube Service scenario, you will need to define a data source connection to the source within Cognos 8 and import them into Framework Manager. In this instance the cube is simply a stub object that is used to reference the cube from the Cubing Services Cube Server. The full set of metadata for the dimensions, hierarchies, and levels remains within Cubing Services.

You can get a detailed list and definition of all supported data sources and software environments from the Cognos 8 BI Software Environments website at:

http://www.ibm.com/support/docview.wss?rs=3442&uid=swg27014432

For more information about how to create and secure data sources, see section 2.2.1 in *Leveraging IBM Cognos 8 BI for Linux on IBM System z*, SG24-7812.

6.6 Application build process overview

The following steps are followed for a typical Cognos build process:

1. Locate and prepare data sources and models.

   IBM Cognos 8 can report from a wide variety of data sources, both relational and dimensional. Database connections are created in the web administration interface, and are used for modeling, authoring, and running the application. To use data for authoring and viewing, the business intelligence studios need a subset of a model of the metadata (called a package). The metadata might need extensive modeling in Framework Manager.

2. Build and publish the content.

   Reports, scorecards, analysis, dashboards, and more are created in the business intelligence studios of IBM Cognos 8. Which studio you use depends on the content, lifespan, and audience of the report, and whether the data is modeled dimensionally or relationally. For example, self-service reporting and analysis are done through Query Studio and Analysis Studio, and scheduled reports are created in Report Studio. Report Studio reports and scorecards are usually prepared for a wider audience, published to IBM Cognos Connection or another portal, and scheduled there for bursting, distribution, and so on. You can also use Report Studio to prepare templates for self-service reporting.

3. Deliver and view the information.

   You deliver content from the IBM Cognos portal or other supported portals, and view information that has been saved to portals or delivered by other

mechanisms. You can also run reports, analyses, scorecards, and more from within the business intelligence studio in which they were created.

For information about tuning and performance, see the *IBM Cognos 8 Administration and Security Guide:*

http://www.ibm.com/software/data/support/cognos_crc.html

6.7 Topology overview with install considerations

Figure 6-2 displays the Cognos technical components for Cognos BI and the platform on which they are installed.

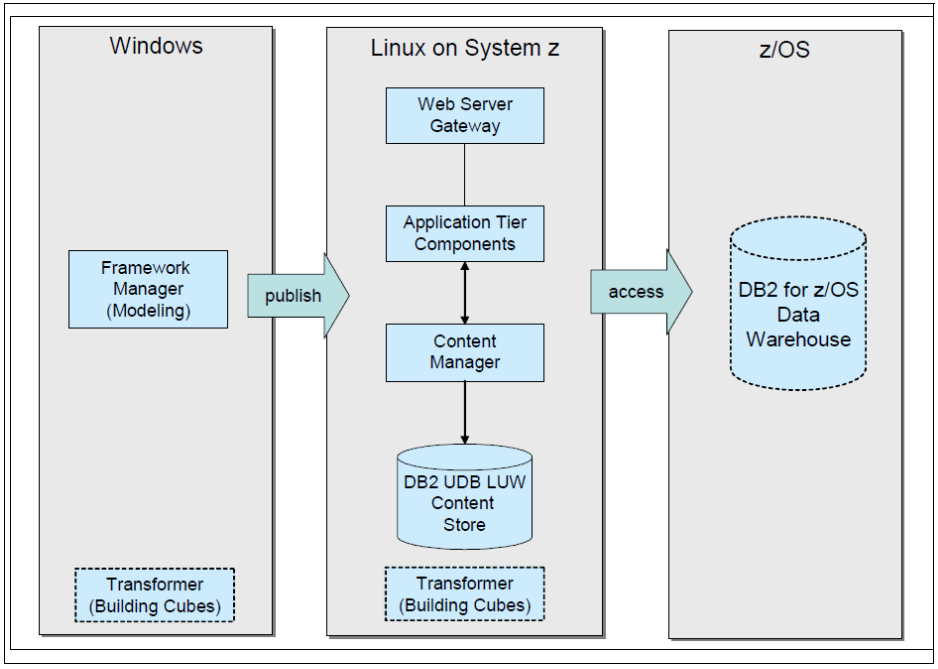

Figure 6-2 Software component view for Cognos installation

The Framework Manager is installed in an MS Windows environment. The Cognos BI modeling tool to define packages that are then published to the Cognos server runs on Linux on System z. The server components include a web server gateway, which accepts HTTP requests from (web) clients, and an application tier layer that processes the requests. The Content Manager accesses the Content Store database, which maintains metadata, such as the published packages and stored reports.

7

System z and the IBM Smart Analytics System 9600

This chapter discusses resource management and performance monitoring for System z with the IBM Smart Analytics System9600 installed.

We discuss some of the System z components and how they interact with the IBM Smart Analytics System 9600 environment.

Resource management and performance monitoring of IT workloads are keys to providing satisfactory service to the business community. This is especially true when transactional and data warehouse workloads are housed within the same System z hardware as their processing needs and the associated resource requirements likely have very different execution characteristics and service delivery objectives. With System z, the infrastructure exists to monitor, manage, and report activities or histories of activities at a very granular level with facilities such as z/OS SMF and additional reporting tools such as RMF™.

In this chapter, the following topics are discussed:

- IBM Smart Analytics System 9600 WLM Policies
- Managing users
- DFSMS
- High availability and backup considerations
- Disaster recovery for System z
- Monitoring on System z

- Capacity management. on System z
- Managing Linux on System z

The IBM Smart Analytics System 9600 installation contains a set of components that require monitoring and tuning for the System z environment.

7.1 IBM Smart Analytics System 9600 WLM Policies

The Workload Manager (WLM) is part of z/OS. Each z/OS system has its own WLM policy, which is the way to classify workloads or tasks. You allocate goals in a business-oriented manager rather than allocation resources to a task. To avoid queries monopolizing your system, it is very important to tune your WLM policy to control parallel query processor consumption. In this section we discuss the customization of the WLM configuration for System z.

For additional information about customization steps for z/OS system performance:

- See the WLM IBM Smart Analytics System service definition high-level overview found in Chapter 11 of *Co-locating Transactional and Data Warehouse Workloads on System z*, SG24-7726.
- Read “Resource management of DB2 data warehouse queries” in section 9.4 of *Co-locating Transactional and Data Warehouse Workloads on System z*, SG24-7726.
- You might need to make additional updates to the WLM Classification rules. Specifically, ensure that you validate the DDF, batch, and TSO classification rules and associated classification groups. Information about how to do this can be found in Chapter 11 of *Co-locating Transactional and Data Warehouse Workloads on System z*, SG24-7726.
- Review service class goals and adjust as necessary for your installation. How to determine your service class goals can be found in Chapter 11 of *Co-locating Transactional and Data Warehouse Workloads on System z*, SG24-7726.
- Once running IBM Smart Analytics System workloads, review performance and tune WLM definitions as necessary.
- Examine the RMF Workload Activity Report to ensure that your work is getting classified as anticipated. The UNCLASS service class and RDDFUNC report class should not have any service consumption.

Service classes

The WLM administrative application was used to define the service classes that z/OS will manage for the IBM Smart Analytics System 9600. These service classes are associated with performance objectives. When a WLM-established stored procedure call originates locally, it inherits the performance objective of the caller, such as TSO or CICS.

If classification rules do not exist to classify some or all of your DDF transactions into service classes, those unclassified transactions are assigned to the default

service class, SYSOTHER, which has no performance goal and is even lower in importance than a service class with a discretionary goal.

For more information about sample WLM service definitions, see Appendix D in the IBM Redbooks publication *Co-locating Transactional and Data Warehouse Workloads on System z*, SG24-7726.

The following service classes are required for DB2 DDF and have already been updated through the WLM panels for the IBM Smart Analytics System 9600:

- Service class DDFHI (Table 7-1): DDF high-priority users and applications. Multiperiod mix of percentile response time and velocity goals, providing higher priority, with more consistent response times, for shorter consumption work (for example, metadata access, operational BI, and trivial reports).

  CPU Critical flag: NO

  Table 7-1 DDFHI

Period	Duration	Importance	Goal description
1	25000	2	90% complete within 00:00:03.000
2	100,000	2	80% complete within 00:00:15.000
3	1,000,000	3	Execution velocity of 30
4		4	Execution velocity of 10

- Service class DDFREFSH (Table 7-2): DDF refresh high-importance daily batch refreshes as well as other intra-day refreshes. This is used during the time of high-importance daily refresh runs.

  CPU Critical flag: NO.

  Table 7-2 DDFREFSH

Period	Duration	Importance	Goal description
1		4	Execution velocity of 10

- Service class DDFSCHED (Table 7-3): DDF scheduled reports.

  CPU Critical flag: NO

Table 7-3 DDFSCHED

Period	Duration	Importance	Goal description
1	2,000,000	4	Execution velocity of 30
		5	Execution velocity of 10

- Service class DDFSTD (Table 7-4): DDF STD high-importance query service class.

  CPU Critical flag: NO

Table 7-4 DDFSTD

Period	Duration	Importance	Goal description
1	25,000	3	90% complete within 00:00:03.000
2	500,000	4	Execution velocity of 10
3			Discretionary

Subsystem type distributed data facility (DDF) work

The data warehouse distributed relational database architecture (DRDA) query key classification rules implemented for the DRDA query service and report classes are outlined in this section. If you do not classify your DDF transactions into service classes, they will be assigned to the default class, SYSOTHER, which is set to a priority even lower than a service class with a discretionary goal.

Classification rules for the data warehouse query workload service classes that have been set up for the IBM Smart Analytics System 9600 are:

- Default service class DDFDEF
- Default report class RDDFDEF

You can classify DDF threads by, among other things, authorization ID. The classification criteria set for the IBM Smart Analytics System 9600 is UI, which means that DDF threads will be classified by the user ID assigned to the transaction (Table 7-5 on page 82). The qualifier name, shown in the table, indicates requests coming from user IDs that start with the letters shown. For example, all user IDs that start with the letter H are assigned to service class DDFHI and report class RDDFHI.

The column labeled as # indicates the level of filtering. The IBM Smart Analytics System 9600 is set with only a first level of filtering. However, if you wanted to

make filtering more granular (for example, all requests coming from the user ID HE* will have a different service and report class than those that start with the user ID of H only), you would assign a second level of filtering.

Table 7-5 DDF classification

#	Qualifier type	Qualifier name	Starting position	Service class	Report class
1	UI	H*		DDFHI	RDDFHI
1	UI	M*		DDFSTD	RDDFSTD
1	UI	N*		DDFLO	RDDFLOW

7.2 Managing users

The IBM Smart Analytics System 9600 will come with the 30 pre-defined LDAP user IDs for Cognos 8 BI:

- hcognos01-10: Critical knowledge workers
- mcgonos01-10: Intermediate knowledge workers
- ncognos01-10: Novice users

These have been pre-defined in order to position the customer for user differentiation right from the start. You may customize this to your naming conventions from z/OS by going into UNIX System Services (USS) and modifying the *coguser.ldif* file found in the `/tmp` directory. In this file, you will find the organizational unit named *DWHzUsers* and the 30 pre-defined user IDs.

To add or modify users:

1. Enter:

   ```
   export PATH=/usr/lpp/ldap/sbin:$PATH
   ```

2. Enter:

   ```
   export NLSPATH=/usr/lpp/ldap/lib/nls/msg/%L/%N:$NLSPATH
   ```

3. Enter:

   ```
   export LANG=En_US.IBM-1047
   ```

4. Enter:

   ```
   ldapmodify -h 129.40.178.5 -D cn=admin -w secret -a -f /tmp/cogus
   ```

Pre-defined for the customer are InfoSphere Warehouse for System z (ISWz) DB2 connections, using user ID ISWZADM. This user ID should only be used for

access to the ISWz metadata repository. For ISWz access of z/OS DB2 data warehouses, set up additional connections—at least one for SQW work and one for Cubing services. These user IDs would go into the WLM service definition. You would want one for SQW refresh data flows, and another for cubing services. You can define whatever use IDs you prefer, although we suggest that you update the DDF classification rules to match.

Pre-defined for the IBM Smart Analytics System 9600 are two IBM Cognos user IDs, one for the IBM Cognos Content Store and one for the sample IBM Smart Analytics System 9600 z/OS DB2 data warehouse database. Both connections utilize COGZADM as the user ID. We suggest continuing to utilize COGZADM as the ID for the content store connection, but create one or more additional connection ids for Cognos DB2 data warehouse access.You can define whatever use IDs that you prefer, although we suggest that you update the DDF classification rules to match.

The DDF subsystem encompasses all the DB2 work that was initiated remotely via DRDA. Any local attached DB2 database processing will be included within the associated service class of the local application (for example, batch, TSO (QMF™), Local WebSphere w/RRS, oMVS, Local CICS-attach, and so on).

7.3 DFSMS

Most of the DB2 for z/OS data sets can be managed with DFSMS storage pools, thus reducing the workload of the DB2 database administrators (DBAs) and storage administrators. Even the most critical data, as defined with service level agreements (SLAs) can be managed by DFSMS with special attention.

With DFSMS, the user can distribute the DFSMS storage groups among storage servers with the purpose of optimizing access parallelism. Another purpose can be managing availability for disaster recovery planning. DFSMS automatically fills in these storage groups with data sets by applying policies that are defined in a set of predefined routines.

The IBM Smart Analytics System 9600 has initially been set up with the following SMS datasets:

- SYS1.DFSMS.SCDS
- SYS1.DFSMS.ACDS
- SYS1.DFSMS.COMMDS

DB2 storage and data classes have been set up as:

- DB2DATA - Application/User Data - PJD001-PJD010
  - GUARANTEED SPACE=YES
  - DATACLAS=DB2EXAD
- DB2SYSTM - DB2 Catalog/System - PJDSC1-PJDSC2
- DB2LOGS - LOGS/BSDS - PJDLG1-PJDLG2
- DB2WORK - DSNDB07 - PJDW01-PJDW10
- NONSMS - Libraries/tools -- PJDDLB

For more information about the enhancements and supported functions for DB2 and DFSMS, see *DB2 9 for z/OS and Storage Management*, SG24-7823, and *IBM System Storage DS8000: Architecture and Implementation*, SG24-6786.

7.4 High-availability and backup considerations

In its most basic configuration, the IBM Smart Analytics System 9600 is built on a highly available platform. The system is delivered with four z/VM guests running Linux. This is the foundation for continuing to leverage the underlying IFL capacity. Multiple z/VM guests can be configured for even further HA.

All Cognos components might have multiple instances through scaling, except Content Manager. As mentioned in 6.3, “Accessing Cognos 8 BI components” on page 73, you could have another Content Manager on standby or even more than one on standby. Of course the standby Content Manager requires minimal disk resources and most of the time they do nothing. In a z/VM virtual environment, the real CPU and memory resources associated with the standby Content Manager will be available to other virtual machines until the standby becomes active.

Having Cognos BI running on another Linux on System z guest on the same System z offers both advantages and disadvantages. The CPU and memory resources for Linux are virtual when running Linux under z/VM, so the associated physical resources will not be needed until the Linux guest becomes active. However, having a second Linux guest up and running the entire time requires additional disks beyond the basic level, and it does not protect data from power failures.

Another issue is data access. Cognos BI might be up and running, but it is still unable to get to your data. A data replication solution might solve this kind of problem and also certain performance problems, as all users would access the data within their geography.

Having one default global Content Manager allows the architecture to appear as a single instance while in fact Cognos services and data sources are geographically widely distributed. Standby Content Managers in all locations ensure high availability. Advanced routing might prevent processing of requests that are remote from the user location.

For further details about high availability, see *DB2 UDB for z/OS: Design Guidelines for High Performance and Availability*, SG24-7134, which covers a large variety of recommendations to increase the availability of data.

For further discussion of DB2 for z/OS and data warehouse transactional processing, see *Co-locating Transactional and Data Warehouse Workloads on System z*, SG24-7726.

7.5 Backup and restore tasks

This section gives an overview of the backup and restore tasks necessary for the IBM System Analytics System 9600. In backing up the data in the DB2 for z/OS data warehouse, use the standard DB2 backup and restore utilities.

7.5.1 Backing up the DB2 catalog and directories

The DB2 system catalog and directories have been backed up using an image copy before the IBM Smart Analytics System 9600 was turned over for use. It is a good practice to periodically run a backup such as this. Example 7-1 is a sample DB2 v9 system image copy job that copies the system catalog and directories and writes the outputs to DASD. You can copy and paste the following text into a member of your own dataset. The sample exists in the DB2I.V9.SDSNSAMP installation dataset.

You will need to change the symbolic &DMMDDYY to the date (for example, 051810).

Note: The SMS STORCLAS named SMSUCLAS must exist on the system

Example 7-1 Sample DB2 v9 system image copy job

```
//ICDB2SYS  JOB ,,MSGLEVEL=1,CLASS=A,MSGCLASS=H,
//    REGION=0M,NOTIFY=&SYSUID
//*************************************************************
//JOBLIB   DD  DISP=SHR,DSN=DB2I.SDSNLOAD
//SYSUTILX  EXEC PGM=DSNUTILB,PARM='DB2I,ICDB2SYS'
```

```
//SYSCOPYX DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB01.SYSUTILX,STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(10,10),RLSE)
//SYSPRINT DD  SYSOUT=*
//SYSUDUMP DD  SYSOUT=T
//SYSIN    DD  *
   COPY TABLESPACE DSNDB01.SYSUTILX  COPYDDN SYSCOPYX
//*************************************************************
//STEP2   EXEC PGM=DSNUTILB,PARM='DB2I,ICDB2SYS',COND=(4,LT)
//SYSPRINT DD  SYSOUT=*
//SYSUDUMP DD  SYSOUT=T
//SYSCOPY1 DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB01.DBD01,
// STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(10,10),RLSE)
//SYSCOPY2 DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB01.SCT02,
// STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(40,10),RLSE)
//SYSCOPY3 DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB01.SPT01,
// STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(450,10),RLSE)
//SYSCOPY4 DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB06.SYSDBASE,
// STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(30,10),RLSE)
//SYSCOPY5 DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB06.SYSDBAUT,
// STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(10,10),RLSE)
//SYSCOPY6 DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB06.SYSGPAUT,
// STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(10,10),RLSE)
//SYSCOPY7 DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB06.SYSGROUP,
// STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(10,10),RLSE)
//SYSCOPY8 DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB06.SYSPLAN,
// STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(40,10),RLSE)
//SYSCOPY9 DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB06.SYSPKAGE,
// STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(150,10),RLSE)
//SYSCOPYA DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB06.SYSUSER,
// STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(10,10),RLSE)
//SYSCOPYB DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB06.SYSSTR,
// STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(10,10),RLSE)
//SYSCOPYC DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB06.SYSVIEWS,
// STORCLAS=SMSUCLAS,
```

```
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(10,10),RLSE)
//SYSCOPYD DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB06.SYSSTATS,
// STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(10,10),RLSE)
//SYSCOPYE DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB06.SYSDDF,
// STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(50,10),RLSE)
//SYSCOPYF DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB06.SYSOBJ,
// STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(10,10),RLSE)
//SYSCOPYG DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB06.SYSSEQ,
// STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(10,10),RLSE)
//SYSCOPYH DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB06.SYSSEQ2,
// STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(10,10),RLSE)
//SYSCOPYI DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB06.SYSHIST,
// STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(10,10),RLSE)
//SYSCOPYJ DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB06.SYSGRTNS,
// STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(10,10),RLSE)
//SYSCOPYK DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB06.SYSJAVA,
// STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(10,10),RLSE)
//SYSCOPYL DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB06.SYSJAUXA,
// STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(10,10),RLSE)
//SYSCOPYM DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB06.SYSJAUXB,
// STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(10,10),RLSE)
//SYSCOPYN DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB06.SYSALTER,
// STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(10,10),RLSE)
//SYSCOPYO DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB06.SYSEBCDC,
// STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(10,10),RLSE)
//SYSCOPYP DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB06.SYSXML,
// STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(10,10),RLSE)
//SYSCOPYQ DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB06.SYSTARG,
// STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(10,10),RLSE)
//SYSCOPYR DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB06.SYSPLUXA,
// STORCLAS=SMSUCLAS,
```

```
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(10,10),RLSE)
//SYSCOPYS DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB06.SYSROLES,
// STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(10,10),RLSE)
//SYSCOPYT DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB06.SYSCONTX,
// STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(10,10),RLSE)
//SYSCOPYU DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB06.SYSRTSTS,
// STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(10,10),RLSE)
//SYSCOPYV DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB01.SYSLGRNX,
// STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(10,10),RLSE)
//SYSCOPYZ DD  DSN=DB2I.IC1.&DMMDDYY.DSNDB06.SYSCOPY,
// STORCLAS=SMSUCLAS,
// DISP=(NEW,CATLG,DELETE),UNIT=3390,SPACE=(CYL,(10,10),RLSE)
//SYSIN    DD  *
   COPY TABLESPACE DSNDB01.DBD01     COPYDDN SYSCOPY1
   COPY TABLESPACE DSNDB01.SCT02     COPYDDN SYSCOPY2
   COPY TABLESPACE DSNDB01.SPT01     COPYDDN SYSCOPY3
   COPY TABLESPACE DSNDB06.SYSDBASE  COPYDDN SYSCOPY4
   COPY TABLESPACE DSNDB06.SYSDBAUT  COPYDDN SYSCOPY5
   COPY TABLESPACE DSNDB06.SYSGPAUT  COPYDDN SYSCOPY6
   COPY TABLESPACE DSNDB06.SYSGROUP  COPYDDN SYSCOPY7
   COPY TABLESPACE DSNDB06.SYSPLAN   COPYDDN SYSCOPY8
   COPY TABLESPACE DSNDB06.SYSPKAGE  COPYDDN SYSCOPY9
   COPY TABLESPACE DSNDB06.SYSUSER   COPYDDN SYSCOPYA
   COPY TABLESPACE DSNDB06.SYSSTR    COPYDDN SYSCOPYB
   COPY TABLESPACE DSNDB06.SYSVIEWS  COPYDDN SYSCOPYC
   COPY TABLESPACE DSNDB06.SYSSTATS  COPYDDN SYSCOPYD
   COPY TABLESPACE DSNDB06.SYSDDF    COPYDDN SYSCOPYE
   COPY TABLESPACE DSNDB06.SYSOBJ    COPYDDN SYSCOPYF
   COPY TABLESPACE DSNDB06.SYSSEQ    COPYDDN SYSCOPYG
   COPY TABLESPACE DSNDB06.SYSSEQ2   COPYDDN SYSCOPYH
   COPY TABLESPACE DSNDB06.SYSHIST   COPYDDN SYSCOPYI
   COPY TABLESPACE DSNDB06.SYSGRTNS  COPYDDN SYSCOPYJ
   COPY TABLESPACE DSNDB06.SYSJAVA   COPYDDN SYSCOPYK
   COPY TABLESPACE DSNDB06.SYSJAUXA  COPYDDN SYSCOPYL
   COPY TABLESPACE DSNDB06.SYSJAUXB  COPYDDN SYSCOPYM
   COPY TABLESPACE DSNDB06.SYSALTER COPYDDN SYSCOPYN
   COPY TABLESPACE DSNDB06.SYSEBCDC COPYDDN SYSCOPYO
   COPY TABLESPACE DSNDB06.SYSXML   COPYDDN SYSCOPYP
   COPY TABLESPACE DSNDB06.SYSTARG  COPYDDN SYSCOPYQ
   COPY TABLESPACE DSNDB06.SYSPLUXA COPYDDN SYSCOPYR
   COPY TABLESPACE DSNDB06.SYSROLES COPYDDN SYSCOPYS
```

```
   COPY TABLESPACE DSNDB06.SYSCONTX COPYDDN SYSCOPYT
   COPY TABLESPACE DSNDB06.SYSRTSTS COPYDDN SYSCOPYU
   COPY TABLESPACE DSNDB01.SYSLGRNX  COPYDDN SYSCOPYV
   COPY TABLESPACE DSNDB06.SYSCOPY   COPYDDN SYSCOPYZ
//
```

7.5.2 Backing up Cognos 8 BI

You will need to regularly back up the IBM Cognos 8 BI data and configuration settings and your Framework Manager projects and models. We suggest that you perform these backups offline during a scheduled outage because it requires stopping the application server. Stopping the application server stops the Cognos application, which prevents it from initiating communications to the content store. Stopping the application server also disrupts running workloads and requires the workloads to be resubmitted for processing.

The Cognos content store is a relational database that stores data that the Cognos application needs to operate, such as report specifications and scheduling, connection information, and information about the external namespace and the Cognos namespace. Periodic backups of the content store capture changes to the operational data (for example, when settings are modified through Cognos Connection or when new applications are added). To back up the content store, you back up the DB2 database named COGZDB using the standard DB2 backup utilities. To restore the content store, you use the DB2 restore utilities. For more information about backup and restore utilities for DB2, see *DB2 9 for z/OS: Using the Utilities Suite*, SG24-6289.

The Cognos configuration backup includes Cognos configuration settings and the `/usr/IBM/cognos/configuration` directory. Periodic backups of the Cognos configuration capture changes made through the cogconfig.sh utility, such as changing the Cognos user's password.

To back up the Cognos configuration using an offline procedure, perform the following steps:

1. Stop the application server:
   a. Log in to the Content Manager guest and stop the Content Manager.
   b. Switch to the Cognos user ID (su - cognos). The command to stop the content manager is:

      ```
      /opt/IBM/WebSphere/AppServer/bin/stopServer.sh server1 -username
      wasadmin -password xxxxxxxxx
      ```

2. Copy the `c8_location/configuration` directory to the backup location. This directory contains the configuration settings. If you must ever restore the configuration settings, you can copy the backed-up directory to the correct location.

To back up framework manager projects and models, copy the Framework Manager project directory and its subdirectories to the backup location. By default, the projects and models are located in My Documents/My Projects.

If you must restore the Framework Manager projects and models, you can copy the backed-up directories to the correct location.

7.5.3 Backing up Linux on System z and important z/VM files

The IBM Smart Analytics System has been installed using the Novell SUSE Linux Enterprise Server 10 sp2.

Backups store incremental changes, such as changes to individual files, and you can do backups during system operations. Backups are part of an on demand file-level recovery system and you should do backups daily.

When backing up Linux on System z, you must ensure that your backup application also backs up the hard and soft links, as well as important files.

The first configuration file read when z/VM IPLs is the SYSTEM CONFIG file. DirMaint™ has been installed and enabled on the IBM Smart Analytics 9600 along with RACF and VM Performance Toolkit (these will be products 6VMDIR10, 6VMRAC10, and 6VMPTK10 in the SYSTEM CONFIG file). You will need to make regular backup copies of the following files from the MAINT user ID using the COPYFILE command:

- `SYSTEM CONFIG`
- `PROFILE EXEC`
- `PROFILE TCPIP`
- `PROFILE XEDIT`

- `USER DIRECT`
- `user withpass a1`
- `lindflt direct a`
- `linux2g protodir a1`
- `linux4g protodir a1`
- `lxisas1 direct a`

Disconnect from MAINT and log onto the system with the 6VMDIR10 user ID. From this user ID, back up the following files:

- `configaa datadvh z`
- `authfor control j`
- `extent control j`
- `user input j`
- `configrc datadvh e`

Full volume backup of systems allows for complete disaster recovery when another data center is available. There are a variety of methods of performing backups with Linux on System z. These include command-line tools included with every Linux distribution, such as dd, dump, cpio, tar, and dar. Also available are text-based utilities, such as Amanda, which is designed to add a more user-friendly interface to the backup and restore procedures.

Additionally, there are backup applications, such as IBM Tivoli Storage Manager (TSM). This solution is for backup management of open systems data. TSM simplifies backup by various means. Its built-in policies allow the backup of data to be largely automated, even in heterogeneous environments, which comprise hundreds or thousands of clients. Furthermore, administrative schedules automate the housekeeping of backup data (for example, creating copies of backup data, moving backup data from primary backup storage (disk) to secondary backup storage (tape), moving tapes to a vaulting location, and the deletion of old backups).

7.6 Disaster recovery for System z

With the emergence of business intelligence and dynamic warehousing, the disaster recovery requirements for a data warehouse environment are similar to that of online transaction processing. Therefore, it is important to consider disaster recovery scenarios before implementing a data warehouse solution.

Disaster recovery is an enterprise-specific plan for recovery. The implementation of the plan requires a thoroughly thought out methodology and dedicated organizational backing. Without a full organizational plan, full recovery of

business data could be jeopardized. For additional information about building a recovery strategy for an IBM Smart Analytics System data warehouse, see:

http://www.ibm.com/developerworks/data/bestpractices/isasrecovery/index.html

More details on disaster recovery related to System z can be found in the following IBM Redbooks publications:

- *Disaster Recovery with DB2 UDB for z/OS*, SG24-6370
- *GDPS Family - An Introduction to Concepts and Capabilities*, SG24-6374

z/VM archiving stores large bodies of data (for example, an entire disk image) for safekeeping, and should be a part of your disaster recovery plan. The data should be mutually consistent, so you can be running, but cannot be making changes. Archive at regular intervals, such as weekly or monthly, or whenever you do major software changes. These archives allow you to restore entire systems quickly. To restore backups, you need a running system, so after a system disaster, use your archive to restore the entire system, then use your backups to restore files. For more information about archives and backups on z/VM, see:

http://publib.boulder.ibm.com/epubs/pdf/hcsx0c00.pdf

z/VM provides two service programs for archiving:

- The DASD Dump Restore (DDR) utility program allows you to create archives of minidisks and complete DASD volumes. The program does not do incremental backups, so all data on a disk is archived whether or not it has changed. There are two versions of the program:
  - The DDR command, which you can issue from CMS
  - A standalone program that you can load (IPL)
- SPXTAPE produces an archive of spool files. Because NSSs (like CMS) and DCSSs are part of the spooling system, archive the spooling system. If problems develop with the spooling system and you need to do a CLEAN start of z/VM, it is much easier to restore archived NSSs and DCSSs instead of rebuild them.

7.7 Capacity management for System z

The IBM Smart Analytics System 9600 was created on 3390 DASD. It might be restored to the same DASD model with equal or greater capacity. The amount of free space varies with product mix, from volume to volume, and with the capacity of the receiving device.

If the system is restored to higher capacity DASD, the free-space indicator in the VTOC will be updated to include the additional space when the first new data set is allocated on each volume.

Capacity planning ensures that adequate resources are available in the future for critical workload to complete in an appropriate time. For capacity planning, you try to predict how changes in workload will change the requirements for all resources.

The focus of capacity management for System z is:

- Ongoing, with system utilization checked against a multi-period plan
- Evaluating impact of new applications
- Identifying and managing workload growth at a business function level
- Tasked to forecast capacity upgrades 3 - 6 months in advance

More information about the traditional steps in capacity planning can be found in section 1.6 in the IBM Redbooks publication *ABCs of z/OS System Programming Volume 11*, SG24-6327.

7.8 System Management Facilities

The z/OS system collects statistical data for each task when certain events occur in the life of the task. The System Management Facility (SMF) formats the information that it gathers into system-related (or job-related) records. System-related SMF records include information about the configuration, paging activity, and workload. Job-related records include information about the CPU time, SYSOUT activity, and data set activity of each job step, job, APPC/MVS transaction program, and TSO/E session. SMF data is written to the SYS1.MAN1, SYS1.MAN2, and SYS1.MAN3 data sets.[1] These data sets have been increased from the default sizing for the IBM Smart Analytics System 9600. The size of the data that the system can write to SMF data sets is constrained by the VSAM control interval size. SMF can only write one control interval at a time. The control interval size for these data sets has been set to 4096.

[1] For a more detailed discussion, see section 3.30 in the *ABCs of z/OS System Programming Volume 11*, SG24-6327.

The volume and variety of information in the SMF records enables the production of many types of analysis reports and summary reports. SMF formats the information that it gathers into system-related records or job-related records, as follows:

- System-related SMF records include information about the configuration, paging activity, and workload.
- Job-related records include information about the CPU time, SYSOUT activity, and data set activity of each job step, job, APPC/MVS transaction program, and TSO/E session.

SMF provides information about:

- System availability
- System or user abends
- VTOC errors
- Tape error statistics
- System configuration
- Device and channel data
- Job activity

Data from SMF records provides information that will enable:

- Billing users
- Reporting reliability
- Analyzing the configuration
- Scheduling jobs
- Summarizing DASD activity
- Evaluating data set activity
- Profiling system resource use
- Maintaining system security

SMF recording

When a subsystem or user program wants to write an SMF record, it invokes the SMF record macro SMFEWTM. This macro takes the user record and invokes SMF code to locate an appropriate buffer in the SMF address space and copy the data there. If the record is full, another SMF program is scheduled to locate full SMF buffers and write them to the SYS1.MANx data set. Each buffer is numbered to correspond to a particular record in the SMF data set. This allows the records to be written in any order and to place them correctly in the data set.

After all records have been written and the SYS1.MANx data set is full, SMF switches to a new SYS1.MANx data set and marks the full one as DUMP REQUIRED. That data set cannot be used again until it is dumped and cleared. Scheduling the SMF dump program must be done in a timely manner to ensure

that the SMF MANx data set is returned to use as soon as possible to ensure that no data is lost due to an *all data sets full* condition.

When the current recording data set cannot accommodate any more records, the SMF writer routine automatically switches recording from the active SMF data set to an empty SMF data set, and then passes control to the IEFU29 SMF dump exit. The operator is then informed that the data set needs to be dumped. When notified by the system that a full data set needs to be dumped, the operator will need to use the SMF data set dump program (IFASMFDP) to transfer the contents of the full SMF data set to another data set, and to reset the status of the dumped data set to empty so that SMF can use it again for recording data.

For more information about how to run the SMF data set dump program in z/OS v1r11, see:

http://publib.boulder.ibm.com/infocenter/zos/v1r11/index.jsp

7.9 Resource Measurement Facility (RMF)

RMF is the IBM product that is used for performance analysis, capacity planning, and problem determination in a z/OS host environment. Many different activities are required to keep the system running smoothly and to provide the best service on the basis of the available resources and workload requirements. This work is done by system operators, administrators, programmers, or performance analysts. RMF produces reports about problems as they occur, so that action can be taken before the problems become critical. RMF can be used to do the following:

- Determine that a system is running smoothly.
- Detect system bottlenecks caused by resource contention.
- Evaluate the service that an installation provides to various groups of users.
- Identify the workload delayed and the reason for the delayMonitor system failures, system stalls, and failures of selected applications.

For more information about RMF, visit the RMF home page:

http://www.ibm.com/servers/eserver/zseries/zos/rmf

7.9.1 RMF monitors

RMF comes with three monitors: Monitor I, Monitor II, and Monitor III. Because Monitor III has the ability to determine the cause of delay, use it to start your system-tuning activities. [2]

Monitor I

Monitor I provides long-term data collection for system workload and resource utilization. The Monitor I session is continuous, and measures various areas of system activity over a long period of time. You can obtain Monitor I reports directly as real-time reports for each completed interval (single-system reports only), or you can let the postprocessor run to create the reports, either as single-system reports or as sysplex reports. Many installations produce daily reports of RMF data for ongoing performance management. In this publication, sometimes a report is called a Monitor I report (for example, the workload activity report), although it can be created only by the postprocessor.

Monitor II

Monitor II provides online measurements on demand for use in solving immediate problems. A Monitor II session can be regarded as a snapshot session. Unlike the continuous Monitor I session, a Monitor II session generates a requested report from a single data sample. Because Monitor II is an ISPF application, you can use Monitor II and Monitor III simultaneously in split-screen mode to get different views of the performance of your system. In addition, you can use the RMF Spreadsheet Reporter to further process the measurement data on a workstation with the help of spreadsheet applications.

Monitor III

Monitor III provides short-term data collection and online reports for continuous monitoring of system status and solving performance problems. Monitor III is useful to begin system tuning because it allows the system tuner to distinguish between delays for important jobs and delays for jobs that are not as important to overall system performance.

7.9.2 RMF Spreadsheet Reporter overview

The RMF Spreadsheet Reporter is a workstation solution for graphical presentation of RMF Postprocessor data. Use it to convert your RMF data to spreadsheet format and generate representative charts for all performance-relevant areas. Performance data derived from SMF records is the basis for z/OS performance analysis and capacity planning. The basic idea of the RMF Spreadsheet Reporter is to exploit the graphical presentation facilities of a workstation for these purposes:

- It extracts performance measurements from SMF records.
- It produces postprocessor report listings and overview records.
- Converts this postprocessor output into spreadsheets.

[2] For a more detailed discussion, see Chapter 3 in the *ABCs of z/OS System Programming Volume 11*, SG24-6327.

Thus, the Spreadsheet Reporter offers a complete solution of enhanced graphical presentation of RMF measurement data.

The Spreadsheet Reporter also provides several sample spreadsheet macros to help you in viewing and analyzing performance data at a glance. For more detailed information, refer to “RMF Spreadsheet Reporter” in the IBM Redbooks publication *ABCs of z/OS System Programming Volume 11*, SG24-6327. For a description of how to use these functions, refer to *z/OS Resource Measurement Facility™ User’s Guide,* SC33-7990.

8

Managing users of the IBM Smart Analytics System 9600

This chapter contains the security and RACF requirements for the following components of the IBM Smart Analytics System 9600:

- TCP/IP and TELNET
- DB2 for z/OS
- InfoSphere Warehouse
- Cognos 8 BI (for reporting, query, analysis, and so on)

The IBM Smart Analytics System 9600 is installed with the Resource Access Control Facility (RACF). RACF was installed with the z/VM system and enabled along with DirMaint.

8.1 TCP/IP and TELNET

Certain RACF modifications are required for TCP/IP and TELNET. RACF makes sure that everyone who accesses the system resources is accountable. This applies to the system tasks as well. For system tasks, RACF associates every started task (STC) with a specific user ID. RACF keeps this information in a resource class called STARTED. For an STC to be started in the system, the STC user ID has to get access to all of the resources used by the STC. The IBM Smart Analytics System 9600 has been set up with the following:

- Group: The STC group created is named RACFSTC.
- Started procedures user IDs:
  - RACFSTCSYS
  - STCUSR

  These user IDs for the MVS STCs have been mapped in the RACF database.

A transport resource list (TRLE) statement is required for the OSA-Express to transfer data using TCP/IP.

8.2 DB2 for z/OS

The RACF access control module is supplied as an assembler source module in the DSNXRXAC member of prefix.SDSNSAMP of DB2 Version 9.1 for z/OS. It requires z/OS Version 1 Release 7 or later. z/OS Version 1 Release 7 provides limited support for DB2 roles. z/OS Version 1 Release 8 provides full support for roles, required for DB2 multi-level security.

The RACF access control module:

- Receives control from the DB2 access control authorization exit point (DSNX@XAC) to handle DB2 authorization checks
- Provides a single point of control for RACF and DB2 security administration
- Provides the ability to define security rules before a DB2 object is created
- Allows security rules to persist when a DB2 object is dropped
- Provides the ability to protect multiple DB2 objects with a single security rule using a combination of RACF generic, grouping, and member profiles
- Eliminates the DB2 cascading revoke
- Preserves DB2 privileges and administrative authorities

- Provides flexibility for multiple DB2 subsystems with a single set of RACF profiles
- Allows you to validate a user ID before giving it access to a DB2 object

RACF support for the RACF access control module includes a set of general resource classes in the RACF module ICHRRCDX (the supplied portion of the RACF class descriptor table). These classes are used when you implement the RACF access control module using the default values.

The RACF access control module checks the RACF profiles corresponding to that set of privileges and authorities:

- Authority checking performed by the RACF access control module simulates DB2 authority checking.
- DB2 object types map to RACF class names.
- DB2 privileges map to RACF resource names for DB2 objects.
- DB2 authorities map to the RACF administrative authority class (DSNADM) and RACF resource names for DB2 authorities.
- DB2 security rules map to RACF profiles.

RACF profiles for DB2 for z/OS

This section details the RACF user IDs and profiles for DB2 for z/OS that have been created for you for the IBM Smart Analytics System 9600.

The following code was used to define DB2 RACF profiles for you:

```
RDEFINE SERVER  (DB2.DB2I.WLMENV)         UACC(NONE)
RDEFINE SERVER  (DB2.DB2I.WLMENVJ)        UACC(NONE)
RDEFINE SERVER  (DB2.DB2I.WLMENV_RACFPC)  UACC(NONE)
RDEFINE SERVER  (DB2.DB2I.WLMUTL1)        UACC(NONE)
RDEFINE SERVER  (DB2.DB2I.REXX_WLMENV)    UACC(NONE)
SETROPTS RACLIST(SERVER) REFRESH
PERMIT DB2.DB2I.WLMENV         CLASS(SERVER) ID(STCGRP) ACCESS(READ)
PERMIT DB2.DB2I.WLMENVJ        CLASS(SERVER) ID(STCGRP) ACCESS(READ)
PERMIT DB2.DB2I.WLMENV_RACFPC  CLASS(SERVER) ID(STCGRP) ACCESS(READ)
PERMIT DB2.DB2I.WLMUTL1        CLASS(SERVER) ID(STCGRP) ACCESS(READ)
PERMIT DB2.DB2I.REXX_WLMENV    CLASS(SERVER) ID(STCGRP) ACCESS(READ)
SETROPTS RACLIST(SERVER) REFRESH
SETROPTS CLASSACT(DSNR) GENERIC(DSNR)
RDEFINE DSNR **.BATCH
RDEFINE DSNR **.DIST
RDEFINE DSNR **.RRSAF
RDEFINE DSNR (DB2I.WLM_REFRESH.WLMENV)
PE DB2I.WLM_REFRESH.WLMENV +
```

```
   CLASS(DSNR) ID(STCGRP) ACCESS(READ)
SETROPTS RACLIST(DSNR) REFRESH
```

8.3 InfoSphere Warehouse

Predefined for the you are the InfoSphere Warehouse for System z (ISWz) DB2 connections, using the connection ID ISWZADM. Use the user ID ISWZADM connection only for access to the ISWz metadata repository. For ISWz access of z/OS DB2 data warehouses, we suggest setting up additional connections, minimally one for SQW work and one for Cubing Services, for example, "ISWZREF" for SQW refresh data flows and "ISWZCUB" for cubing services. You can use whatever IDs you prefer, but we suggest updating the DDF classification rules to match.

The InfoSphere Warehouse RACF user ID ISWZADM with the password ISWZADM has been set up for the IBM Smart Analytics System 9600. This ID is the InfoSphere Warehouse administration user ID and has TSO privileges and access to SDSF. This user ID should only be used to access the InfoSphere Warehouse for System z metadata repository. User IDs for SQW work and cubing services should be set up, as well as any other user IDs needed for general access of the DB2 for z/OS data warehouses.

The password for ISWZADM can be changed from in TSO with the following RACF command:

```
==>TSO ALU ISWZADM PASSWORD(newpassword) RESUME
```

The following commands were run in DB2 SPUFI to grant authority to ISWZADM:

```
GRANT BINDADD     TO ISWZADM WITH GRANT OPTION ;
GRANT CREATEALIAS TO ISWZADM WITH GRANT OPTION ;
GRANT CREATEDBA   TO ISWZADM WITH GRANT OPTION ;
GRANT CREATEDBC   TO ISWZADM WITH GRANT OPTION ;
GRANT CREATESG    TO ISWZADM WITH GRANT OPTION ;
GRANT CREATETMTAB TO ISWZADM WITH GRANT OPTION ;
GRANT CREATE ON COLLECTION * TO ISWZADM ;
```

In addition, SELECT has been granted (**`Grant SELECT ON`**) to all of the DB2 system catalog and directory tables (except IPLIST, IPNAMES, LOCATIONS, LULIST, LUMODES, LUNAMES, and SYSDBAUTH) to allow the InfoSphere administrator to build the DB2-related objects for the InfoSphere warehouse. The following commands were used to do this:

```
GRANT SELECT ON SYSIBM.SYSAUXRELS         TO ISWZADM ;
  GRANT SELECT ON SYSIBM.SYSCHECKDEP        TO ISWZADM ;
```

```
GRANT SELECT ON SYSIBM.SYSCHECKS          TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSCHECKS2         TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSCOLDIST         TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSCOLDISTSTATS    TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSCOLDIST_HIST    TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSCOLSTATS        TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSCOLUMNS         TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSCOLUMNS_HIST    TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSCONSTDEP        TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSCONTEXT         TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSCOPY            TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSCTXTTRUSTATTRS  TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSDATABASE        TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSDATATYPES       TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSDBRM            TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSDEPENDENCIES    TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSDUMMY1          TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSDUMMYA          TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSDUMMYE          TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSDUMMYU          TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSENVIRONMENT     TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSFIELDS          TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSFOREIGNKEYS     TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSINDEXES         TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSINDEXES_HIST    TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSINDEXPART       TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSINDEXPART_HIST  TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSINDEXSPACESTATS TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSINDEXSTATS      TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSINDEXSTATS_HIST TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSKEYCOLUSE       TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSKEYS            TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSKEYTARGETS      TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSKEYTARGETSTATS  TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSKEYTARGETS_HIST TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSKEYTGTDIST      TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSKEYTGTDISTSTATS TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSKEYTGTDIST_HIST TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSLOBSTATS        TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSLOBSTATS_HIST   TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSOBDS            TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSOBJROLEDEP      TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSPACKAGE         TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSPACKDEP         TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSPACKLIST        TO ISWZADM ;
```

```
GRANT SELECT ON SYSIBM.SYSPACKSTMT        TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSPARMS           TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSPKSYSTEM        TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSPLAN            TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSPLANDEP         TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSPLSYSTEM        TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSRELS            TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSROLES           TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSROUTINES        TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSSEQUENCES       TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSSEQUENCESDEP    TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSSTMT            TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSSTOGROUP        TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSSTRINGS         TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSSYNONYMS        TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSTABCONST        TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSTABLEPART       TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSTABLEPART_HIST  TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSTABLES          TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSTABLESPACE      TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSTABLESPACESTATS TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSTABLES_HIST     TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSTABSTATS        TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSTABSTATS_HIST   TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSTRIGGERS        TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSVIEWDEP         TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSVIEWS           TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSVLTREE          TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSVOLUMES         TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSVTREE           TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSXMLRELS         TO ISWZADM ;
GRANT SELECT ON SYSIBM.SYSXMLSTRINGS      TO ISWZADM ;
```

Buffer pool use privilege for regular tablespaces, LOB tablespaces, and indexes has also been granted to the ISWZADM user ID.

The privilege to call the following stored procedures has also been granted:

- DSNUTILU, DSNWZP
- ADMIN_JOB_SUBMIT
- ADMIN_JOB_QUERY
- ADMIN_JOB_FETCH
- ADMIN_JOB_CANCEL
- ADMIN_DS_BROWSE

User ISWZUSR has been created with the password ISWZUSR as well. The DBADM authority under the following owners has been granted to this user on metadata database tables:

- DBDRV_RES
- COM_CONFIG
- DB_RES, DWEREPOS
- ISWSCHED, SQWMETA_V2
- PROCMGMT, SQWSTAT_V2)

In addition, ALTERIN, CREATEIN, and DROPIN have been granted to tables using the schema SQWMETA-V2.

8.4 IBM Cognos 8 BI

Two IBM Cognos BI connections have been created for the IBM Smart Analytics System 9600:

- One for the IBM Cognos content store
- One for the sample IBM Smart Analytics System z/OS DB2 data warehouse database

Both use the DB2 for z/OS data warehouse. Both connections use the COGZADM connection ID. Use the COGZADM user ID for the content store connection, but create one or more additional connection IDs for one Cognos DB2 data warehouse access. You can define/utilize whatever IDs you prefer, just remember to update your DDF classification rules to match.

8.4.1 DB2 customization for IBM Cognos 8 BI

Some DB2 specifics for IBM Cognos 8 BI are:

- IBM Cognos 8 BI content store database name: COGZDB
- Audit or logging database: COGZAUD
- Content store tablespace name: COGZTS
- Content store stogroup: COGZSG
- Content store large bufferpool: BP32K1
- Content store regular bufferpools: BP1, B2, BP4, BP8

Member COGZGRAN was created with the grants shown in Figure 8-1.

```
SET CURRENT SQLID = 'ISASD1' ;
-
- EXECUTE THE FOLLOWING GRANTS FOR THE COGNOS ADMINISTRATOR USER
- NOTE THAT THE Z/OS GROUP HAS THE RESPONSIBILITY OF SETTING UP
- USERID IN RACF.
-
GRANT DBADM ON DATABASE COGZDB TO COGZADM WITH GRANT OPTION ;
GRANT DBADM ON DATABASE COGZAUD TO COGZADM WITH GRANT OPTION ;
GRANT USE OF BUFFERPOOL BP2 TO COGZADM WITH GRANT OPTION ;
GRANT USE OF BUFFERPOOL BP4 TO COGZADM WITH GRANT OPTION ;
GRANT USE OF BUFFERPOOL BP1 TO COGZADM WITH GRANT OPTION ;
GRANT USE OF BUFFERPOOL BP32K1 TO COGZADM WITH GRANT OPTION ;
GRANT USE OF BUFFERPOOL BP8K1 TO COGZADM WITH GRANT OPTION ;
GRANT USE OF STOGROUP COGZSG TO COGZADM WITH GRANT OPTION ;
-
```

Figure 8-1 DB2 GRANTs for Cognos Admin user ID

If you enter `DB2I.DSNDBD.COGZDB.CZ*` in ISPF panel 3.4, you will be able to see where that the Cognos Content Store tablespaces were allocated on the SMS STORCLAS DB2DATA volumes. Additionally, a list of Cognos DB2 datasets should exist that look similar to:

`DB2I.DSNDBD.COGZDB.CZCML***.I0001.A001`

8.4.2 Cognos 8 Security

Cognos 8 Security is designed to meet the need for security in various situations.

The security model can be easily integrated with the existing RACF security setup on z/OS and DB2 for z/OS system. Cognos security is built on top of your existing RACF/DB2 security model. You use RACF and DB2 to define and maintain users, groups, and roles, and to control access. Each authentication provider known to Cognos 8 is referred to as a namespace.

In addition to the external namespaces that represent the RACF and DB2 authentication model, Cognos 8 has its own namespace called Cognos. The Cognos namespace enhances organization security policies and deployment ability of applications.

Figure 8-2 demonstrates the relationships between DB2 and Cognos.

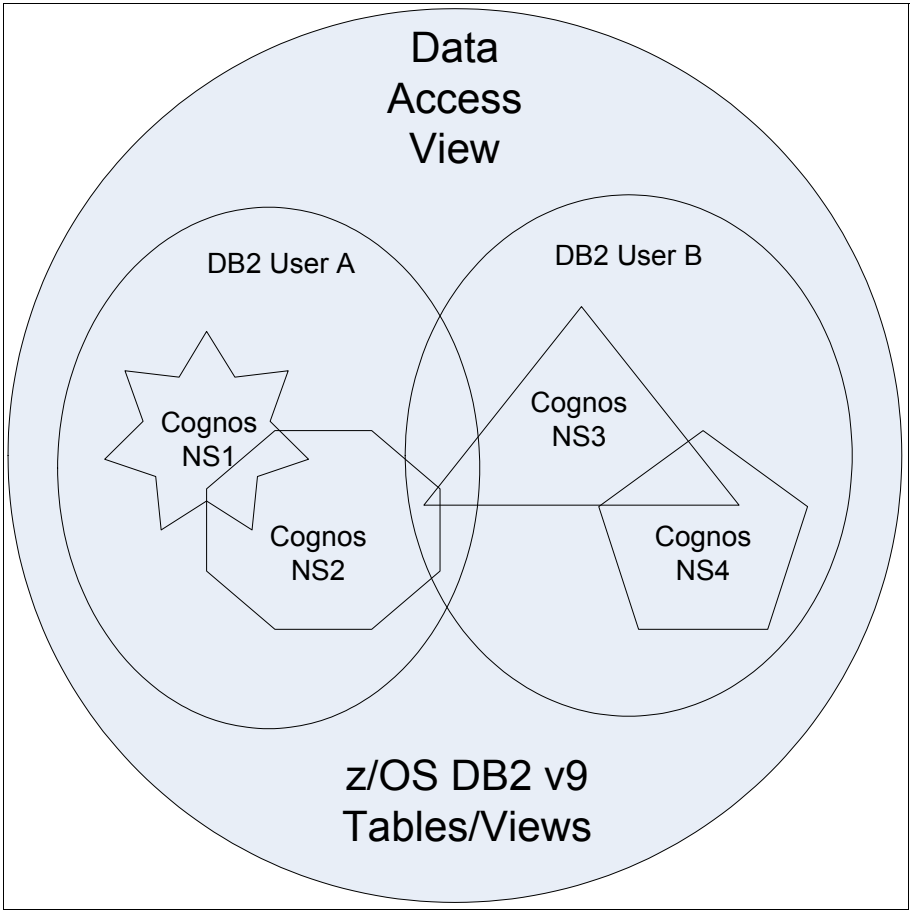

Figure 8-2 RACF/DB2/Cognos data access view

8.4.3 Authentication providers

User authentication to the reporting interface by users is controlled by Cognos. User authentication to a data source in Cognos 8 is managed by RACF and DB2. Cognos users log in to a namespace, which uses a login to RACF with further security to DB2 through a DB2 user profile. Figure 8-3 shows an overview of this.

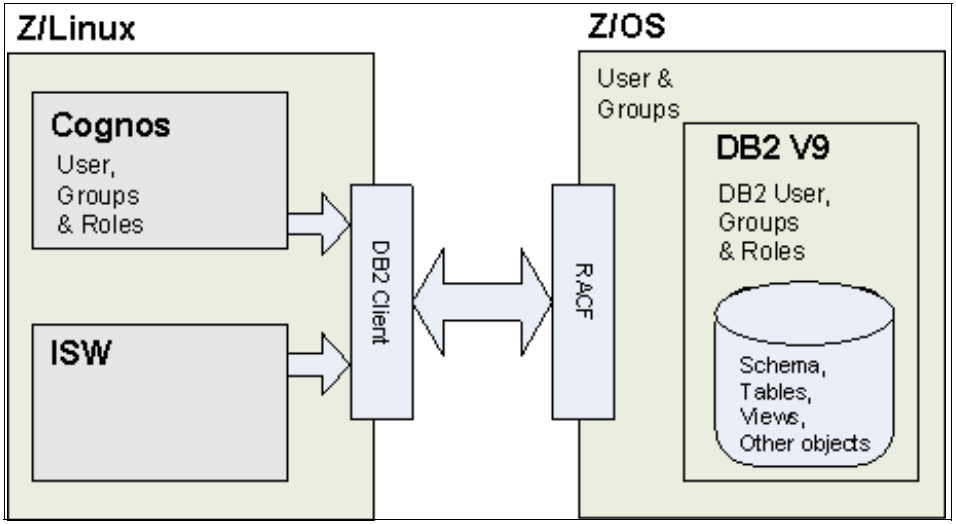

Figure 8-3 Security overview diagram

RACF defines users, then DB2 defines users, groups, and roles used for authentication. If you set up authentication for Cognos 8, users must provide valid credentials, such as a user ID and password, at logon time. The RACF user is configured in the Cognos administration and transparent to the report user. Each namespace will use one of the RACF users, which will limit access to data on z/OS according to the rules set for that user as a DB2 user.

If multiple namespaces have been configured for your system, at the start of a session you must select one namespace that you want to use. However, this does not prevent you from logging on to other namespaces later in the session. For example, if you set access permissions, you might want to reference entries from different namespaces. To log on to a different namespace, you do not have to log out of the namespace that you are currently using. You can be logged on to multiple namespaces simultaneously.

Your primary logon is the namespace and the credentials that you used to log on at the beginning of the session. The namespaces that you log on to later in the session and the credentials that you use become your secondary logons.

Cognos 8 does not replicate the users, groups, and roles defined in your RACF and DB2 configuration.

However, you can reference them in Cognos 8 when you set access permissions to reports and other content. They can also become members of Cognos groups and roles.

You configure authentication providers using Cognos configuration. For more information, see the *Installation and Configuration Guide* which can be found at:

http://publib.boulder.ibm.com/infocenter/c8bi/v8r4m0/index.jsp?topic=/com.ibm.swg.im.cognos.inst_cr_winux.8.4.0.doc/inst_cr_winux.html

8.4.4 Authorization

Authorization is the process of granting or denying users access to data, and permission to perform activities on that data, based on their signon identity.

Cognos 8 authorization assigns permissions to users, groups, and roles that allow them to perform actions, such as read or write, on content store objects, such as folders and reports. The content store can be viewed as a hierarchy of data objects. These objects include not only folders and reports, but packages for report creation, directories, and servers.

When Cognos 8 administrators distribute reports to users, they can set up folders in which reports and other objects can be stored. They can then secure those folders so that only authorized personnel can view, change, or perform other tasks using the folder contents.

For information about the Content Manager hierarchy of objects and the initial access permissions, see "Initial Access Permissions" in the *IBM Cognos 8 Administration and Security Guide*, which can be found at:

http://download.boulder.ibm.com/ibmdl/pub/software/data/cognos/documentation/docs/en/8.4.0/ug_cra.pdf

For information about setting access permissions to the Cognos 8 entries, see "Access Permissions" in the *Cognos 8 Administration and Security Guide*.

8.4.5 Cognos namespace

The Cognos namespace is the Cognos 8 built-in namespace. It contains the Cognos objects, such as groups, roles, data sources, distribution lists, and contacts.

During the content store initialization, built-in and predefined security entries are created in this namespace. You must modify the initial security settings for those entries and for the Cognos namespace immediately after installing and configuring Cognos 8.

You can rename the Cognos namespace using Cognos Configuration, but you cannot delete it.

When you set security in Cognos 8, you might want to use the Cognos namespace to create groups and roles that are specific to Cognos 8. In this namespace, you can also create security policies that indirectly reference the third-party security entries so that Cognos 8 can be more easily deployed from one installation to another.

The Cognos namespace always exists in Cognos 8, but the use of Cognos groups and roles that it contains is optional. The groups and roles created in the Cognos namespace repackage the users, groups, and roles existing in the authentication providers to optimize their use in the Cognos 8 environment. For example, in the Cognos namespace, you can create a group called HR Managers and add to it specific users and groups from your corporate IT and HR organizations defined in your authentication provider. Later, you can set access permissions for the HR Managers group to entries in Cognos 8.

8.4.6 Optimizing users, groups, and roles in Cognos Namespace

If you are maintaining groups and roles in the Cognos namespace for ease of deployment, it is best to populate groups and roles with users in RACF, and then add those groups and roles to the Cognos groups and roles that are appropriate. Otherwise, you might have trouble managing large lists of users in a group in the Cognos namespace.

8.4.7 Application security

To supplement the existing Cognos 8 security and to further prevent inadvertent and malicious attacks, Cognos Application Firewall is enabled by default.

Cognos Application Firewall is a security tool designed to supplement the existing Cognos 8 security infrastructure, at the application level. Cognos Application Firewall acts as a smart proxy for the Cognos product gateways and dispatchers and works to prevent the Cognos 8 products from processing malicious data. HTTP and XML requests are analyzed, modified, and validated before the gateways or dispatchers process them, and before they are sent to the requesting client or service.

Cognos Application Firewall is configured using the Cognos 8 configuration tool. For more information about its features, see "Cognos Application Firewall" in the *Cognos 8 Administration and Security Guide*.

8.5 Cognos users, groups, and roles

Users, groups, and roles are created for authentication and authorization purposes. In Cognos 8, you can use users, groups, and roles created in third-party authentication providers, and groups and roles created in Cognos 8. The groups and roles created in Cognos 8 are referred to as Cognos groups and Cognos roles.

8.5.1 Users

A user entry is created and maintained in a third-party authentication provider to uniquely identify a human or a computer account. You cannot create user entries in Cognos 8.

Information about users, such as first and last names, passwords, IDs, locales, and email addresses, is stored in the providers. However, this might not be all the information required by Cognos 8. For example, it does not specify the location of the users' personal folders, or format preferences for viewing reports. This additional information about users is stored in Cognos 8, but when addressed in Cognos 8, the information appears as part of the external namespace.

8.5.2 Deleting and recreating users

If you use an LDAP server, the stability of My Folders objects depends on how you use the IDs. If the configuration of the LDAP provider uses with the default attribute of dn for the unique identifier parameter, a reinstated user with the same name keeps the My Folders objects of the original user.

You can delete, copy, and change user profiles. For more information, see "Managing User Profiles" in the *Cognos 8 Administration and Security Guide*.

8.5.3 User locales

A locale specifies linguistic information and cultural conventions for character type, collation, format of date and time, currency unit, and messages. You can

specify locales for individual products, content, servers, authors, and users in Cognos 8.

User locale refers to the product and content locales for each Cognos 8 user. Requests from users arrive with an associated locale. Cognos 8 must determine the language and locale preferences of users and enforce an appropriate response locale when you distribute reports in different languages.

A user locale specifies the default settings that a user wants to use for formatting dates, times, currency, and numbers. Cognos 8 uses this information to present data to the user.

Cognos 8 obtains a value for user locale by checking these sources, in the order listed:

1. User preference settings

   If the user sets the user preference settings in Cognos Connection, Cognos 8 uses these settings for the user's product and content locale and for default formatting options. The user preference settings override the values obtained from the authentication provider.

2. Authentication provider

   If the authentication provider has locale settings that are configured, Cognos 8 uses these values for the user's product and content locale.

3. Browser setting

   Anonymous and guest users cannot set user preference settings. For these users, Cognos 8 obtains a user locale from the browser stored on the user's computer.

8.5.4 Groups and roles

Users can become members of groups and roles defined in third-party authentication providers, and groups and roles defined in Cognos 8. A user can belong to one or more groups or roles. If users are members of more than one group, their access permissions are merged.

Groups and roles represent collections of users that perform similar functions, or have a similar status in an organization. Examples of groups are employees, developers, or sales personnel.

Members of groups can be users and other groups. When users log on, they cannot select a group that they want to use for a session. They always log on with all the permissions associated with the groups to which they belong.

Roles in Cognos 8 have a similar function as groups. Members of roles can be users, groups, and other roles.

Figure 8-4 shows the structure of groups and roles.

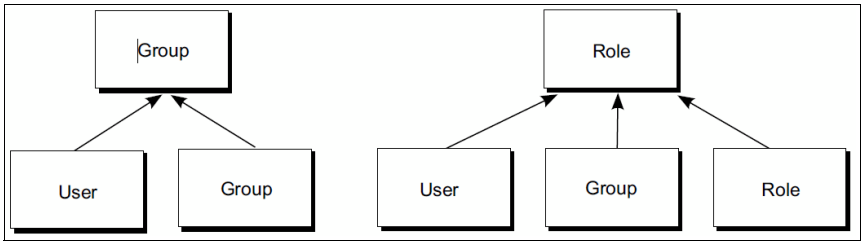

Figure 8-4

You create Cognos groups and roles when:

- You cannot create groups or roles in your authentication provider.
- Groups or roles are required that span multiple namespaces.
- Portable groups and roles are required that can be deployed.

In this case, it is best to populate groups and roles in the third-party provider, and then add those groups and roles to the Cognos groups and roles to which they belong. Otherwise, you might have trouble managing large lists of users in a group in the Cognos namespace. Two key things to keep in mind when adding groups and roles to the Cognos groups and roles are:

- Address the specific needs of the administration of Cognos 8.
- Avoid cluttering your organization security systems with information used only in Cogno 8.

8.5.5 Access permissions

In Cognos 8, you can secure your organization's data by setting access permissions for the entries. You specify which users and groups have access to a specific report or other content in Cognos 8. You also specify the actions that they can perform on the content.

When you set access permissions, you can reference RACF users, groups, DB2 users, groups and roles, and Cognos groups and roles. However, if you plan to deploy your application in the future, use only the Cognos groups and roles to set up access to entries in Cognos 8 to simplify the process.

8.5.6 Cognos Application Firewall

Business intelligence solutions are frequently critical to your operations. Cognos Application Firewall is a tool designed to supplement the existing Cognos 8 security infrastructure. By default, this supplemental security is enabled.

Cognos Application Firewall acts as a smart proxy for the Cognos product gateways and dispatchers. HTTP and XML requests are analyzed, modified, and validated before the gateways or dispatchers process them, and before they are sent to the requesting client or service.

Cognos Application Firewall works to protect the Cognos 8 products from processing malicious data. The most common forms of malicious data are buffer overflows and cross-site scripting attacks (XSS links), either through script injection in valid pages or redirection to other websites. For information about enabling the Cognos Application Firewall, see the *Installation and Configuration Guide*.

The following objects must be created on DB2 for z/OS for the Cognos:

- DB2 content store database, stogroup, and user ID

  These are required to access DB2 and create and delete Cognos databases.
- DB2 notification database
- DB2 logging database

8.6 Configuring IBM Cognos 8 components to use LDAP

You can configure IBM Cognos 8 components to use an LDAP namespace as the authentication provider. You can use an LDAP namespace for users that are stored in an LDAP user directory, Active Directory Server, IBM Directory Server, Novell Directory Server, or Sun Java System Directory Server.

You can also use LDAP authentication with DB2 and Essbase OLAP data sources by specifying the LDAP namespace when you set up the data source connection. For more information, see the *Administration and Security Guide*.

You also have the option of making custom user properties from the LDAP namespace available to IBM Cognos 8 components.

To bind a user to the LDAP server, the LDAP authentication provider must construct the distinguished name (DN). If the Use external identity property is set

to True, it uses the External identity mapping property to try to resolve the user's DN. If it cannot find the environment variable or the DN in the LDAP server, it attempts to use the User lookup property to construct the DN.

If users are stored hierarchically within the directory server, you can configure the User lookup and External identity mapping properties to use search filters. When the LDAP authentication provider performs these searches, it uses the filters that you specify for the User lookup and External identity mapping properties. It also binds to the directory server using the value that you specify for the Bind user DN and password property or using anonymous if no value is specified.

When an LDAP namespace has been configured to use the External identity mapping property for authentication, the LDAP provider binds to the directory server using the Bind user DN and password or using anonymous if no value is specified. All users who log on to IBM Cognos 8 using external identity mapping see the same users, groups, and folders as the Bind user.

If you do not use external identity mapping, you can specify whether to use bind credentials to search the LDAP directory server by configuring the Use bind credentials for search property. When the property is enabled, searches are performed using the bind user credentials or using anonymous if no value is specified. When the property is disabled, which is the default setting, searches are performed using the credentials of the logged-on user. The benefit of using bind credentials is that instead of changing administrative rights for multiple users, you can change the administrative rights for the bind user only.

8.7 Cognos security model

The security model can be easily integrated with the existing security infrastructure in your organization. It is built on top of one or more third-party *authentication providers.* You use the providers to define and maintain users, groups, and roles, and to control the authentication process. Each authentication provider known to Cognos 8 is referred to as a namespace.

In addition to the external namespaces that represent the third-party authentication providers, Cognos 8 has its own namespace called *Cognos*. The Cognos namespace enhances your organization security policies and deployment ability of applications.

Security in Cognos 8 is optional. If security is not enabled it means that no third-party authentication providers are configured, and therefore all user access is anonymous. Anonymous users have limited, read-only access.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks

For information about ordering these publications, see “How to get Redbooks” on page 118. Note that some of the documents referenced here might be available in softcopy only.

- *DB2 9 for z/OS Technical Overview*, SG24-7330
- *DB2 9 for z/OS Performance Topics*, SG24-7473
- *Enterprise Data Warehousing with DB2 9 for z/OS*, SG24-7637
- *DB2 9 for z/OS Stored Procedures: Through the CALL and Beyond*, SG24-7604
- *Co-locating Transactional and Data Warehouse Workloads on System z*, SG24-7726
- *50 TB Data Warehouse Benchmark on IBM System z*, SG24-7674
- *InfoSphere Warehouse: A Robust Infrastructure for Business Intelligence*, SG24-7813
- *Using IBM System z As the Foundation for Your Information Management Architecture*, REDP-4606
- *IBM z/OS Application Connectivity to DB2 for z/OS and OS/390*, TIPS0356

Other publications

The following publication is also relevant as an information source:

- *DB2 Data Warehouse Edition V9.1,* GC18-9800

Online resources

The following websites are also relevant as further information sources:

- *IBM Cognos 8 Administration and Security Guide*

  http://download.boulder.ibm.com/ibmdl/pub/software/data/cognos/documentation/docs/en/8.4.0/ug_cra.pdf

- *IBM Cognos 8 Installation and Configuration Guide*

  http://publib.boulder.ibm.com/infocenter/c8bi/v8r4m0/index.jsp?topic=/com.ibm.swg.im.cognos.inst_cr_winux.8.4.0.doc/inst_cr_winux.html

How to get Redbooks

You can search for, view, or download Redbooks, Redpapers, Technotes, draft publications and Additional materials, as well as order hardcopy Redbooks publications, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

E

F

G

H

I

L

M

O

P

Q

R

S

T

U

W

X

Z

Getting Started with the IBM Smart Analytics System 9600

Getting the most from the IBM Smart Analytics System 9600

Understanding System z and the IBM Smart Analytics System 9600

Managing the components

The IBM Smart Analytics System 9600 is a single, end-to-end business analytics solution to accelerate data warehousing and business intelligence initiatives. It provides integrated hardware, software, and services that enable enterprise customers to quickly and cost-effectively deploy business-changing analytics across their organizations.

As a workload-optimized system for business analytics, it leverages the strengths of the System z platform to drive:

- Significant savings in hardware, software, operating, and people costs to deliver a complete range of data warehouse and BI capabilities
- Faster time to value with a reduction in the time and speed associated with deploying Business Intelligence
- Industry-leading scalability, reliability, availability, and security
- Simplified and faster access to the data on System z

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-7902-00 ISBN 0738435651