IBM® Business Analytics

Using zEnterprise for Smarter Analytics
Volume 2 Implementation

- Explore the use of COGNOS, SPSS, and DB2
- Learn how to set up analytics tools on System z
- Determine how to test the workload

Mike Ebbers
Alex Louwe Kooijmans
Colin Page
Patrick Hempeler
Qiong Liu
LK Swift
Lih Wang

Redbooks

ibm.com/redbooks

International Technical Support Organization

Using zEnterprise for Smarter Analytics: Volume 2 Implementation

April 2013

SG24-8008-00

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

First Edition (April 2013)

This edition applies to IBM Cognos Business Intelligence V10, SPSS Modeler 14.2, and DB2.

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 on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

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

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

AIX®
BladeCenter®
Cognos®
DB2®
DB2 Connect™
eServer™
FICON®
HACMP™
HiperSockets™
IBM®
IMS™
Informix®
OMEGAMON®
Parallel Sysplex®
POWER®
POWER7®
PowerHA®
PowerVM®
PR/SM™
RACF®
Rational®
Redbooks®
Redbooks (logo) ®
Smarter Analytics™
Smarter Banking®
SPSS®
System x®
System z®
Tivoli®
TM1®
WebSphere®
z/OS®
z/VM®
zEnterprise®
zSeries®

The following terms are trademarks of other companies:

Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

ITIL is a registered trademark, and a registered community trademark of The Minister for the Cabinet Office, and is registered in the U.S. Patent and Trademark Office.

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

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

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

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

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

Preface

This IBM® Redbooks® publication series explains the assessment and implementation of a workload, integrated within IBM Smarter Banking® Showcase, and hosted at IBM Montpellier, France. Intended for decision-makers, consultants, architects, administrators, and specialists, this book is the second volume in a series of two:

- Assessment: Volume 1 (SG24-8007) describes how to evaluate the requirements of a new IBM Smarter Analytics™ workload, addressing the user, system resources, and data processing profiles to identify the most optimal configuration by using IBM methodologies, such as fit-for-purpose. Given that the existing showcase is based on the IBM zEnterprise® System, deployment options include IBM z/OS®, Linux on IBM System z®, IBM AIX® running on IBM POWER® processor-based blades within the zEnterprise IBM BladeCenter® Extension (zBX), and Windows Server 2008 running on IBM System x® and BladeCenter blades also within zBX.
- Implementation: Volume 2 (SG24-8008), which you are reading, describes the setups that are involved in deploying the Smarter Analytics workload within the showcase. With multiple components, including IBM Cognos® BI, IBM Cognos TM1®, Cognos Metric Studio, IBM DB2® for z/OS, and a number of application design tools, the workload spans multiple operating environments. The use of application clustering, setting up performance policies by using Unified Resource Manager, and simulation test execution results are included.

If you want to see this solution in action, contact your IBM representative or send email to mailto:mopbcoe@fr.ibm.com

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 (ITSO), Poughkeepsie Center.

Mike Ebbers is a Consulting IT Specialist and Project Leader at the International Technical Support Organization, Poughkeepsie Center. He has worked with IBM mainframe hardware and software products since 1974 in the field, in education, and in the ITSO.

Alex Louwe Kooijmans is a Solution Architect with the Center of Excellence for Financial Services in the IBM Systems & Technology division. He specializes in core banking transformation, with a focus on the IBM System z platform. Until recently, he was a Project Leader with the ITSO, specializing in service-oriented architecture (SOA) technology and solutions on System z, in addition to application modernization and transformation on IBM z/OS. Previously, he worked as a Client IT Architect in the financial services sector with IBM in The Netherlands. Alex also worked at the Technical Marketing Competence Center for IBM eServer™ zSeries® and Linux in Boeblingen, Germany. From 1997 to 2000, he completed his first assignment with the ITSO, managing IBM Redbooks publications projects and delivering workshops on IBM WebSphere®, Java, and e-business technology on System z. Before 1997, Alex held positions in application design and development, product support, and project management, mostly in relation to the IBM mainframe.

Colin Page is the Lead Architect for the Smarter Banking Showcase with IBM in Montpellier, France, although originally from the UK. He is a Certified IT Specialist & System zChampion. He has over 20 years of experience working on mainframe projects, mostly at customer sites, specializing in DB2 for z/OS and now across all middleware and infrastructure components.

Patrick Hempeler is a Technical Sales Specialist with an IBM Integrated Operating Team (IOT) in Europe, located in Stuttgart, Germany. He holds a Bachelor of Science in Business Information Technologies from Corporate State University Baden Württemberg and a Bachelor of Arts in Business Administration from Open University London. His areas of expertise include a wide range of information management topics (data integration, data warehousing, business analytics) and also client/server and mainframe computing.

Qiong Liu is a senior IT Specialist with the System z Advanced Technical Skills team at IBM Beijing, China. She has six years of experience in System z, z/OS and IBM Parallel Sysplex®, zBX, and Unified Resource Manager. She holds a PhD degree in Optical-Electronical Engineering from Beijing Institute of Technology. Her areas of expertise include System z, z/OS and Parallel Sysplex, capacity planning and sizing, performance testing and tuning, and also major z/OS software such as DB2 and MQ.

LK Swift is an Executive IT Architect in Atlanta, GA, USA. He is a Senior Certified IT Architect, and has 25 years of experience in IT, most of which were spent creating global solutions for the IBM CIO organization. He holds a degree in Computer Science and Engineering from the University of Florida. His areas of expertise include IT Infrastructure, distributed systems, and the migration of computer systems. LK's most recent project is the migration of thousands of stand-alone servers in IBM to strategic, virtualized platforms including Linux on System z.

Lih Wang is an IBM certified Consulting IT Specialist in Chicago for the client-facing technical sales organization. She has over 25 years of experience in enterprise systems management and product management areas. She is an advocate for the new zEnterprise Innovation and specializes in performance optimization for z/OS, DB2, IBM IMS™, WebSphere MQ, and mainframe connectivity. She has authored performance monitoring topics and presented at DB2, IMS, and MQ technical conferences and users groups during the last two decades. Lih is an ITIL certified member and holds a Master's degree in Computer Science.

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

Guillaume Arnould
Stephane Faure
Fabrice Jarassat
Marie-Laure Pessoa Caucat
Nigel Williams
Stephane Winter
IBM Product and Solutions Support Center (PSSC), Montpellier, France

Yann Kindelberger
IBM European Design Center, Montpellier, France

Dennis Mosby, for providing examples of Cognos sizing
IBM Techline

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

Introduction

Within IBM Financial Services, several divisions have come together to create a set of business analytics assets of specific use to the banking, insurance, and financial market industries. These assets are referred to as *blueprints*.

In a previous exercise[1], a blueprint, which is designed to provide profitability information for both the customer and commercial operational aspects of a retail bank, is assessed for integration with an existing bank's information technology environment. The blueprint is known as the *Customer Profitability Analytics* (CPA) blueprint.

The CPA blueprint consists of a number of models that were developed by using a common suite of IBM products. The principal components are as follows:

- IBM Cognos TM1 Server (9.5.0 or later)
- IBM Cognos 10.1 Business Intelligence Server
- IBM Cognos 10.1 Metrics Server
- IBM DB2 Database 9.5 (Express-C or later)
- IBM SPSS® Modeler 14.2

The blueprint was developed and made available to IBM sales support personnel to demonstrate to customers. The package contains sample data and is based upon a Windows 32-bit runtime environment. However, based on the findings of the previously mentioned assessment, we implemented the IBM CPA blueprint on several platforms within an IBM zEnterprise System (zEnterprise).

[1] Described in *Using zEnterprise For Smart Analytics: Volume 1 Assessment*, SG24-8007 (not yet available)

1.1 Summary of the assessment phase

Before we started the actual implementation of CPA and its required middleware components on zEnterprise, we conducted an assessment study to learn about various aspects:

- Whether zEnterprise would provide business value for this specific workload
- Which platform topology would be the correct one within zEnterprise
- How many resources would be needed to run CPA on zEnterprise

At the end of the assessment study, a business case was produced as the starting point for the implementation phase, described in this book. Figure 1-1 shows an overview of the steps and decision points involved.

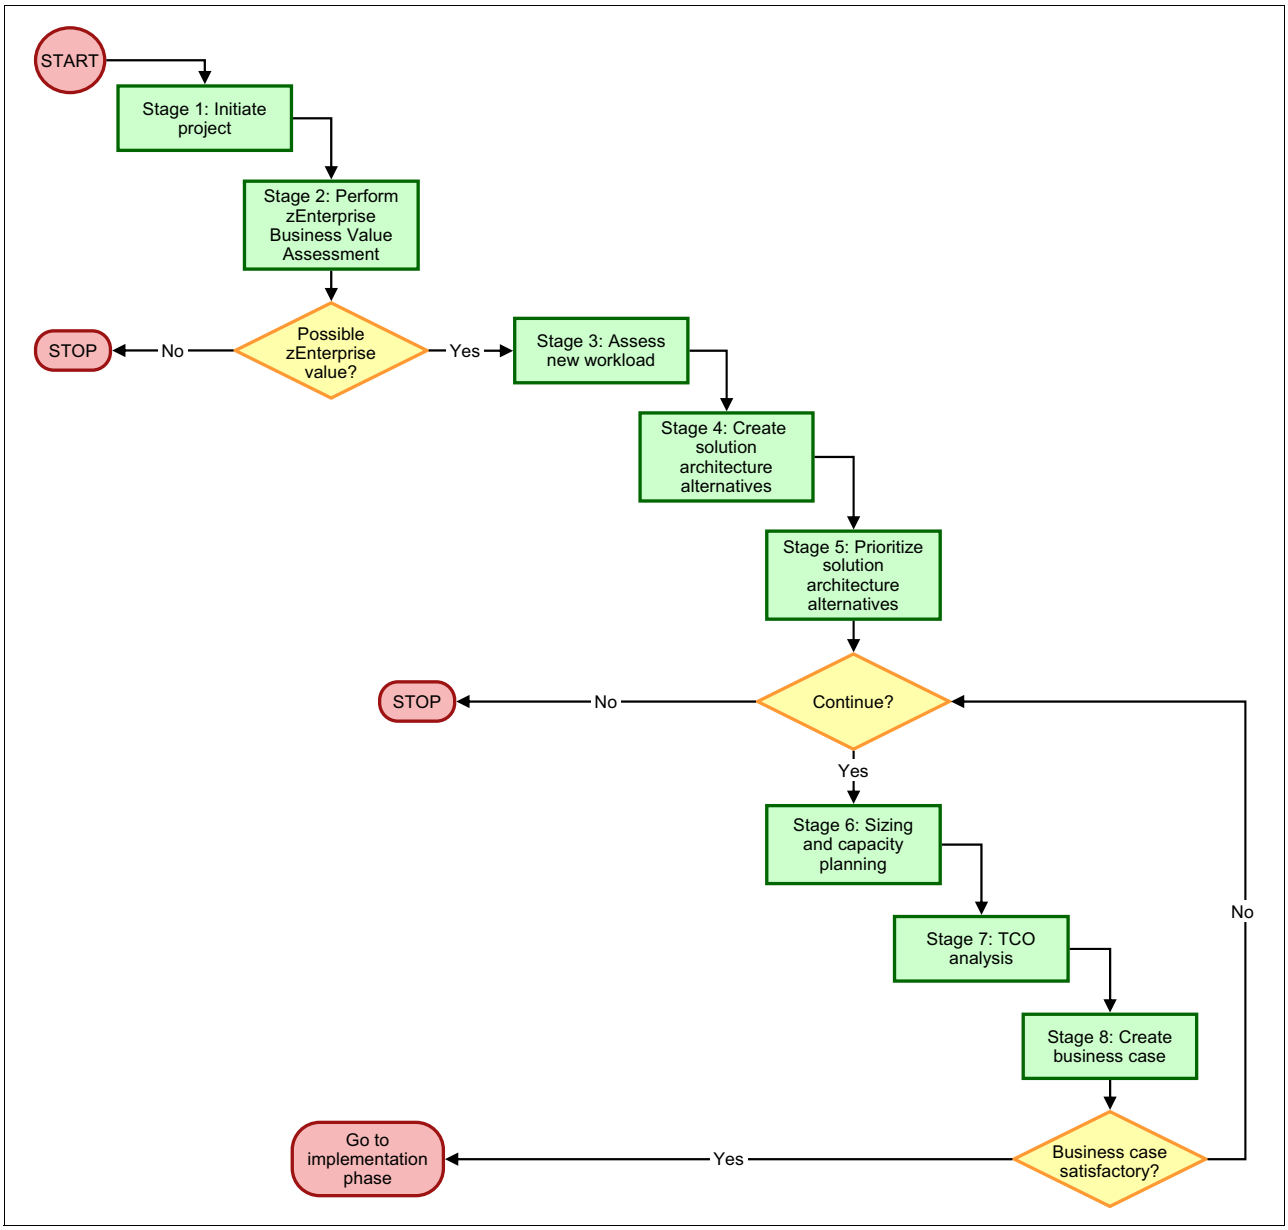

Figure 1-1 Overview of steps and decisions in the assessment phase

The primary decision points during the assessment phase are as follows:

- Stage 2: Perform zEnterprise Business Value Assessment
- Stage 5: Prioritize solution architecture alternatives
- Stage 6: Sizing and capacity planning
- Stage 7: Total cost of ownership (TCO) analysis

1.2 Process overview of implementation phase

To be sure that we thought through the process, we established a plan in advance of the implementation. We created an implementation plan and a technical design by using IBM Rational® Software Architect, which provided guidance to everyone working on the project. By using this design documentation, we configured the environment and prepared the system for test.

Figure 1-2 shows the steps and decisions of the implementation phase.

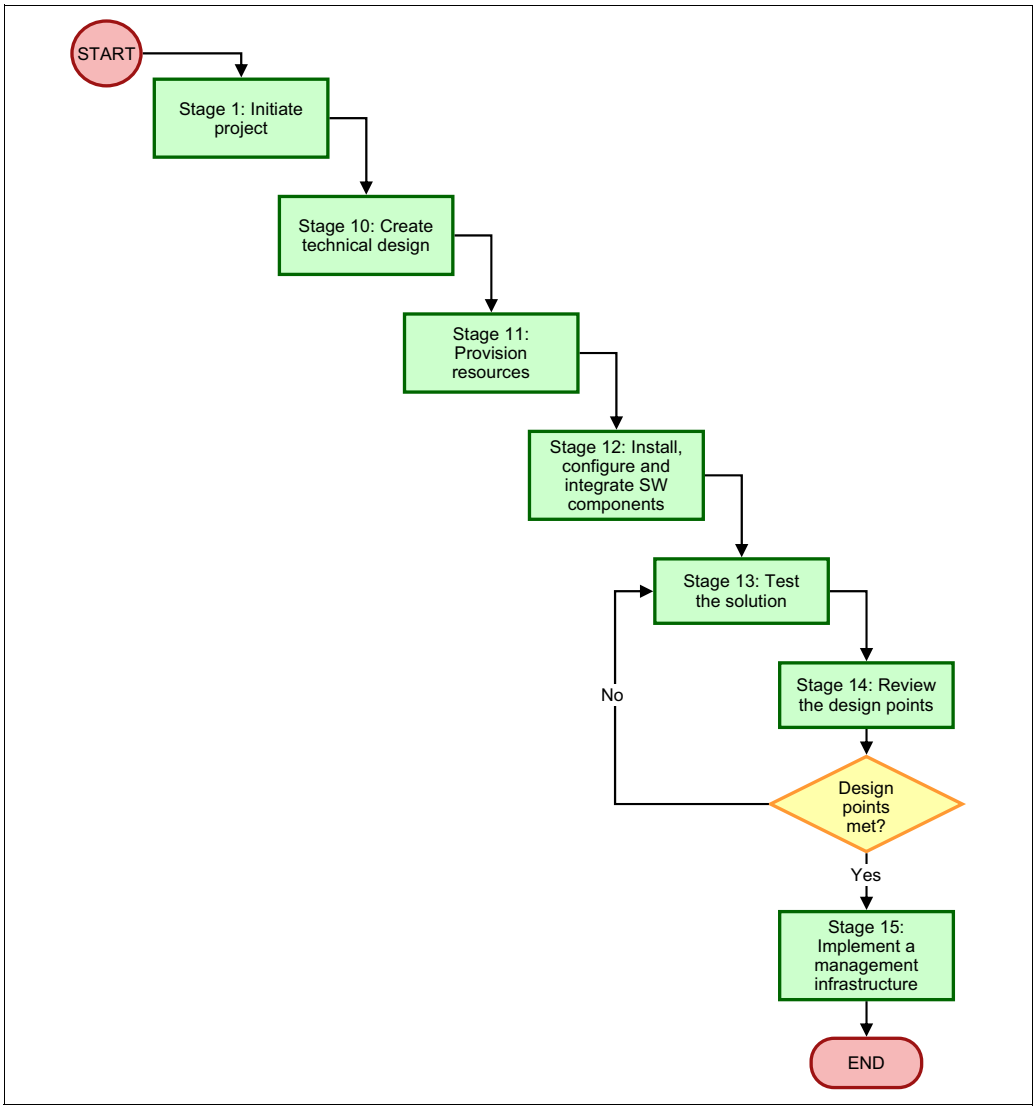

Figure 1-2 Overview of steps and decisions in the implementation phase

The numbering for the implementation stages is 9 - 15. In the following sections, we describe the steps in more detail. The numbering for the stages begins *after* stages 1 - 8, which were done in the assessment phase.[2]

1.2.1 Stage 9: Create implementation plan

This stage is the first activity in the implementation phase; the output is a plan that describes the tasks to be executed throughout the implementation phase and their dependencies.

1.2.2 Stage 10: Create technical design

The technical design described the environment to be implemented and configured at various levels. At a minimum the following information must be included:

- Physical architecture

  Describe the physical servers, network adapters, routers, switches, desktop computers, consoles, and so on.

- Network topology

  Describe the IP network, network addresses, gateways, and so on.

- Middleware topology

  Describe web servers, application servers, database servers, transaction managers, and so on. Especially important is to show how the middleware components integrate with each other.

Good diagrams are critical in stage 10.

1.2.3 Stage 11: Provision resources

Because we decided to deploy CPA on zEnterprise, we used the Unified Resource Manager to provision the resources. The first step we did was to ensure that connectivity existed between all the components of the zEnterprise, using both IBM HiperSockets™ and intraensemble data network (IEDN). Next, we provisioned the resources for Linux on System z, followed by configuration of capacity on the zBX for the POWER/AIX and Intel/Windows blades.

1.2.4 Stage 12: Install, configure, and integrate software components

This stage was probably the most time-consuming. All required middleware products had to be installed and configured from the beginning on the chosen platforms. In fact, for each product, we did the following tasks:

1. Planning and preparation of the installation and configuration
2. Basic installation
3. Configuration
4. Unit test and installation verification test (IVP)
5. Tuning and performance optimization

Most of this activity, however, is not specific to zEnterprise; it would be the same on other hardware by using the same chosen system.

[2] Described in *Using zEnterprise For Smart Analytics: Volume 1 Assessment*, SG24-8007 (not yet available)

1.2.5 Stage 13: Test the solution

After the infrastructure was set up, the end-to-end testing was started. We developed a test plan and test cases by using a varying number of users and requests to test the scalability of the solution and to determine whether any bottleneck exists. We used an established set of tools and monitors to run these tests.

1.2.6 Stage 14: Review design points

In the assessment phase, we made certain decisions on platform placement of the workload components. We also did a sizing and capacity planning study to estimate the required resources and amount needed (central processors, Integrated Facility for Linux, memory, disk space, and so on).

We decided to include a step to review the actual results and compare them with the decisions that were made in the assessment phase.

1.2.7 Stage 15: Implement a management infrastructure

After the applications were installed and running, we needed a framework that allowed us to monitor and manage the system and application components. We describe how to use a suite of IBM Tivoli® applications to establish a monitoring solution that provides a complete end-to-end view of the solution while running.

2

Implementation planning

Before we embarked on the deployment project, the first step was to create an implementation plan and assign tasks to the individuals who were involved in the project. This chapter provides an outline of the implementation plan and the dependencies.

2.1 Overview of the implementation steps

The following steps are a high-level overview of implementation:

1. Infrastructure provisioning
   a. Acquire new required hardware components.
   b. Install and connect new required hardware components.
   c. Configure the network.
   d. Provision new resources.
   e. Amend and reconfigure existing resources where required.
2. Middleware provisioning
   a. Install new software components.
   b. Configure new software components.
   c. Amend and reconfigure existing software components where required.
   d. Deploy the application.
3. Testing implementation
   a. Perform unit tests of the infrastructure components.
   b. Perform end-to-end tests of the infrastructure.
   c. Perform online and batch test scenarios with the application.
   d. Review aspects of non-functional requirements against test cases.
4. Results implementation

   Review components and adjust configuration as appropriate

2.2 More details of implementation steps

The high-level steps are captured in Table 2-1 through Table 2-4 on page 10. The first table represents the prerequisite tasks for establishing the development environment, which are described further in Chapter 3, “Technical design” on page 11 and Chapter 5, “Install, configure, and integrate the software components” on page 37.

Table 2-1 Infrastructure provisioning

Platform	Key steps	Description
z/OS	Verify DB2 for z/OS BDW settings	Verify the DB2 Data Warehouse tables, access permissions, and data volumes.
Linux on System z	Create Instance 1: Cognos BI Create Instance 2: SPSS Modeler Server	► Set up the BI environment, capable of supporting 2000 regular users, 1 TB per day volumes. I/O profile (multiuser configuration). This I/O profile (a multiuser configuration) has low CPU usage but high throughput. ► Set up AIX environment, capable of 300 power users. This CPU profile has high CPU (IFL) usage but low throughput, i.e. High CPU/IFL req/low throughput.
AIX	Create new analytics server.	Set up AIX environment for Cognos TM1 and Metrics Manager components.
Windows on zBX	Create new windows instance.	Set up Windows environment for development tools, with access to DB2 for z/OS.
Unified Resource Manager	Provision Unified Resource Manager resources.	Create virtual servers, hypervisors, Application Response Measurement (ARM). and performance policies.

Platform	Key steps	Description
Network	Establish new internal and external network configurations.	Set up new virtual LANs for access to each system and intrasystem communication that use intraensemble data network (IEDN) or intranode management network (INMN).
Security	Create IBM RACF® & LDAP user IDs.	Create user IDs for the Redbooks team and middleware product IDs and permissions.
Storage	Provision disk space.	Assign disk space from existing DS8800 estate to each environment.

Table 2-2 contains the Middleware implementation activities, which are described in detail in Chapter 5, “Install, configure, and integrate the software components” on page 37.

Table 2-2 Middleware provisioning

Component	Platform
Cognos BI Server	Linux on System z
Cognos BI Gateway	AIX
Cognos TM1 Server	AIX
Cognos Metrics Manager	AIX
Cognos Data Access Pack	AIX
SPSS Modeler Server	Linux on System z
SPSS Modeler Client	Windows 2008
WebSphere Application Server	Linux on System z & AIX
DB2 for Linux, UNIX, Windows (LUW)	AIX

Table 2-3 contains the testing implementation activities that are described in detail in Chapter 6, “Testing the application” on page 65.

Table 2-3 Test case implementation

Test requirements	Test cases	Notes
Systems testing	Security access Ping tests Tivoli agents views Unified Resource Manager access DB2 for z/OS access	Configuration tests to ensure availability of various services
Online user testing	Cognos BI, TM1, Metrics & SPSS use-case test	Log on; execute dashboards, BI reports, TM1 reports, Metric Studio scorecards and SPSS models
Batch user testing	Multiuser cognos BI, TM1, Metrics & SPSS use-case test	Log on; execute dashboards, BI reports, TM1 reports, Metric Studio scorecards and SPSS models
Non-functional testing	Quality of service testing: for example, availability, scalability and security	Proof points regarding high and continuous availability, scalability, and security across the enterprise components

Table 2-4 represents the follow-up activities to be implemented based upon the test results. These activities can involve system, middleware, or configuration changes. These results are explained in Chapter 7, “Review the design points” on page 73.

Table 2-4 Results implementation activities

Results	Component	Action
Systems design	Servers, blades, and LPARs	Review and amend configuration as required.
Middleware deployment	All software	Review software placement and redeploy.
Network design	Intranet and extranet services	Review network configurations
Middleware versioning	All software	Check software levels, fix packs, and so on.
Storage deployment	DS8800 server	Check usage, deployment, and I/O response times.
Integration design	All software and platforms	Review “hybrid” compatibility across the middleware and platform components.
Data integrity	Local files, system and user data	Verify expected results, and quality and integrity of data.
Quality of service	All	Verify that application deployment and testing meets the original goals.

3

Technical design

We are implementing the IBM Customer Profitability Analytics (CPA) blueprint solution introduced in Chapter 1, “Introduction” on page 1. With the IBM CPA solution, you can maximize the value of your customer base by understanding the interconnected factors that influence profitability through a single, at-a-glance view.

3.1 Application architecture

The IBM Customer Profitability Analytics solution consists of three major components on an integrated software platform, which provides the bank with credible, insightful, and actionable profitability information to drive better business outcomes. The components are a profitability calculator, a reporting capability, and a predictive analysis.

3.1.1 Profitability calculator

The first and foundational component of the solution is to provide the bank with a profitability calculator, which can efficiently produce customer account profitability statements. The solution calculates and applies best-practice management-accounting theory to calculate measurements such as: cost of funds, risk-adjusted return on capital (RAROC), provision for losses, non-interest expenses using both activity-based costing and allocation methodology. It also provides the calculation of many key performance ratios (RAROC, return on equity, expense ratio, loan yields, deposit rates, and return on assets).

Cognos TM1 and Metric Studio provide an advanced analytics design and delivery toolkit to deliver these calculations.

3.1.2 Reporting functionality

The second major component is a robust reporting capability that exposes the profitability information to the correct business user at the correct time to enhance business outcomes. The solution includes executive dashboards and scorecards for the bank's major line of businesses and also financial and marketing analyst dashboards. The reporting capabilities include many functions, highlighting key performance indicators pertinent to the business user's role.

The reporting capability is delivered through the IBM Cognos 10 technology software platform.

3.1.3 Predictive analytics

The third major component is the ability to apply predictive analytics to the historical profitability information ad customer attribute data to predict what will happen next. The information you captured and calculated can now be analyzed by predictive models that help you understand and anticipate what customers want and will do next. These models use advanced analytics to uncover patterns in the data and then predict the likelihood of future events. For example, you might use predictive analytics to segment your customers, based on profitability, and then target each customer within a segment with the most relevant recommendation. This approach can increase customer profitability and improve customer loyalty and satisfaction.

IBM SPSS predictive analytics technology provides a comprehensive tool to support this requirement.

3.1.4 Summary of application environment

The IBM CPA blueprint consists of a number of models, developed by using a common suite of IBM products. These principal components are as follows:

- IBM Cognos TM1 Server (9.5.0 or greater)
- IBM Cognos 10.1 Business Intelligence Server
- IBM Cognos 10.1 Metrics Server
- IBM DB2 Database 9.5 (Express-C or greater)
- IBM SPSS Modeler 14.2

Figure 3-1 shows how all the components come together into a logical application architecture.

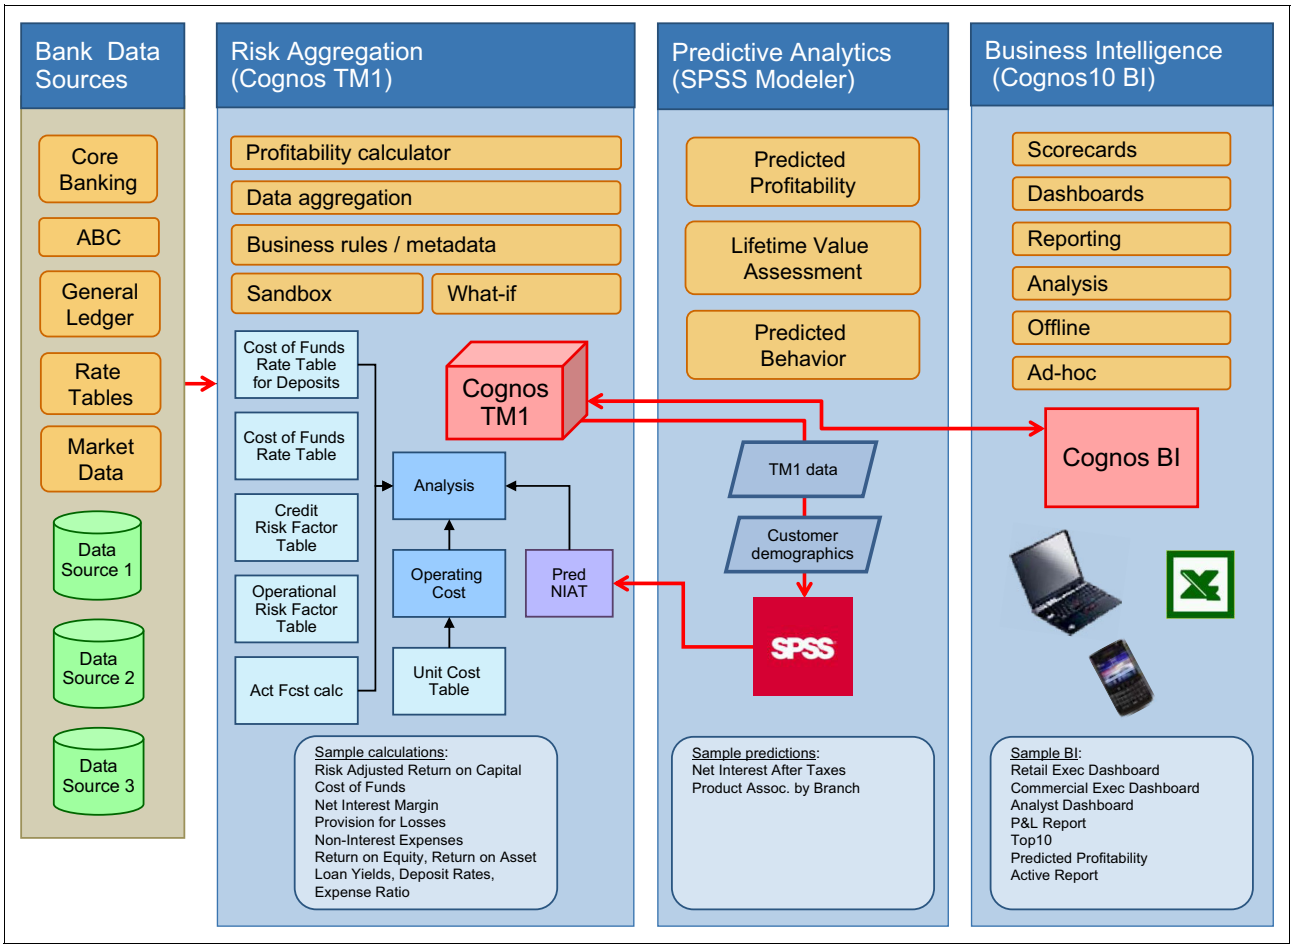

Figure 3-1 Customer Profitability Analytics blueprint logical architecture

3.2 Physical environment

We are implementing the CPA solution onto an existing zEnterprise system, with an attached IBM zEnterprise BladeCenter Extension (zBX). Figure 3-1 on page 13 shows the configuration of the base infrastructure from which we allocate our resources.

More detail about the zEnterprise system and zBX is in the following resources:

- *IBM zEnterprise 196 Technical Guide*, SG24-7833
- zEnterprise BladeCenter Extension web page:

  http://www.ibm.com/systems/z/hardware/zenterprise/zbx.html

Table 3-1 Hybrid infrastructure

Component	Model	CPUs/Cores	Memory
zEnterprise	z196 M80	41 CPUs	600 GB
zBX	002	N/A	N/A
Power processor-based blade	PS701	1/8	32 GB
Intel processor-based blade	HX5	2/16	64 GB

A previous activity, described in *Using zEnterprise For Smart Analytics: Volume 1 Assessment*, SG24-8007, evaluated architectural alternatives for hosting the application components of the CPA blueprint. The conclusion of that exercise is that we should place the components as follows:

- Data marts:
  - Cognos BI Reports: DB2 for z/OS with IBM DB2 Analytics Accelerator
  - Predict profitability: DB2 for z/OS with IBM DB2 Analytics Accelerator
  - CPA Metric Studio: DB2 for z/OS with IBM DB2 Analytics Accelerator

  More information about IBM DB2 Analytics Accelerator can be found at:

  http://www.ibm.com/software/data/db2/zos/analytics-accelerator/
- Software components:
  - Cognos 10 BI: Linux on System z
  - Cognos TM1 with Cognos Data Access Pack: zBX connected POWER blade running AIX
  - SPSS Modeler Server: Linux on System z

The following components were not in scope, but were included as administration tools:

- Cognos 10 Framework Manager: zBX connected x86 blade running Microsoft Windows
- SPSS Modeler Client: zBX connected x86 blade running Microsoft Windows

Product availability: During our implementation of the CPA blueprint, we realized that not all of the required Cognos BI components are available on Linux on System z, and therefore, we adjusted the implementation to match product availability. Details of our findings are in Chapter 5, “Install, configure, and integrate the software components” on page 37.

See the following figure, table, and chapter:

- Figure 3-2 shows how we map the operating system (OS) and hardware requirements for these components to our zEnterprise system.
- Table 3-2 on page 16 shows the resources that are allocated for each.
- Chapter 4, “Provisioning our resources” on page 25 describes how these resources are provisioned in the hybrid environment.

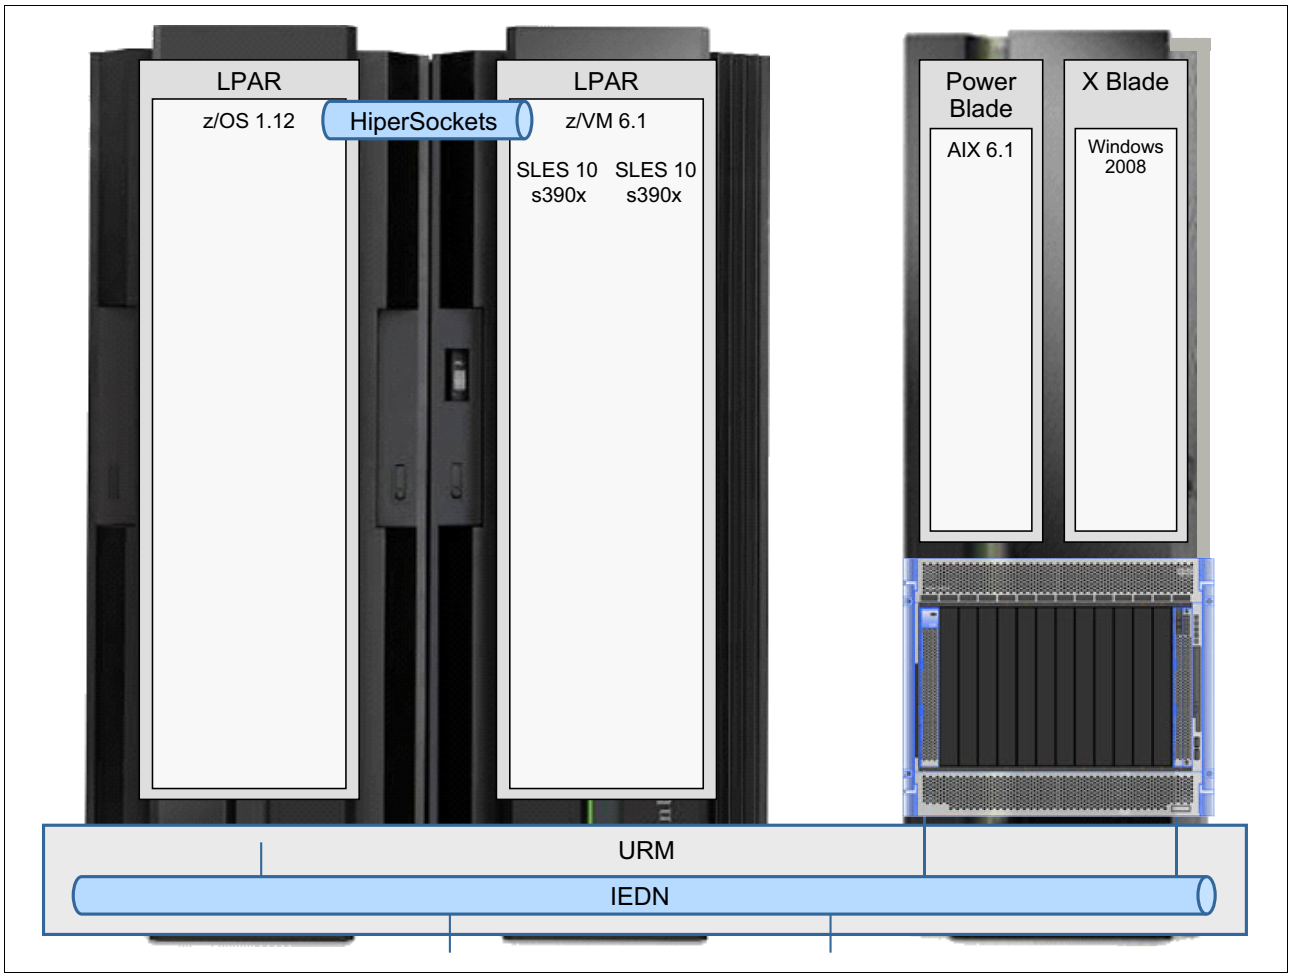

Figure 3-2 Physical Architecture for Hybrid Workload Implementation (CPA OS)

Table 3-2 Resource allocation for each software component

Software component	Image name	Hardware placement	Operating system	Allocated processors or cores	Allocated memory
Data Marts	BDW01DB	zEnterprise 196	z/OS 1.12	5 CPs (shared)	32 GB
Cognos BI	cognbi	zEnterprise 196	SLES[a] 10 s390x	2 IFLs[b]	8 GB
Cognos TM1	cogntm1	zBX POWER blade	AIX 6.1	6 cores	8 GB
SPSS Modeler Server	cognspss	zEnterprise 196	SLES 10 s390x	2 IFLs	8 GB
Cognos 10 Framework Manager and SPSS Modeler Client	cogndev	zBX Intel Blade	Windows Server 2008	4 cores	4 GB

a. SUSE Linux Enterprise Server (SUSE)
b. Integrated Facility for Linux (IFL)

3.3 Network architecture

The network architecture of the new system provides multiple networks with differing characteristics for bandwidth, security, and reliability.

Because IBM HiperSockets technology is available only between systems on the z196 server, we configure the middleware on Linux on System z to use this network when calling middleware running on z/OS. The use of HiperSockets over other network options provides the best possible response time and allows DB2 data on z/OS to appear “in close proximity” to the Linux on System z servers.

Communication between servers on the z196 and the zBX blades occurs over the 10 GbE IEDN.

Because some portions of the mainframe will be used for services other than CPA, having an isolated network for testing is preferable. In this manner, we can simulate a large number of users without affecting production network performance. When not used for testing, this network can perform other administrative functions, such as backups, that might otherwise create a network bottleneck.

These networks plus the network for intranet users are depicted in Figure 3-3.

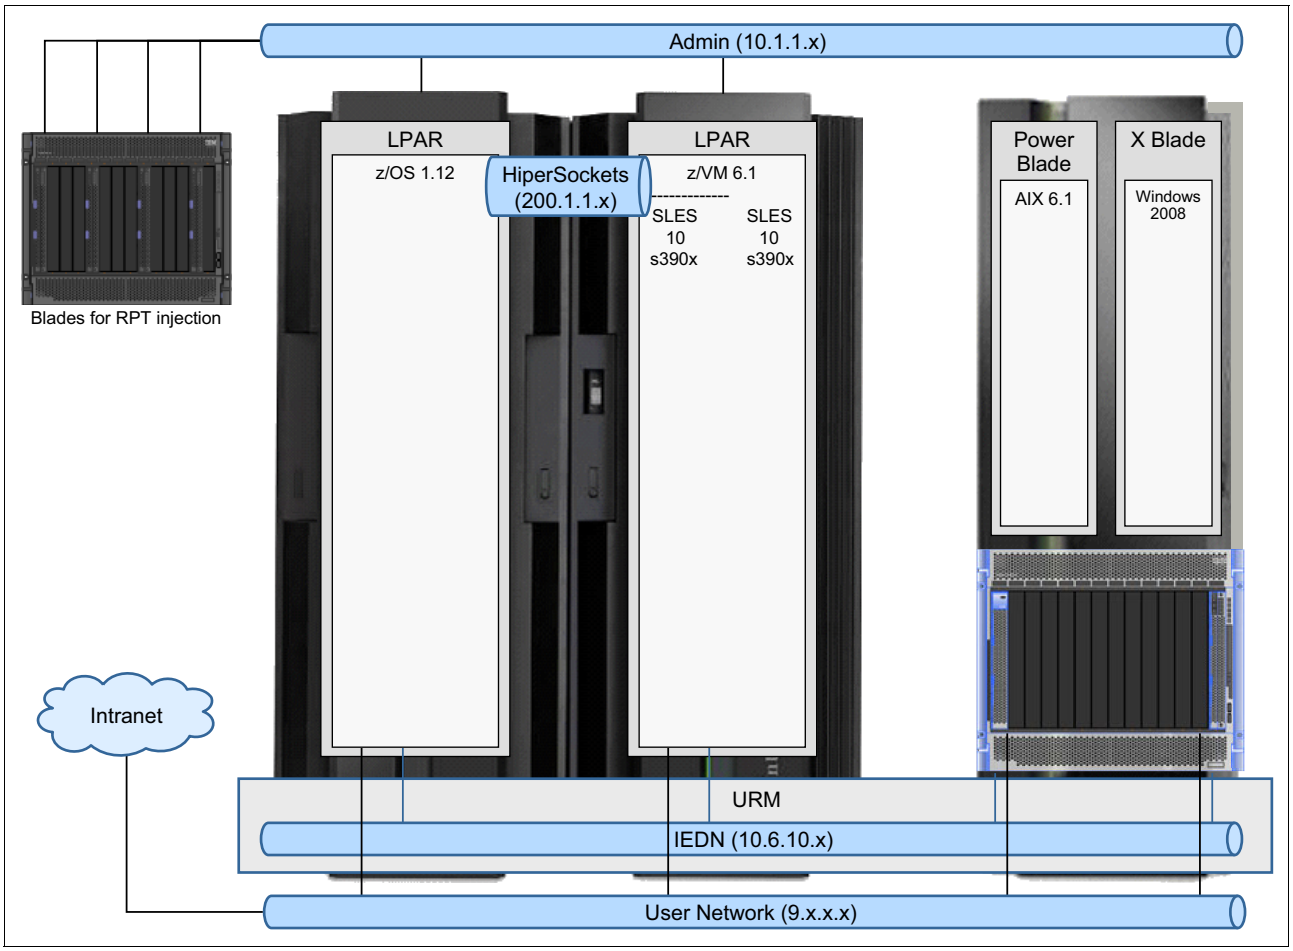

Figure 3-3 Network architecture for hybrid workload implementation (CPA network)

Table 3-3 lists the IP address for each image on the each of the networks, plus an alias that can be used to ensure that network traffic is directed to a specific network (where desired). No alias is required for the user network, because the image name is used for that network.

Table 3-3 IP address or alias for each image on each network

Image name	HiperSockets	IEDN	Admin	User
BDW01DB	200.1.1.20 z/OS1hyp	10.6.10.20 BA01	10.1.1.20 BA01_priv	9.212.128.20
cognbi	200.1.1.44 cognbi_hyp	10.6.10.44 cognbi_iedn	10.1.1.44 cognbi_priv	9.212.128.44
cogntm1	N/A	10.6.10.36 cogntm1_iedn	Not allocated	9.212.135.36
cognspss	200.1.1.78 cognspss_hyp	10.6.10.78 cognspss_iedn	10.1.1.78 cognspss_priv	9.212.130.78
cogndev	N/A	10.6.10.38 cogndev_iedn	Not allocated	9.212.135.38

In our network, a Domain Name System (DNS) is not yet established, so we distribute a hosts file to the systems that will require it (all of them except BDW01DB), as shown in Example 3-1.

Example 3-1 Hosts file for name resolution

```
# Hosts file for CPA Deployment
#
127.0.0.1       localhost
# z/OS image for DB2 datamarts
9.212.128.20    BDW01DB
200.1.1.20      z/OS1hyp
10.6.10.20      BA01
10.1.1.20       BA01_priv
# Linux on system z image for Cognos 10 BI
9.212.128.44    cognbi
200.1.1.44      cognbi_hyp
10.6.10.44      cognbi_iedn
10.1.1.44       cognbi_priv
# AIX image for Cognos TM1
9.212.135.36    cogntm1
10.6.10.36      cogntm1_iedn
# Linux on system z image for SPSS Modeler
9.212.130.78    cognspss
200.1.1.78      cognspss_hyp
10.6.10.78      cognspss_iedn
10.1.1.78       cognspss_priv
# Windows 2008 image for development tools
9.212.135.38    cogndev
10.6.10.38      cogndev_iedn
```

3.4 Middleware topology

Figure 3-4 shows how we map the various middleware components of the CPA blueprint onto the physical architecture defined in the previous section. Note that only those components on z/OS and Linux on System z will be able to communicate through HiperSockets. All other intraensemble communications will use IEDN.

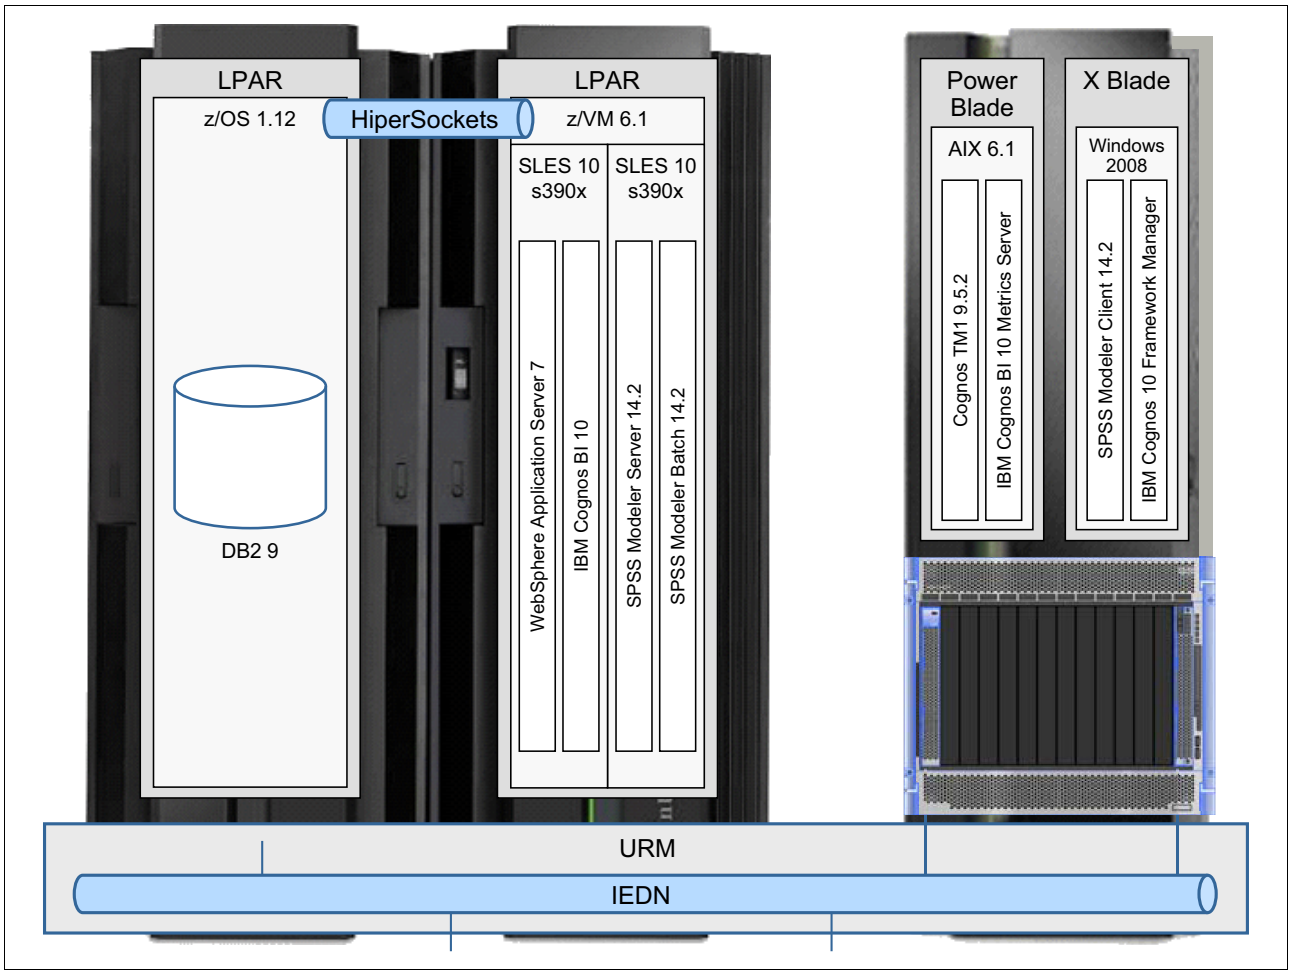

Figure 3-4 Physical architecture overlaid with middleware placement (CPA model)

3.4.1 Data management or data source connectivity

While implementing the CPA blueprint, we integrate with an existing Showcase environment, with a large amount of existing DB2 data on z/OS.

Repositories

Each individual middleware component has its own repository for configuration information. In addition, some have additional shared repositories. The placement of each of these is presented in the following sections.

Cognos BI Configuration Management Database

Cognos BI stores its configuration and content data in a Configuration Management Database (CMDB). In general, IBM suggests that the CMDB be hosted on a server separate from the Cognos BI server itself. Because our environment already has a supported database platform (DB2 10 for z/OS), this is where we will allocate our CMDB tables.

Metric Studio CMDB

Metric Studio Manager requires a Configuration Management Database separate from the one used for Cognos BI; however, DB2 for z/OS is not a supported platform at this time. For this reason, we place the Metric Studio CMDB in DB2 UDB 9.5 for AIX on a IBM POWER 7 zBX blade. Our preference was to use DB2 for z/OS, if it had been supported.

DB2 for z/OS

The DB2 10 for z/OS environment we plan to use is already installed and is supporting the core banking, payments, and other banking services. There are separate subsystems for the OLTP database, CB01, and the warehouse tables, DW01. The content store will be located next to the warehouse tables.

3.4.2 SPSS Modeler Server

IBM SPSS Modeler is a powerful, versatile data mining workbench that helps you build accurate predictive models quickly and intuitively, without programming. The SPSS Modeler Server consists of the Modeler itself, the SPSS Data Access Pack, a batch function, the Modeler Client, and a License Server.

The assessment phase of this project, described in *Using zEnterprise For Smart Analytics: Volume 1 Assessment*, SG24-8007, concludes that the SPSS Modeler Server should be placed on Linux on System z, and the SPSS Modeler Client should be placed on a zBX blade running Windows. Because no mention is made of the Licensing Server, we decided to put it on the AIX Server. The Licensing server was supported only on AIX or Windows, and because we did not place any production workload on the zBX blade running Windows, this choice seemed to be the best.

Figure 3-5 is a logical diagram that shows placement of the SPSS functions. Additional information about IBM SPSS is at the following site:

http://www.ibm.com/software/analytics/spss/

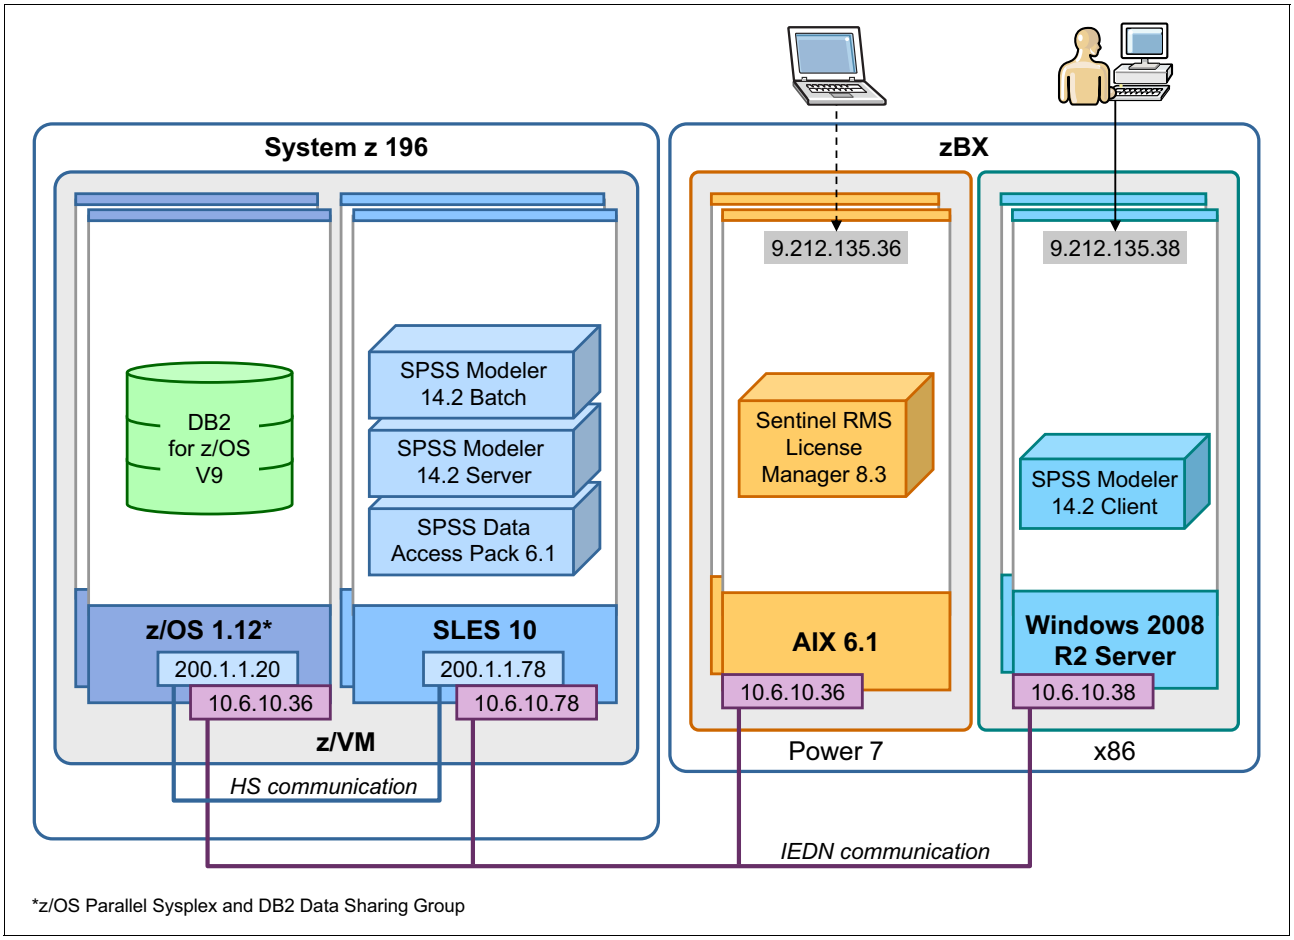

Figure 3-5 IBM SPSS Modeler Logical Architecture

3.4.3 Cognos 10 BI

IBM Cognos BI provides a unified workspace for business intelligence and analytics that an entire organization can use to answer key business questions. Components of IBM Cognos BI include IBM Cognos Business Insight, IBM Cognos Report Studio, IBM Cognos Query Studio, IBM Cognos Metric Studio, and IBM Cognos Framework Manager, among others.

The IBM Cognos Platform delivers the capabilities to manage business intelligence applications with centralized, web-based administration. It provides a complete view of system activity and also system metrics and thresholds so that organizations can resolve potential issues before there is a business impact.

The IBM Cognos platform is built on a web-based service-oriented architecture (SOA) that is designed for scalability, availability, and openness. This n-tiered architecture consists of three server tiers:[1]

- Web tier: The IBM Cognos gateway provides user session connectivity to IBM Cognos BI applications.
- Application tier: This tier consists of the IBM Cognos Dispatcher, the IBM Cognos Report Server, and the IBM Cognos Content Manager.
- Data tier: A relational, online analytical processing (OLAP), or other repository is where application or business data is stored.

Figure 3-6 shows the logical architecture for our Cognos BI components. The web tier is split across a Linux for System z and an AIX server. We find that this split is necessary, reasons for which are described in Chapter 5, “Install, configure, and integrate the software components” on page 37. The application tier is also split across these two platforms, and the data tier adds z/OS to the others.

We use WebSphere 7 to host the BI components on both AIX and Linux for System z. However, a 32-bit server is required for the metrics function on AIX. On Linux for System z, all functions are compatible with 64-bit WebSphere.

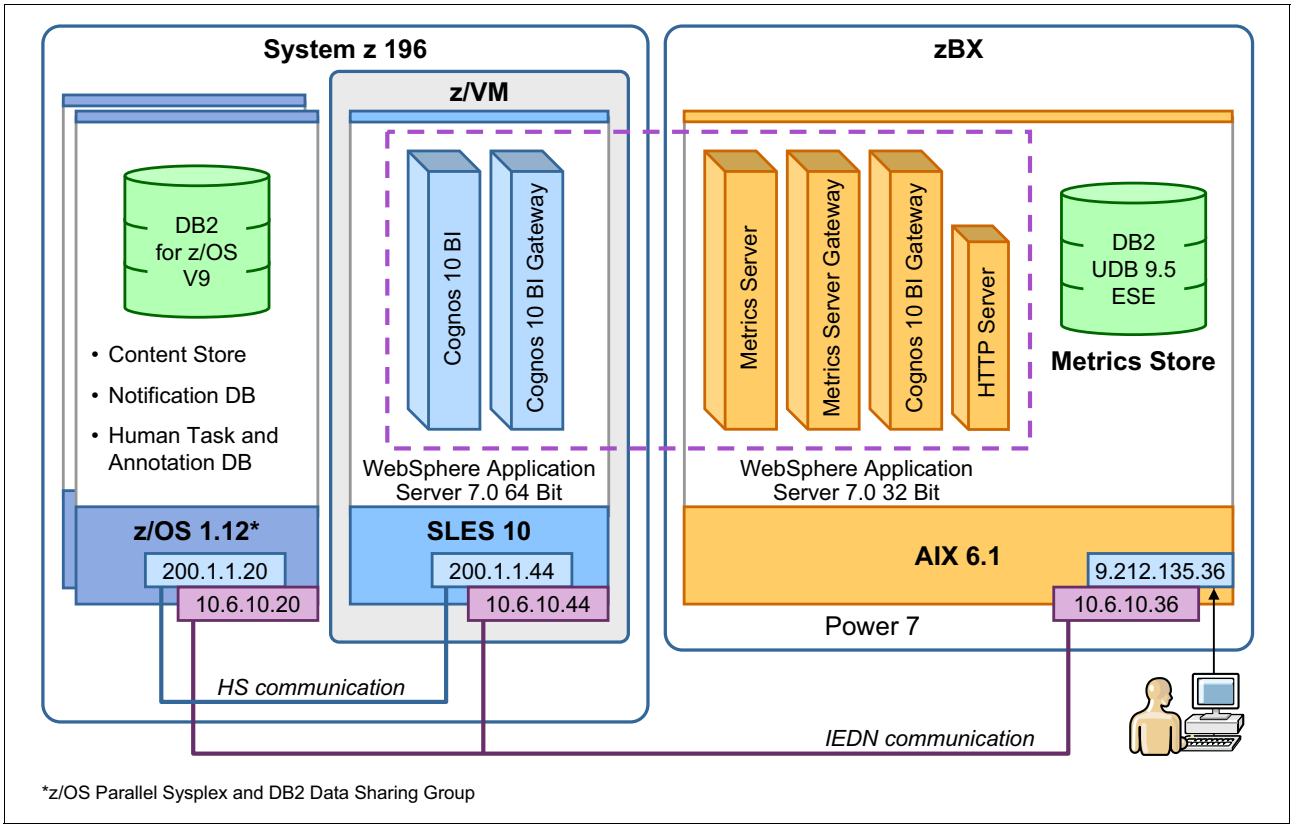

Figure 3-6 Cognos 10 architecture

[1] “Cognos Jump-Start” by Ann Jackson and Dean Browne; http://bit.ly/UPbK2C

3.4.4 Cognos TM1

IBM Cognos TM1 is enterprise planning software that provides a complete and dynamic environment for developing timely, reliable, and personalized forecasts and budgets. It consists of a multidimensional OLAP engine, and model design client software. See Figure 3-7.

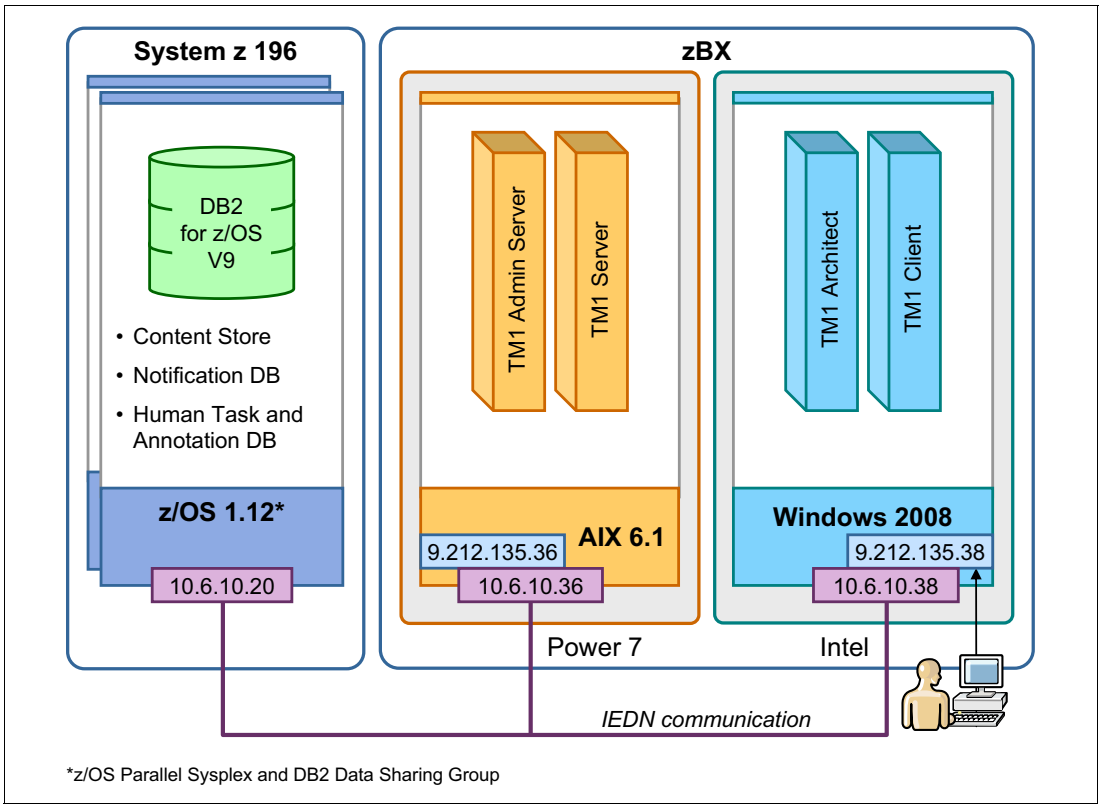

Figure 3-7 Cognos TM1 logical architecture

3.5 Availability aspects

This current architecture is not designed for high availability. However, several easy measures can help improve availability.

The Cognos 10.1 BI servers run under WebSphere Application Server. WebSphere Application Server can easily scale, both vertically (within the server) and horizontally (across multiple servers). Multiple Cognos BI and metrics servers can be added to increase both availability and scalability.

Cognos TM1 has no built-in cluster awareness, but may be configured by using products such as Tivoli System Automation or High Availability Cluster Multi-Processing (IBM HACMP™) to provide a higher level of availability.

3.6 Summary

We describe a technical design that details a hybrid environment for this business analytics workload. The design begins with a physical view, continues to network, OS and middleware.

Implementing this package across a range of operating systems, running on multiple processors, required greater understanding of the components and supported interfaces than originally planned. For example, Cognos Metric Studio is currently supported only within a 32-bit WebSphere run time; DB2 for z/OS is not a supported repository. Therefore we had to install DB2 UDB for AIX.

However, despite these technical glitches, the components were installed and were able to communicate between each other. The overall IBM CPA package was implemented across the products and the user reporting interface could deliver dashboards, scorecards, and predictive analytics using each of the environments seamlessly.

4

Provisioning our resources

This chapter describes the hybrid environment for this project. We describe the design and implementation process of the total environment, including network, storage, virtual servers, and Unified Resource Manager.

4.1 Our hybrid environment

Figure 4-1 shows the hybrid environment in our project.

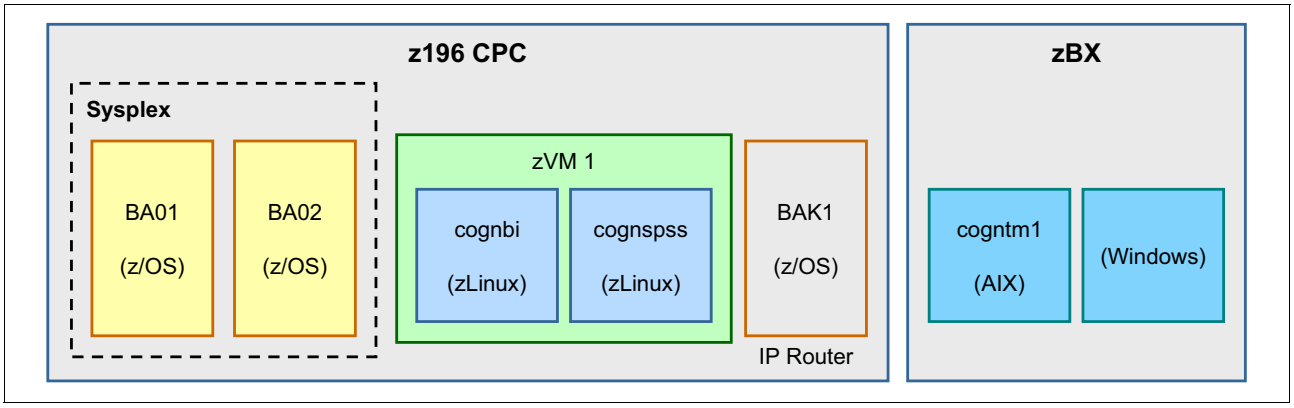

Figure 4-1 Hybrid environment

The following details apply:

- A z/OS Parallel Sysplex with two members (BA01 and BA02) is used to hold major data, such as IBM Cognos BI content store and user data source in DB2 on z/OS. The z/OS is version 1.12 and the DB2 on z/OS is V9. The z/OS Parallel Sysplex is an existing environment that we do not need to provision.
- z/OS system (BAK1) works as an IP router for external access to virtual servers that are running on blades in zBX. This z/OS also is an existing environment.
- An IBM z/VM® Version 6 Release 1.0, service level 1003 (64-bit), is running on z196. On the existing z/VM system, we need to create two more Linux on System z images to support IBM Cognos BI and SPSS software respectively. The Linux version is SUSE Linux Enterprise Server (SLES) 10 SP3 (64-bit).
- On the zBX side, we need to create a virtual server on a POWER blade. This AIX will be used to run Cognos TM1, Metric Manager, and Metric Studio. The AIX version is V6.1.0.0 SP 05 (64-bit).
- We need to create a virtual server on an x blade and install a Windows system. This configuration will be used to run an SPSS client and Cognos Metrics Manager. The Windows version is Windows Server 2008.

So, we must prepare two Linux instances, plus one AIX instance on a POWER blade, and one Windows instance on an x blade for this project. Next, we describe the detailed configuration design and process steps for network, storage, and virtual servers.

4.2 Configuring the network

This section describes the network infrastructure to support the Smarter Banking Showcase business analytical workload.

4.2.1 Types of networks

Figure 4-2 shows the network infrastructure of our project environment.

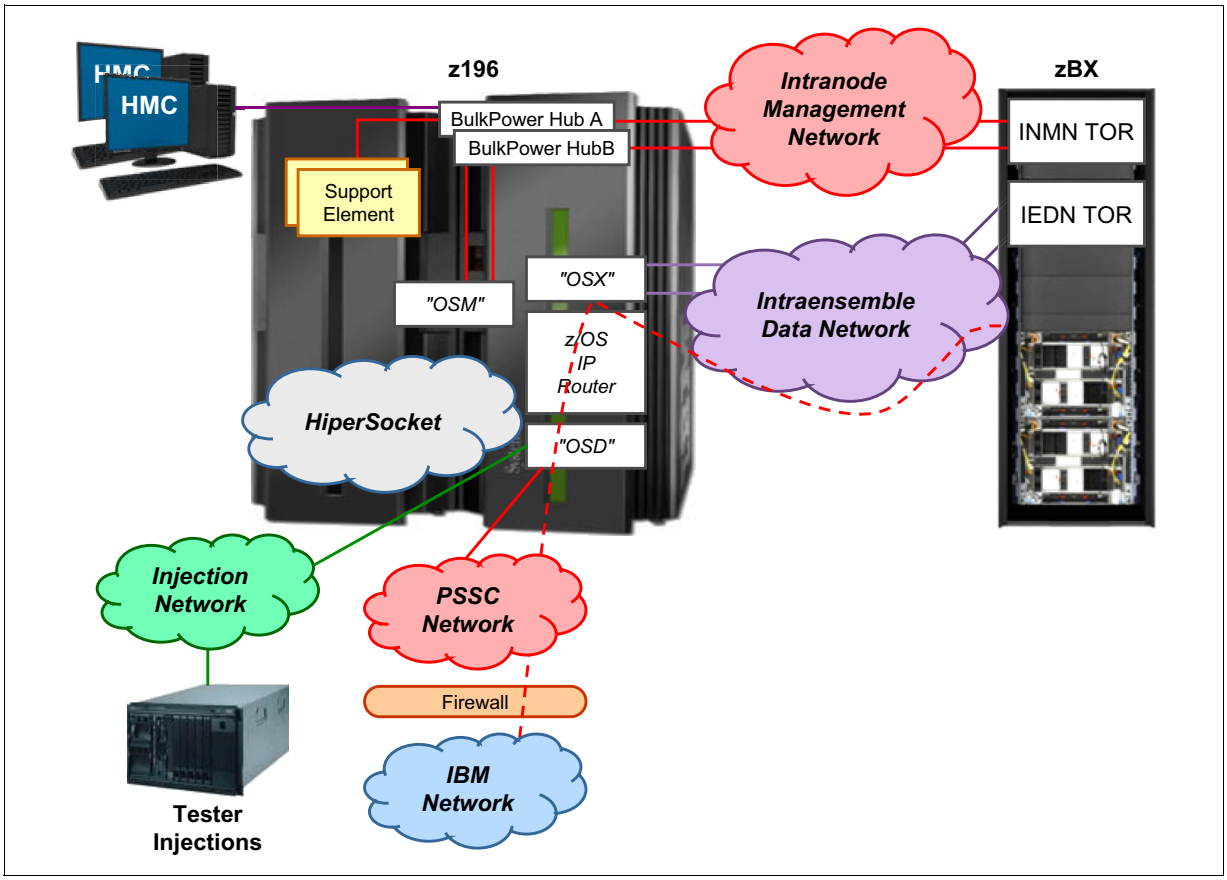

Figure 4-2 Network infrastructure

Four types of networks are in our project environment:

- Intranode management network (INMN)
- Intraensemble data network (IEDN)
- HiperSockets
- Customer management data network

Intranode management network (INMN)

An INMN is required for platform management within a node. The network allows the HMC to communicate to the hypervisors within the managed ensemble.

An INMN requires two ports from two separate OSA-Express3 1000BaseT Ethernet adapters (ensuring redundancy), connected to the bulk power hub (BPH) ports. BPH ports also provide the connectivity from the BPH to the top-of-rack (TOR) switches for the INMN. The INMN supports only the running IPv6 protocols. IPv6 link local addresses are assigned by the system.

Intraensemble data network (IEDN)

An IEDN is sometimes known as a private data network. An IEDN is visible to the operating systems that are running in virtual servers. It is required for application data communications within an ensemble that contains one or more IBM zEnterprise BladeCenter Extension (zBX) modules.

The IEDN is a 10 GbE flat data link (Layer 2) network. Virtual servers and optimizers connect to the IEDN for communication in support of customer workloads. An IEDN is part of the ensemble, and the HMC manages the IEDN through its user interface.

In our project environment, the z196 has two OSA-Express3 10 GbE adapters (CHPID type OSX) connected with zBX by using top-of-rack (TOR) switches. All data communication between POWER blade, x blade, and z196 is based on IEDN. Meanwhile the virtual servers for this project are isolated into groups on the physical network by defining virtual local area networks (VLANs).

HiperSockets

HiperSockets is a technology that provides high-speed Transmission Control Protocol/Internet Protocol (TCP/IP) connectivity between servers within a System z. This technology eliminates the requirement for any physical cabling or external networking connection among these *virtual* servers. It works like an internal local area network (LAN). HiperSockets is useful if you have a high volume of data that is flowing among these *virtual* servers.

HiperSockets uses internal Queued Direct Input/Output (iQDIO) at memory speeds to pass traffic among these *virtual* servers. HiperSockets is a Licensed Internal Code (LIC) function that emulates the Logical Link Control (LLC) layer of an OSA-Express QDIO interface. In our project environment, HiperSockets was set up for data communication inside the z196.

Customer-managed data network

You use your own data network for external network communication. You provide and manage the customer data network. It is not managed by an HMC or its network virtualization function tasks.

In our project environment, our z196 is connected to a customer data network named PSSC through an OSD OSA channel path. The zBX is connected to a customer data network through a z/OS IP router running in the z196. A customer can access the project environment from the IBM network, and then go to the PSSC lab environment.

Also, there is an internal network named Injection network. This network is only for data communication between injection servers that drive workloads. Customers cannot access this network externally.

4.2.2 Configuring INMN

Figure 4-3 shows the INMN configuration in our project environment.

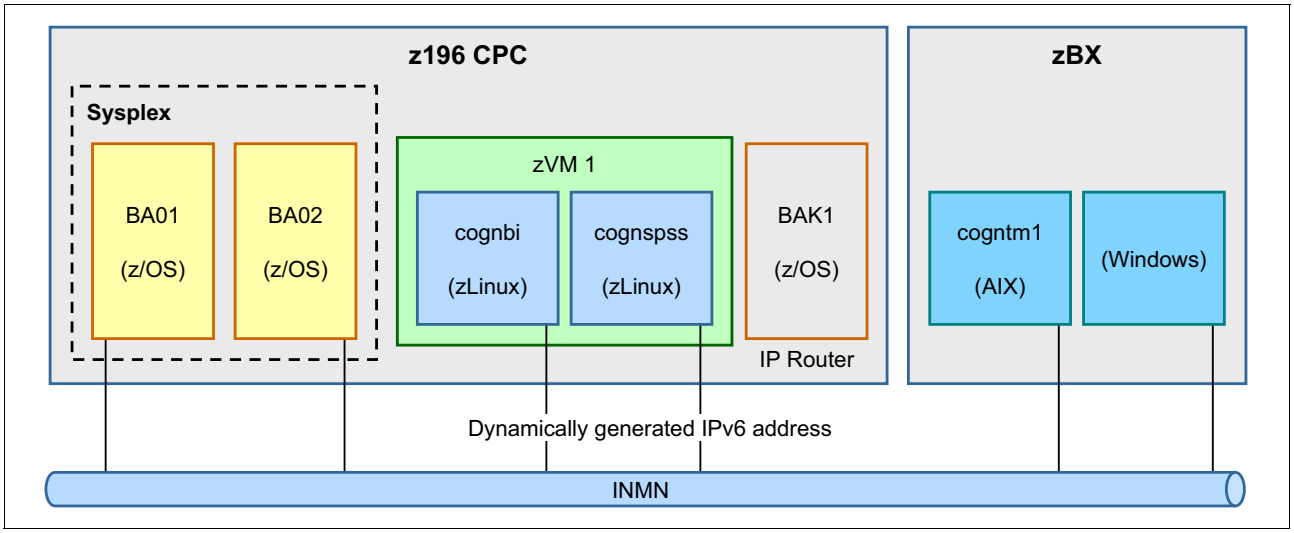

Figure 4-3 INMN configuration

Our z196 and zBX comprise an ensemble node. All system images within this ensemble node are connected to the INMN network to provide management and monitor information. They provide point-to-point connection to Unified Resource Manager and Support Element (SE), which cannot be accessed by any other application except Unified Resource Manager.

INMN supports running only IPv6 protocols. So for each system image, we need to enable IPv6 protocol. Then IPv6 link local addresses are dynamically assigned by the system.

4.2.3 Configuring IEDN

Figure 4-4 shows the IEDN configuration in our project environment.

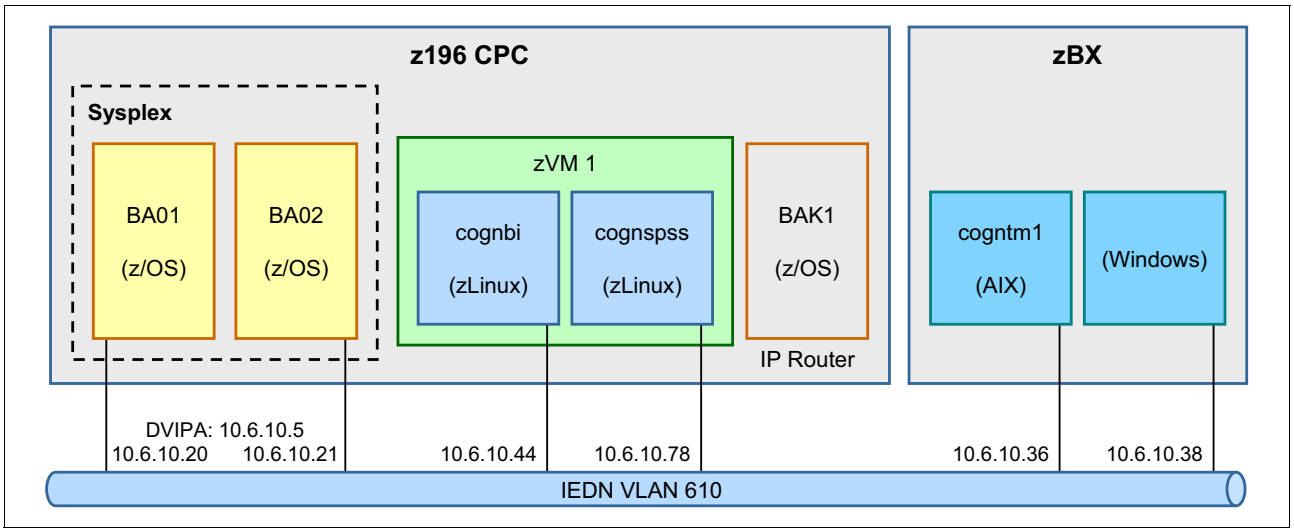

Figure 4-4 IEDN configuration

All data communication between z196 and zBX is based on IEDN. Each system image is defined with IEDN interface and an IP address for data communication. These IP addresses are used only for data communication within IEDN while cannot be accessed from external.

These system images are isolated into a VLAN subset for secure data communication.

One dynamic virtual IP address (VIPA) is configured for the z/OS Parallel Sysplex used in the project. If a virtual server that is running on zBX wants to communicate with z/OS Parallel Sysplex, it can point to the dynamic VIPA (DVIPA) address to get network redundancy.

4.2.4 Configuring HiperSockets

Figure 4-5 shows the HiperSockets configuration in our project environment.

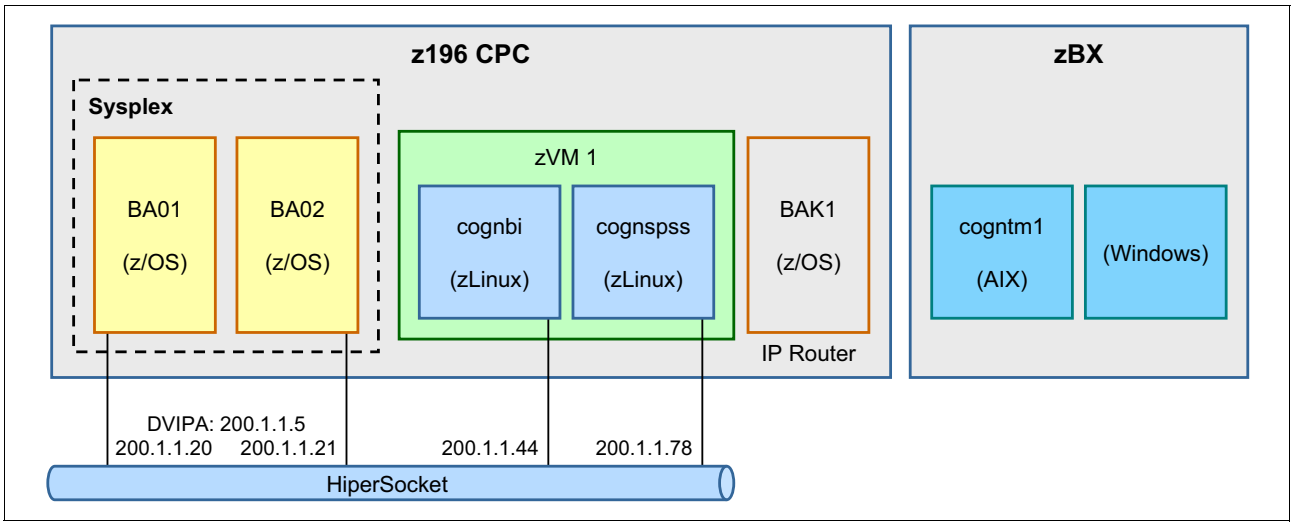

Figure 4-5 HiperSockets configuration

Each z/OS member and Linux virtual server is configured with one HiperSockets IP address. Also one dynamic VIPA is configured for the z/OS Parallel Sysplex used in the project. The detailed configuration process of HiperSockets for z/OS and Linux on System z is in *HiperSockets Implementation Guide*, SG24-6816.

All data communication in the z196 system will go through the HiperSockets network. Meanwhile, if Linux wants to communicate with z/OS Parallel Sysplex, it can point to its DVIPA address to get network redundancy.

4.2.5 Configuring customer-managed data network

Figure 4-6 shows the customer-managed data network configuration in our project environment.

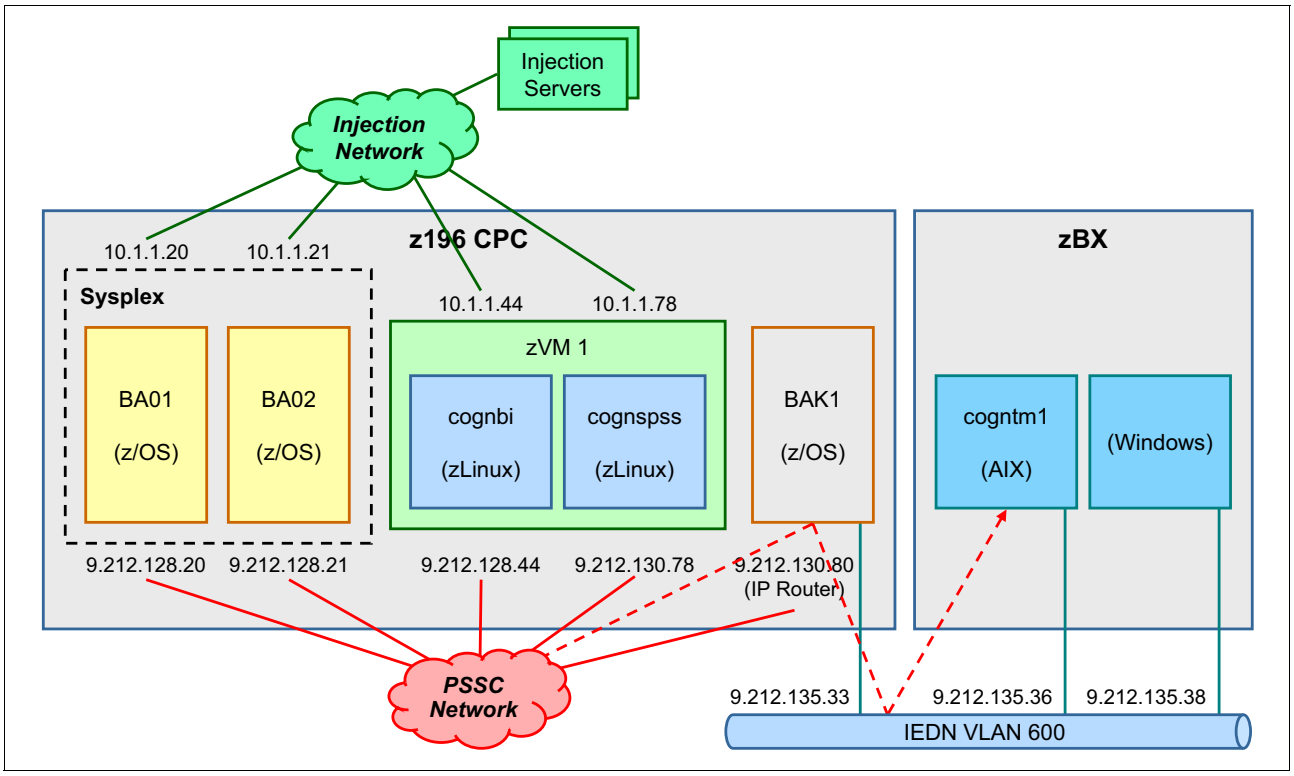

Figure 4-6 Customer-managed data network configuration

There are two customer-managed data networks:

- PSSC network is the trunk network infrastructure of Montpellier Products and Solution Support Center (PSSC). A customer can access the PSSC network through the IBM network by having PSSC firewall authority.

  z196 is connected to PSSC network through OSD OSA channel path. Customer can directly access z/OS and Linux through the IP address of the PSSC network.

  zBX is connected to customer data network through a z/OS IP router running in z196. Virtual servers running in zBX and a z/OS system working as an IP router are grouped to a separate IEDN VLAN 600. When the customer tries to access virtual servers running in zBX, the connection first goes to the z/OS IP router through OSD OSA channel path, and then routes to the target virtual servers through the IEDN network.

- Internal network, named the Injection network, is only for data communication between injection servers that are driving the workload and z196. The customer cannot access this network externally.

4.3 Configuring storage

Figure 4-7 shows the storage infrastructure in our project environment. The following details apply:

- z196 and zBX share storage in a physical DS8800 unit.
- z196 is connected to DS8800 through IBM FICON® directors, based on FICON protocol. Then, the z/OS, z/VM, and Linux systems that are running on z196 can use extended count key data (ECKD) volumes that are defined in the DS8800.
- zBX is connected to DS8800 through SAN switches, based on Fibre Channel Protocol. Then, the virtual servers that are running on POWER blades and x blades in zBX can use SCSI disks that are defined in the DS8800.

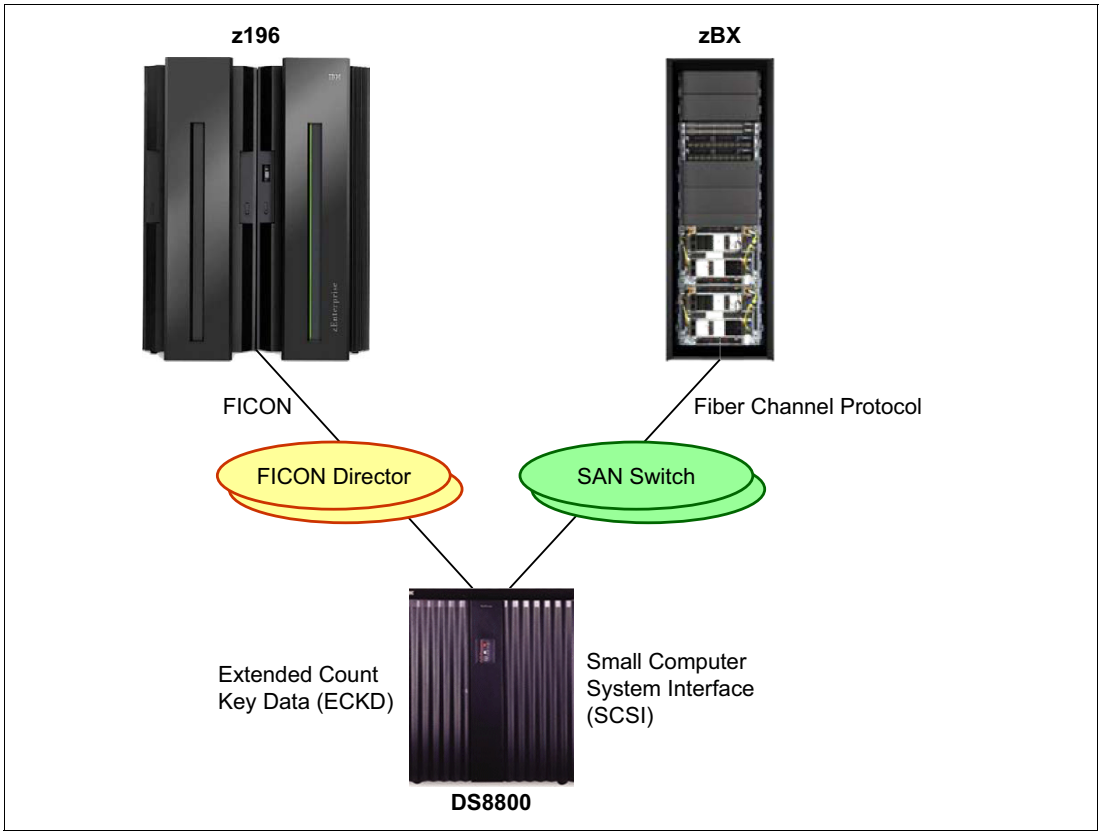

Figure 4-7 Storage infrastructure

For this project environment preparation, the z/OS Parallel Sysplex is an existing environment that has attached storage. Therefore, our only steps are to apply ECKD volumes for two Linux servers, and SCSI disks for virtual servers running on blades in zBX:

- ECKD volumes: Four 3390-9 volumes for two Linux guests
- SCSI disks: One 120-GB disk for one AIX system that is running on POWER blade, and another 120-GB disk for one Windows system that is running on an x blade in the zBX.

A storage administrator can then help to prepare the following storage resources:

1. z/VM can attach to all ECKD volumes that are defined in z196 input/output configuration data set (IOCDS), and then assign to Linux. So, the storage administrator needs to assign the ECKD volumes in an existing DS8800 for this project.
2. For SCSI disks, a storage administrator needs to create a logical unit number (LUN) on your existing DS8800, create zones in the SAN switch, and create and import a storage access list for each hypervisor. Then, we have storage resources to install new virtual servers on these hypervisors. For details of these operations, see *Building an Ensemble Using IBM zEnterprise Unified Resource Manager*, SG24-7921.

4.4 Configuring z/VM and Linux on System z

In our project environment, we need to create two more Linux virtual servers on a z/VM that is running in z196.

In our case, the z/VM is an existing environment, and we have controller and template images of SLES 10, so we can clone Linux from template images by using controller Linux as follows:

1. Use the prepared ECKD volumes, and format DASD for minidisks.
2. Define a new user ID for virtual server in z/VM.
3. Log on to the controller Linux to copy from the template image mini disks to the target Linux mini disks.
4. Start the target Linux, and change the network configuration, such as IEDN and INMN.

The detailed configuration process is in *z/VM and Linux on IBM System z: The Virtualization Cookbook for SLES 10 SP2*, SG24-7493.

4.5 Configuring POWER blades

In our project environment, we need to create one more AIX virtual servers on a POWER blade in zBX.

After planning network and storage, we can prepare the AIX virtual server as follows:

1. Use Unified Resource Manager to create a virtual server on a designated POWER blade.
2. Assign the planned network and storage to this virtual server.
3. Install AIX from the Network Installation Manager (NIM) server.
4. Enable IPv6 to join INMN.

The detailed configuration process is in *Building an Ensemble Using IBM zEnterprise Unified Resource Manager*, SG24-7921.

4.6 Configuring x blades

In our project environment, we need to create one more Windows virtual server on an x blade in zBX. After planning network and storage as described in previous sections, we can prepare the Windows virtual server as follows:

1. Use Unified Resource Manager to create a virtual server on the designated x blade.
2. Assign the planned the network and storage to this virtual server.
3. Install the Windows Server 2008.

4.7 Configuring workload performance management using Unified Resource Manager

We define the BI workload, configure the guest platform management provider, and then enable Application Response Measurement.

4.7.1 Defining the BI workload

One of the strengths of the System z platform is the ability to run multiple workloads at the same time across multiple z/OS images, and to manage those workloads according to performance goals that you set. The IBM zEnterprise System (zEnterprise) extends this performance management capability to both traditional System z and BladeCenter hardware environments.

The Unified Resource Manager provides specific performance management functions:

- Performance monitoring and reporting functions provide system administrators and performance analysts with the data they need to understand whether performance goals are being met. If these goals are not being achieved, detailed performance data helps identify the source of performance problems.
- The virtual server processor management function manages processor (CPU) resources across virtual servers based on a goal-oriented performance policy. For example, Unified Resource Manager can dynamically adjust processor resource allocations between the virtual servers on z/VM and between those on each POWER blade.

In our project environment, we enabled INMN for each virtual server, enabled process management for each hypervisor and virtual server. Then, we define a workload for this BI application. The detailed configuration process of performance management can be found in *zEnterprise System Ensemble Performance Management Guide*, GC27-2607.

Figure 4-8 shows the virtual servers that are included for this BI workload. The Windows virtual server is used only for running some development and modelling tools and clients, but is not involved in run time. So we do not need to include it in the workload.

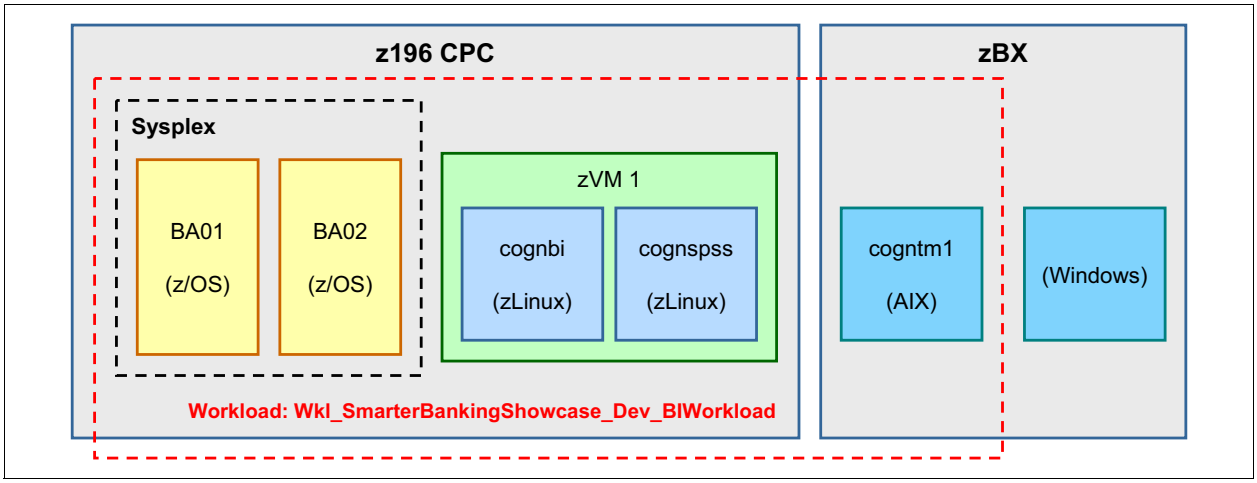

Figure 4-8 SmarterBankingShowcase BI workload

Figure 4-9 shows the detailed definition of the workload, including performance policy, service classes, and classification rules.

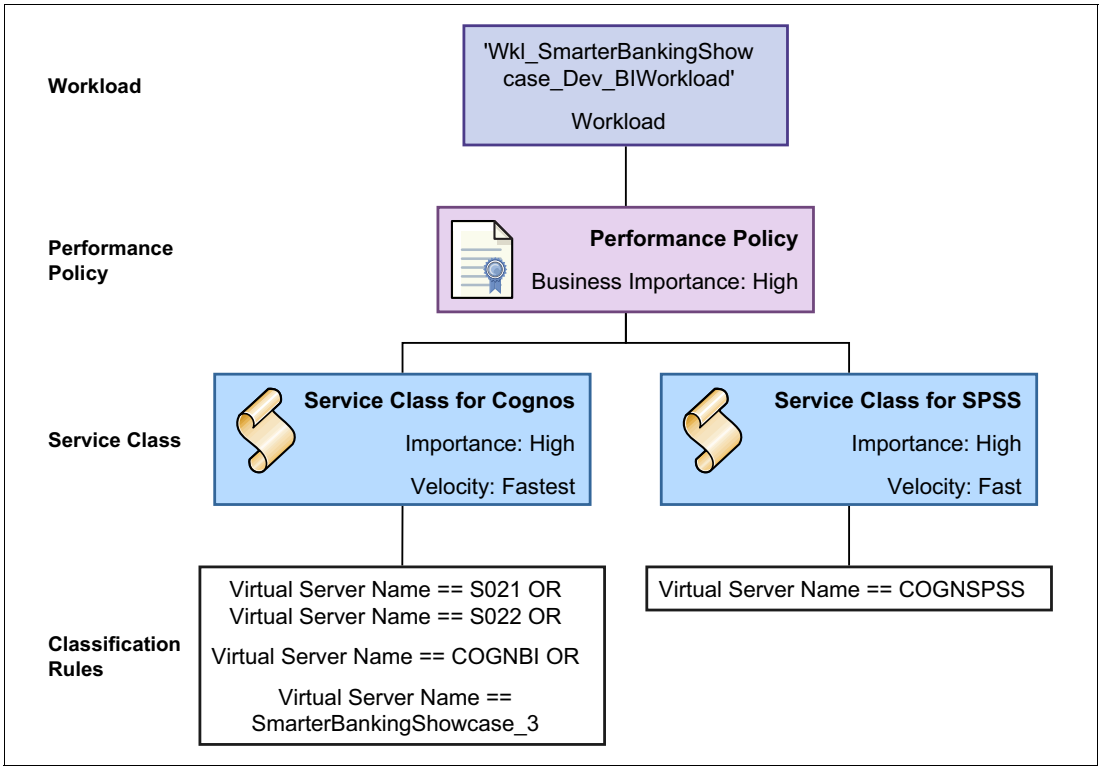

Figure 4-9 SmarterBankingShowcase BI workload detailed definition

4.7.2 Configuring GPMP and enabling ARM

We then configured the guest platform management provider (GPMP) and enabled Application Response Measurement (ARM) for supported middleware. With customization that includes installing or starting guest platform management providers, you can request additional performance reports that include application data from IBM middleware products, which are instrumented with the Open Group Application Response Measurement 4.0 standard.

The additional reports include the following items:

- Virtual Server Topology report: This report provides a graphical view of the relationships between virtual servers that are running the workload and providing the resources to complete the work.
- Hops Report: This report illustrates information for each hop involved with handling a specific service class within a specific workload.

In our project environment, we installed and configured GPMP for z/OS, Linux, and AIX virtual servers. The ARM-instrumented IBM middleware products includes DB2 UDB, DB2 on z/OS, WebSphere Application Server, and HTTP Server. We enabled ARM for these middleware components that were involved in our test.

Figure 4-10 shows the middleware we enabled in ARM.

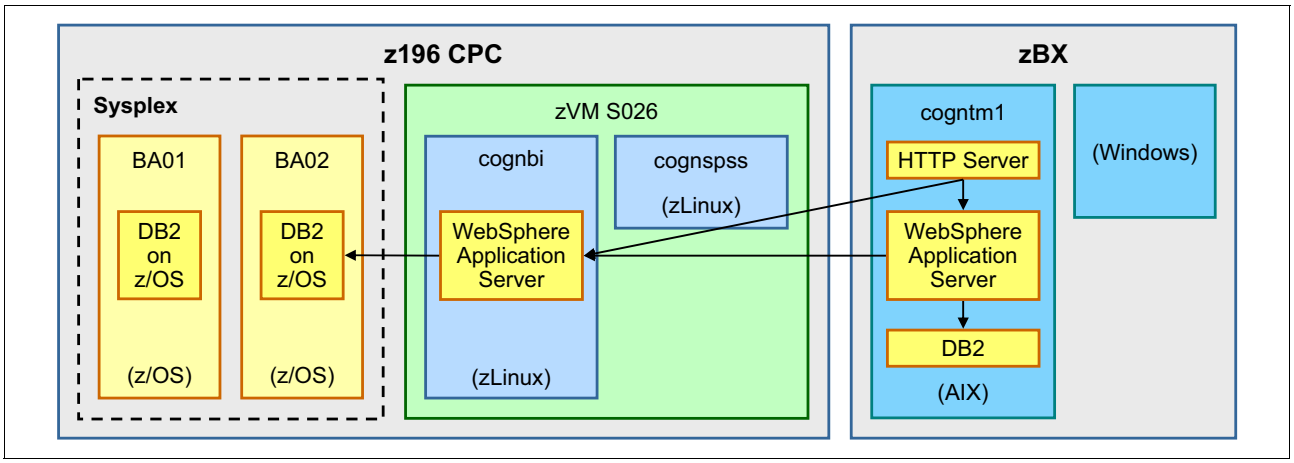

Figure 4-10 Enable ARM for supported middleware

5

Install, configure, and integrate the software components

In this chapter, we describe how to prepare your previously installed hardware components for a successful installation of the software products. We also explore what you have to consider during the software installation processes, the CPA blueprint deployment, and the integration of the individual parts.

Although this book does not provide detailed steps of product installation instructions, because they are described in the product installation guides, we do examine the possible difficulties and settings you should to be aware of.

5.1 IBM Cognos BI installation considerations

Cognos BI installation is covered extensively in other publications.

We installed IBM Cognos Metrics Manager and shared resources with the IBM Cognos Business Intelligence 10.1.1. Because Cognos Metrics Manager is not currently supported on Linux on System z, we installed Metrics Manager and shared resources on separate computers. See 5.2, “IBM Cognos Metric Studio deployment and integration considerations” on page 42.

5.1.1 Overview of the IBM Cognos BI installation

Cognos BI is installed as follows:

1. Define the content store, notification database, Human Task service, and annotation database in DB2 on z/OS.
2. Install the Cognos BI application tier components and Content Manager on Linux on System z. These components run on the WebSphere Application Server on Linux on System z.
3. Install Cognos BI gateway in the same directory as Cognos Metrics Manager that is running in AIX.

5.1.2 Installing Cognos BI components

To install the Cognos BI, follow the process described in the *Cognos Business Intelligence Installation and Configuration Guide:*

http://pic.dhe.ibm.com/infocenter/cbi/v10r1m0/index.jsp?topic=%2Fcom.ibm.swg.im.cognos.inst_cr_winux.10.1.0.doc%2Finst_cr_winux.html

1. Start the installation wizard.
2. Install Cognos BI application tier components and Content Manager on Linux on System z.
3. After setting several environment variables, such as JAVA_HOME, start the Cognos configuration tool to configure Cognos BI:
   a. Go to *c10_installation*/bin64 directory.
   b. Type ./cogconfig.sh to start Cognos configuration tool.

5.1.3 Configuring Cognos BI data access

We define the content store, notification database, Human Task service, and annotation database in DB2 on z/OS. Then, the Content Manager running on the WebSphere Application Server on Linux on System z will connect to this database through a JDBC type 4 driver.

Define database, table space, and tables in DB2 on z/OS

A database administrator must run a script to create a set of table spaces that are required for the content store database. The script must be modified to replace the placeholder parameters with parameters that are appropriate for your environment. The tables of the content store will be created and initiated by Cognos BI when it starts.

If you are using the same DB2 database on z/OS for both the content store, notification databases, Human Task service, and annotation databases, run the scripts to create these database, table spaces, and tables at the same time that you create the content store database table spaces.

These schemas can be found in the following location. Be sure to use the schemas for DB2 on z/OS:

c10_installation/configuration/schemas

Prepare the JDBC driver

IBM Cognos Business Intelligence uses Java Database Connectivity (JDBC) to access the database that is used for the content store. If you are using a DB2 database on z/OS for the content store, you must use a type 4 JDBC connection.

Copy the following files from the *DB2_installation*/sqllib/java directory to the *c10_location*/webapps/p2pd/WEB-INF/lib directory:

- The universal driver file: db2jcc.jar
- The license file for DB2 on z/OS: db2jcc_license_cisuz.jar

Configure the DB2 connection in the Cognos configuration

In the Cognos configuration tool, we must configure several advanced properties based on our naming convention and environment on DB2 on z/OS. These values are symbols that the Cognos BI will use to replace parameters in schema to define content store tables and do testing. Figure 5-1 shows the advanced properties that are required in our environment.

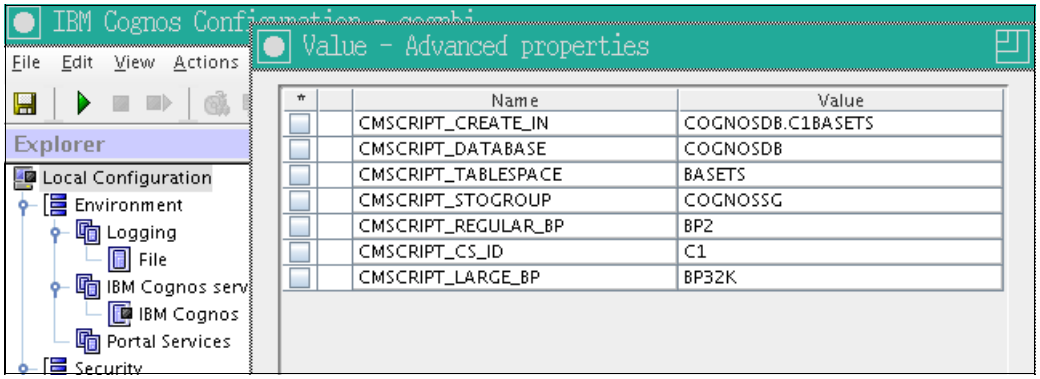

Figure 5-1 Cognos configuration advanced properties

Next, we can configure the connection from Content Manager to DB2 on z/OS. In the content store database resource properties, we use the definitions shown in Figure 5-2 on page 40:

- The database server and port number point to DB2 on z/OS. The IP address is the HiperSockets dynamic VIPA IP address of a z/OS Parallel Sysplex to provide high availability. The port number is a DB2 data sharing distributed data facility (DDF) listening port.
- Enter the user ID and password that has authority to create tables in the content store database and table spaces that we previously defined.
- For DB2 on z/OS, the database name should be the location name of the DDF.

After the definition, we click **Actions** → **Test** to test the connection to DB2 on z/OS.

Because the application tier components and Content Manager are both installed on the same system, with the notification database, Human Task service, and annotation database

in the same DB2 subsystem on z/OS as the content store, defining a DB2 connection to these databases is unnecessary.

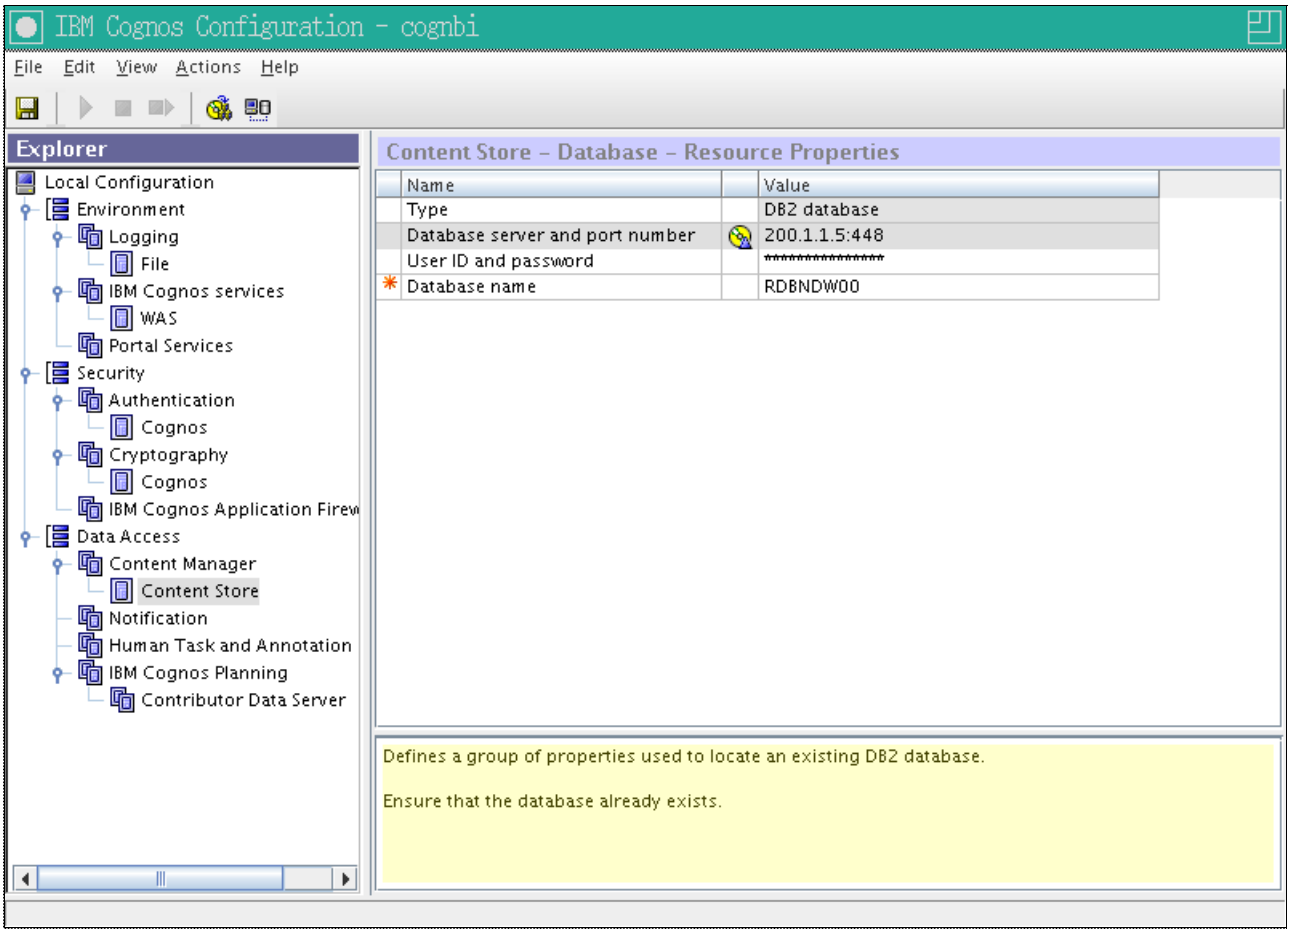

Figure 5-2 Content store database resource properties

5.1.4 Preparing to run Cognos BI on WebSphere Application Server

IBM Cognos BI installs and uses Tomcat as the application server, by default. We chose to run IBM Cognos BI within IBM WebSphere Application Server for better availability, scalability, and so on. Therefore, we needed to create a WebSphere Application Server cluster in advance, and deploy Cognos BI in a WebSphere Application Server instance within a managed node.

We backed up the configuration files for running on Tomcat, then configured for running on WebSphere Application Server. Figure 5-3 on page 41 shows environment group properties for running Cognos BI on WebSphere Application Server:

- We use the IEDN network IP address for communication between Linux on System z and AIX on POWER blade in zBX.
- The Cognos BI gateway can run on AIX on a POWER blade, so the gateway URI points to an HTTP server that is running on AIX on POWER blade.
- The dispatcher is an application tier component that will run on WebSphere Application Server on Linux on System z, so the dispatcher URI for the gateway points to WebSphere Application Server on Linux on System z.

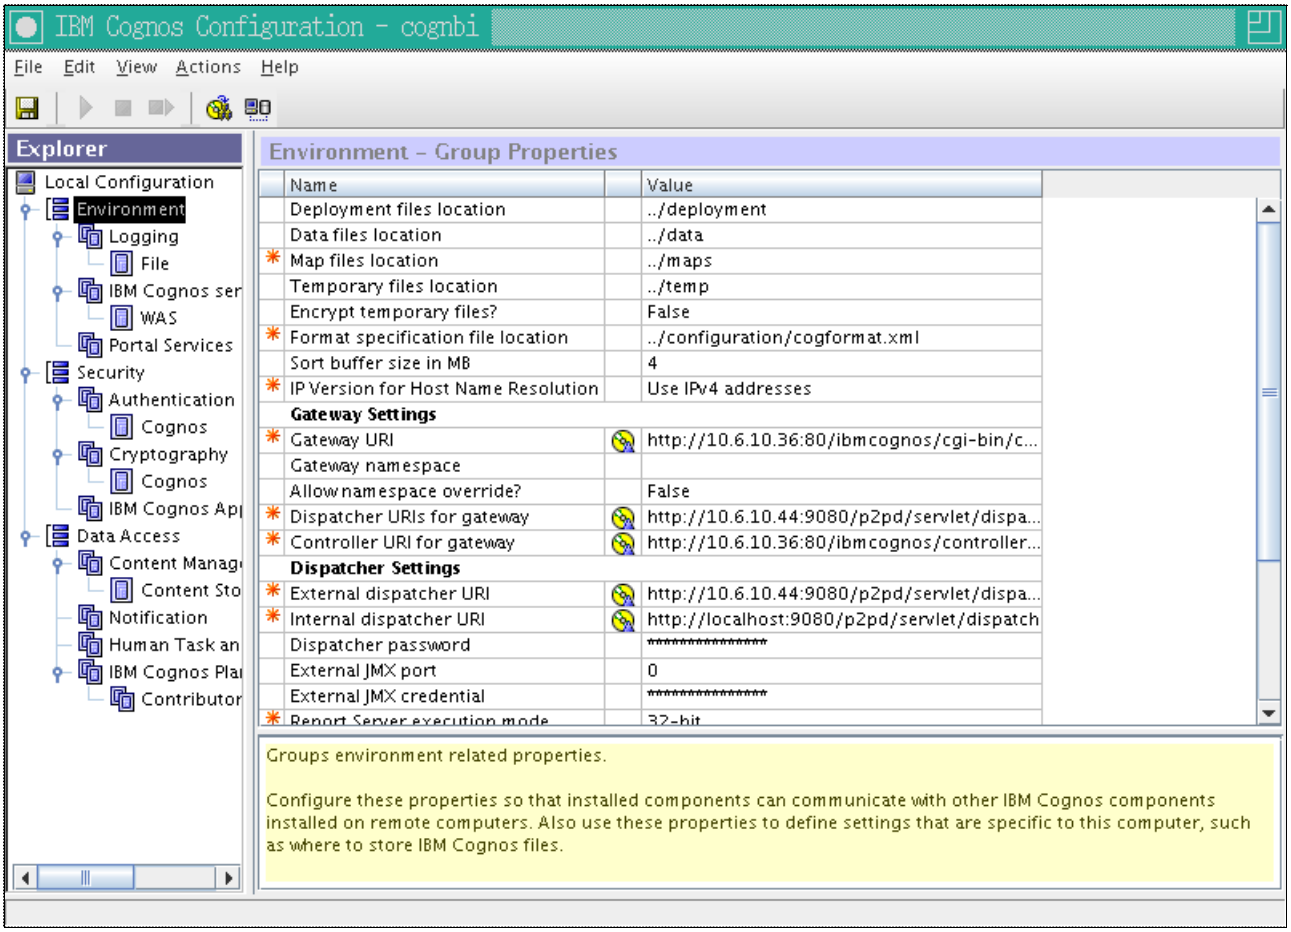

Figure 5-3 Environment group properties for running Cognos BI on WebSphere Application Server (WAS)

Expand **Environment** → **IBM Cognos services**, right-click **IBM Cognos**, and then click **Delete** to delete the entry for the IBM Cognos service running under Tomcat, then add the new name of WebSphere Application Server for IBM Cognos service running in WebSphere Application Server.

An environment variable should be defined in **WebSphere Application Server instance Server** → **Java and Process Management** → **Process Definition** → **Environment Entries for LD_LIBRARY_PATH**, which points to the `c10_location/bin64` location.

After these configuration settings and various environment variable settings were updated, we used the Build Application Wizard to build the EAR file that will be deployed in WebSphere Application Server. We then deployed the exported application to a WebSphere Application Server instance from the WebSphere Application Server administrative console.

We tested the Content Manager and dispatcher from these URIs in case we use the context root as the default of p2pd:

- Content Manager: `http://9.212.128.44:9080/p2pd/servlet`
- Dispatcher: `http://9.212.128.44:9080/p2pd/servlet/dispatch`

5.1.5 Installing Cognos BI gateway

IBM Cognos Metrics Manager is available only in a 32-bit version. You must install it on a 32-bit system. If sharing resources with IBM Cognos BI server, install the gateway from both products in the same directory on a 32-bit system.

Therefore, we install the Cognos BI gateway in the same directory as the Cognos Metrics Manager in AIX on POWER blade.

5.2 IBM Cognos Metric Studio deployment and integration considerations

IBM Cognos Metric Studio as part of IBM Cognos Business Intelligence is not yet available in a version for Linux on System z. The Metric Studio icon is available and visible only when a Metrics Server is installed in your Business Intelligence infrastructure. If you want to use scorecarding capabilities in your reports within the Business Intelligence infrastructure, consider a distributed deployment of the involved components.

Determining where to put your IBM Cognos Metrics Server is not the only consideration; you also must consider the following aspects:

- IBM Cognos Metrics Server application tier installation
- IBM Cognos Metrics Server gateway installation
- Metric Store placement
- IBM Cognos Business Intelligence integration

5.2.1 IBM Cognos Metrics Server application tier and gateway installation

As described previously, IBM Cognos Metrics Server is not yet available on Linux for System z, so we had to deploy across a distributed Business Intelligence infrastructure, and the Metrics Server was placed on our POWER 7 zBX blade that was running AIX 6.1.

Because the Metrics Server is available in only a 32-bit version, we used a WebSphere Application Server v7 32-bit for the deployment (current patch level is 7.0.0.13). To integrate Metrics Server capabilities into our infrastructure, we adjusted its configuration, pointing to our Cognos Business Intelligence Infrastructure. We describe where to point in a later section. For now, we describe installing the components.

To be able to access the installed Metrics Server application tier, we must install the IBM Cognos Metrics Server gateway and refer it to our HTTP Server and the appropriate dispatcher. Because it is a 32-bit installation and part of Metrics Server, we deploy it on AIX in our 32-bit WebSphere Application Server also. Installing Metrics Server application tier and gateway together is involves selecting a check box in the installation wizard (when using the GUI-based installation).

Because we already installed the IBM Cognos Content Manager on Linux on System z, we can reuse it for Metrics Server, therefore we do not need to install it again. For the content database, we access the DB2 for z/OS of IBM Cognos Business Intelligence that we already defined. Therefore, we do not need to configure it again.

Variable setting: Set the AIXTHREAD_SCOPE variable to `S`:

```
export AIXTHREAD_SCOPE=S
```

Figure 5-4 shows the window for selecting the components to install. For the complete installation information, see the *IBM Cognos Business Intelligence Installation and Configuration Guide*, GC19-2964.

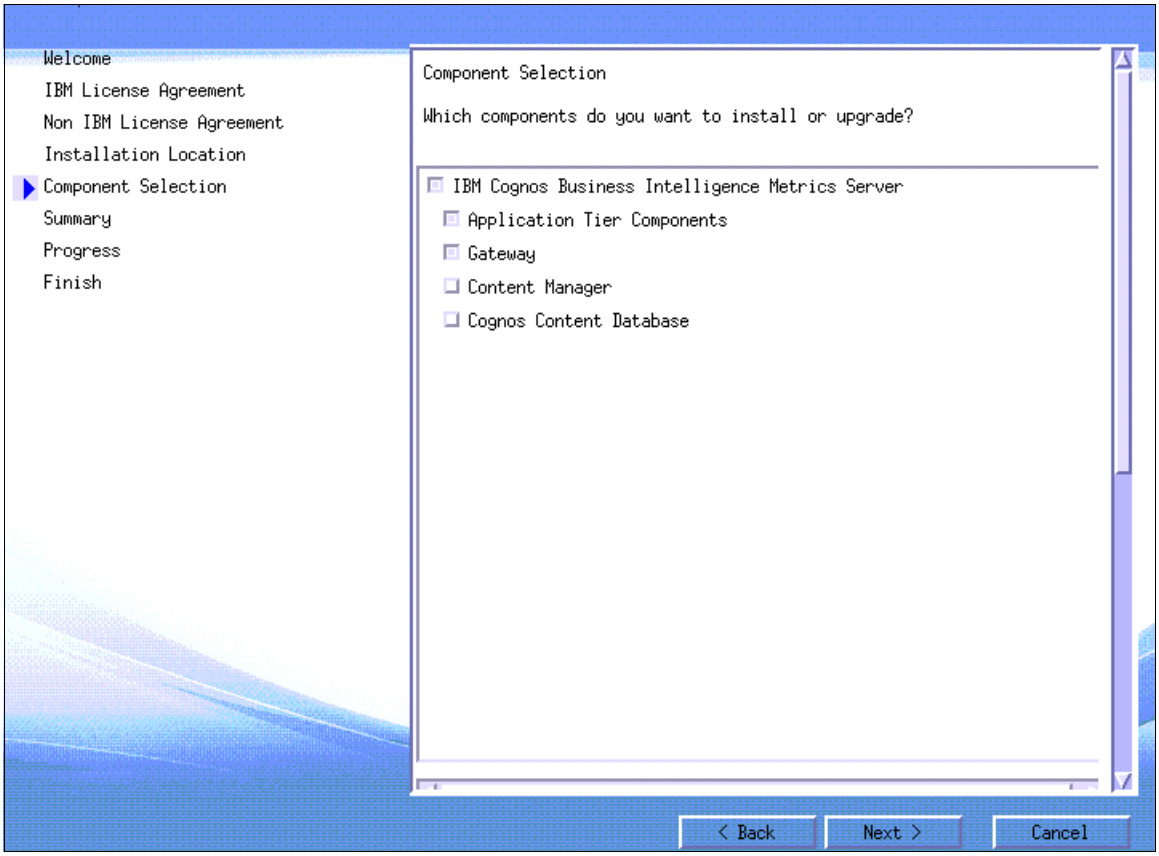

Figure 5-4 Component selection during IBM Cognos Metrics Server installation

Append the `c10_location/bin` directory to the appropriate library path environment variable.

```
export LIBPATH=/opt/Cognos/c10BIGW/bin:$LIBPATH
```

5.2.2 Metric store placement

IBM Cognos Metrics Server requires a metric store database that will contain content for metric packages. It also contains scorecarding application settings, such as user preferences. The metric store has not yet been created, so we have to do so.

IBM Cognos Metrics Server does not support DB2 for z/OS, so we cannot use the same database that we use as Content Store Database. IBM Cognos Metrics Server currently supports only DB2 on LUW, Microsoft SQL Server, and Oracle databases. Therefore, we place a DB2 9.5 Enterprise Server Edition (ESE) on the same system as the application tier that is hosting WebSphere Application Server on our AIX blade to use it as metric store.

For detailed information about supported metric store databases, see the following location:

https://www.ibm.com/support/docview.wss?uid=swg27019126#metric_stores

Before you can use a DB2 or Oracle database as a metric store, you must specify environment variables.

The proper syntax for creating environment variables is shell-dependent. For IBM DB2 databases, you must set the database variables by running the environment setup scripts that are included with the IBM DB2 installation. For Bourne or Korn shells, run the following command or add it to the `profile` script:

```
DB2_installation_path/cfg/db2profile
{cogntm1:root}/opt/IBM/db2/V9.5/cfg # ./db2profile
```

You can use a prepared script to create the required metric store database, which is provided with the standard IBM Cognos installation. Run the *`cmm_create_db`* `script` in the `c10_location/configuration/schemas/cmm/db2` directory with the following syntax:

```
cmm_create_db <instance><userid><password><dbname><diskdrive><alias>
```

See the *IBM Cognos Metrics Manager Getting Started Installation Guide* for further information, and see the following web resource for detailed attributes of the statement:

http://ibm.co/RPEAT6

5.2.3 IBM Cognos Business Intelligence gateway deployment and integration

At the end of this project, we plan to have an integrated business scenario that is based on a hybrid infrastructure. Therefore, after installing Cognos BI application tier components on Linux for System z, and also Metrics Server application tier components and Metrics Server gateway on AIX, you should integrate components such as the following topics.

IBM Cognos Business Intelligence gateway placement

If you want to share resources between IBM Cognos BI server and IBM Cognos Metrics Server in a distributed environment, you must install both gateways in the same directory. Because Metrics Server gateway is available only in a 32-bit system, you have to install both gateways in the same 32-bit directory, in our scenario on the AIX-based zBX blade.

WebSphere Application Server clustering

For centralized management, and to be able to use one single deployment manager, you can cluster the WebSphere Application Servers. For our scenario, we clustered the 64-bit WebSphere Application Server installed on Linux on System z, and the 32-bit WebSphere Application Server installed on the AIX zBX blade. It might look similar to Figure 5-5.

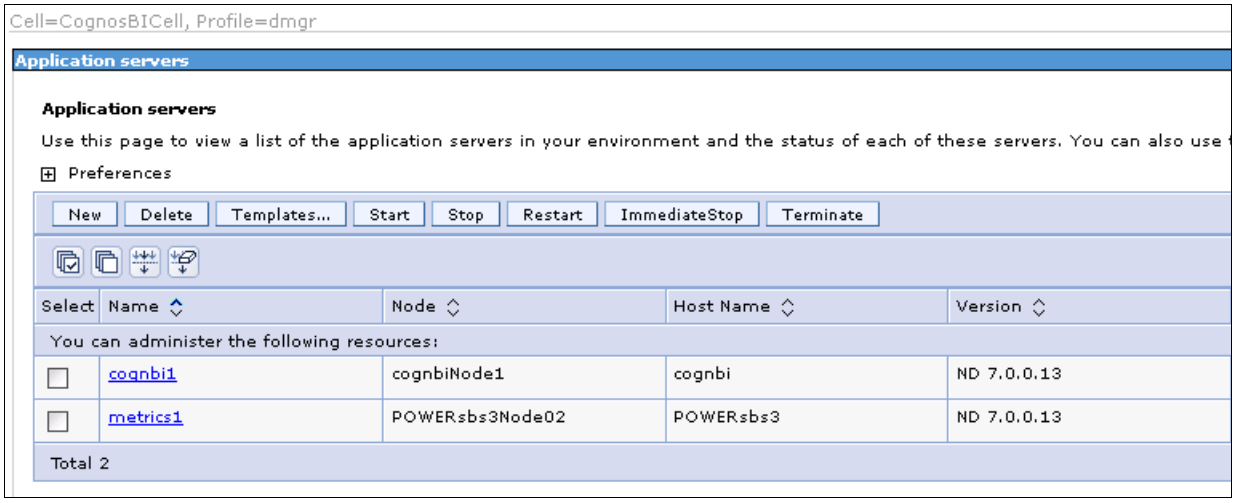

Figure 5-5 WebSphere Application Server cluster nodes

HTTP Server placement

Consider the entry point for external requests. The HTTP Server that provides the entry point to your hybrid business intelligence infrastructure should be placed next to your gateways and at the smallest common denominator. In our scenario, this approach means deploying the HTTP Server on the 32-bit WebSphere Application Server on AIX.

All these installation considerations express the architecture shown in Figure 5-6.

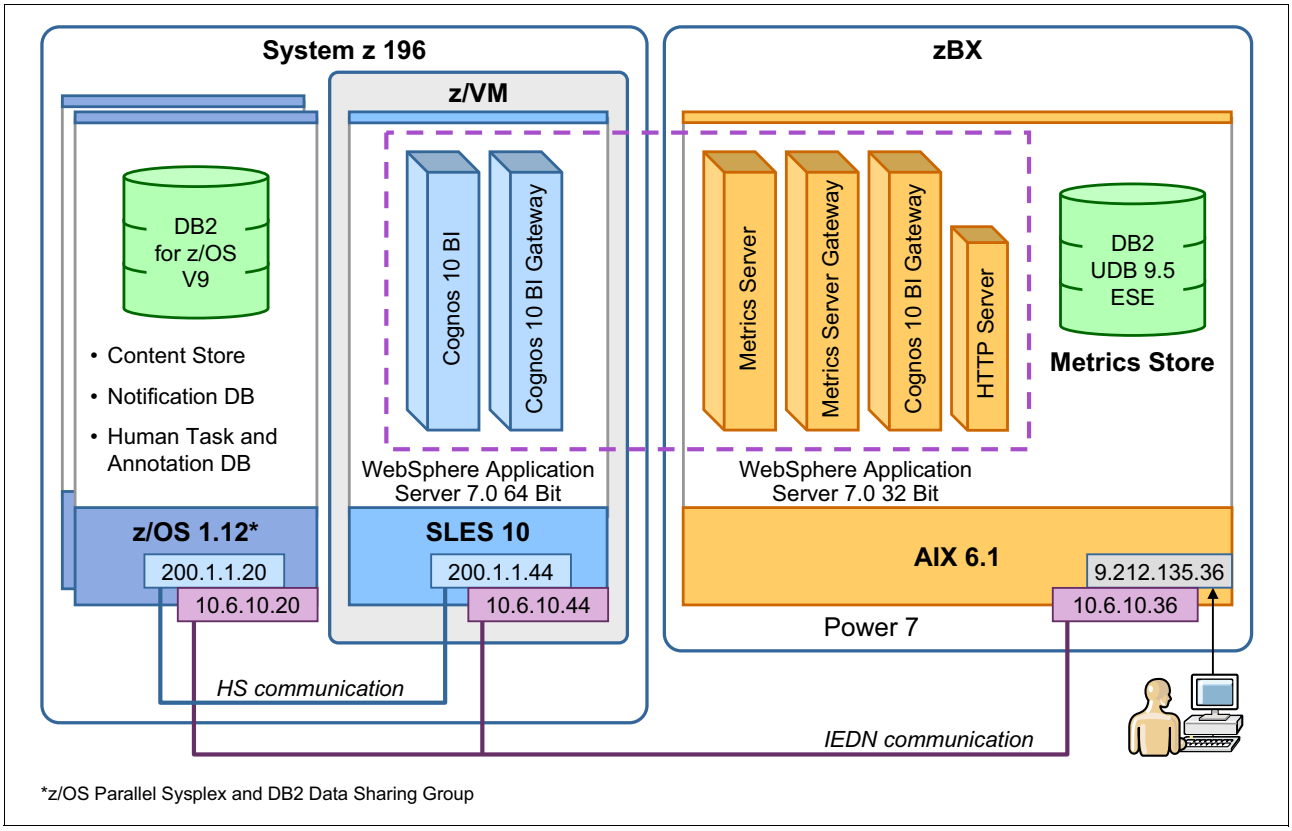

Figure 5-6 IBM Cognos Business Intelligence and Metrics Server integration architecture

5.2.4 Metrics Server configuration and package deployment

In this section, we prepare for the deployment, and then configure, create, and deploy the package.

Preparing for deployment

- Set environment variable JAVA_HOME pointing to the AppServers JVM:

  ```
  export JAVA_HOME=/opt/IBM/WebSphere32/AppServer/java/jre
  ```

- Set environment variable LIBPATH:

  ```
  export LIBPATH=/opt/Cognos/c10Metrics/bin
  ```

- Change the default connection time-out.

Open the `BIBusTK_Config.xml` file, which is in the `<c10_location/configuration/>` location. Change the `<BIBUSTK_CONNECTION_TIMEOUT>` to 90% of the value that is specified for your application server. WebSphere Application Server uses a default connection time-out value of 30 seconds. We use 27 seconds. See Example 5-1.

Example 5-1 BIBusTK_Config.xml adjusted for WebSphere Application Server deployment

```
<?xml version="1.0" encoding="UTF-8"?>
<!--
Licensed Materials - Property of IBM

IBM Cognos Products: cclbibustk

© Copyright IBM Corp. 2003, 2010

US Government Users Restricted Rights - Use, duplication or disclosure restricted
 by GSA ADP Schedule Contract with IBM Corp.
-->
<BIBusTKConfig>
        <!-- BIBUSTK_CONNECTION_TIMEOUT: Overrides the timeout value for BIBusTK
connection pool. -->
        <!-- 25000 ms as default. -->
        <!-- -->
        <BIBUSTK_CONNECTION_TIMEOUT>27000</BIBUSTK_CONNECTION_TIMEOUT>
        <!-- -->
        <!-- BIBUSTK_SOCKET_TCP_NODELAY configuration options -->
        <!-- 0: Turn off TCP_NODELAY for BIBusTkSocket. -->
        <!-- 1: Turn on TCP_NODELAY for all BIBusTkSocket. -->
        <BIBUSTK_SOCKET_TCP_NODELAY>1</BIBUSTK_SOCKET_TCP_NODELAY>
        <!-- -->
</BIBusTKConfig>
```

Configuring, creating, deploying the package

Before you create an appropriate package of IBM Cognos Business Intelligence and Metrics Server for a deployment on the WebSphere Application Servers, you must configure the packages. First, launch each configuration wizard by using the **`/cogconfig.sh`** command.

> **Set JAVA_HOME:** Make sure JAVA_HOME is set and pointing to the 32-bit WebSphere Application Server JVM before you start the configuration wizard.

Adjust the following settings:

- Gateway URI
- Dispatcher URIs for gateway
- Controller URI for gateway
- External dispatcher URI
- Internal dispatcher URI
- Report Server execution mode
- Dispatcher URI for external applications
- Content Manager URIs

Figure 5-7 shows settings in our scenario. You see that we used HiperSockets communication between our IBM Cognos Business Intelligence and Metrics Server components.

Environment – Group Properties

	Name		Value
	Data files location		../data
*	Map files location		../maps
	Temporary files location		../temp
	Encrypt temporary files?		False
*	Format specification file location		../configuration/cogformat.xml
	Sort buffer size in MB		4
*	IP Version for Host Name Resolution		Use IPv4 addresses
	Gateway Settings		
*	Gateway URI		http://10.6.10.36:80/ibmcognos/cgi-bin/c...
	Gateway namespace		
	Allow namespace override?		False
*	Dispatcher URIs for gateway		http://10.6.10.36:9080/p2pd/servlet/dispa...
*	Controller URI for gateway		http://10.6.10.36:80/ibmcognos/controller...
	Dispatcher Settings		
*	External dispatcher URI		http://10.6.10.36:9080/p2pd/servlet/dispa...
*	Internal dispatcher URI		http://localhost:9080/p2pd/servlet/dispatch
	Dispatcher password		**********
*	Report Server execution mode		32-bit
	Other URI Settings		
*	Dispatcher URI for external applica...		http://10.6.10.36:9080/p2pd/servlet/dispa...
*	Content Manager URIs		http://10.6.10.44:9080/p2pd/servlet
	Font Settings		
*	Physical fonts locations		../bin/fonts
	Physical fonts map		<click the edit button>
	Fonts to embed (batch report servi...		<click the edit button>
	Fonts to embed (report service)		<click the edit button>

Figure 5-7 IBM Cognos Metrics Server configuration settings - Environment configuration

After the environment configuration is completed, go to *IBM Cognos services* and create a configuration that matches your WebSphere Application Server installation. See Figure 5-8.

WAS – Configuration – Resource Properties

	Name		Value
	Type		WebSphere
	WebSphere Application Server loca...		/opt/IBM/WebSphere32/AppServer
	Profile		AppSrv01
	Server instance		metrics1

Figure 5-8 IBM Cognos Metrics Server configuration settings: WebSphere Application Server configuration

When the wizard finishes the package creation, you can switch to your 32-bit WebSphere Application Server and deploy the created package.

Naming conflicts: When you deploy your Metrics Server in a clustered WebSphere Application Server environment, give the application a name separate from the existing application. Otherwise there could be naming conflicts.

Your WebSphere Application Server cluster now has two deployed applications, similar to what is shown in Figure 5-9.

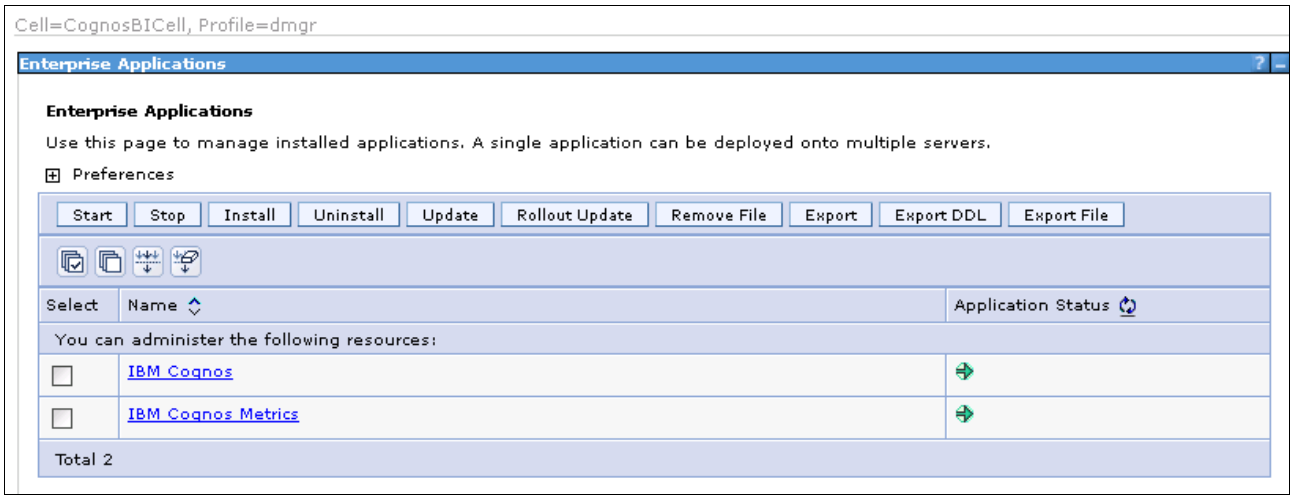

Figure 5-9 Enterprise Applications on the WebSphere Application Server cluster

Configure your web server and, to your HTTP Server configuration, add the `Alias` entries shown in Example 5-2.

Example 5-2 Aliases to be added to the HTTPServer configuration

```
ScriptAlias /ibmcognos/cgi-bin /opt/Cognos/c10Metrics/cgi-bin
<Directory "/opt/Cognos/c10Metrics/cgi-bin">
    AllowOverride None
    Options None
    Order allow,deny
    Allow from all
</Directory>

Alias /ibmcognos /opt/Cognos/c10Metrics/webcontent
<Directory "/opt/Cognos/c10Metrics/webcontent">
    Options Indexes MultiViews
    AllowOverride None
    Order allow,deny
    Allow from all
</Directory>
```

Finally, you have to set the appropriate library path in the application servers settings to make sure it uses the correct settings. See Figure 5-10.

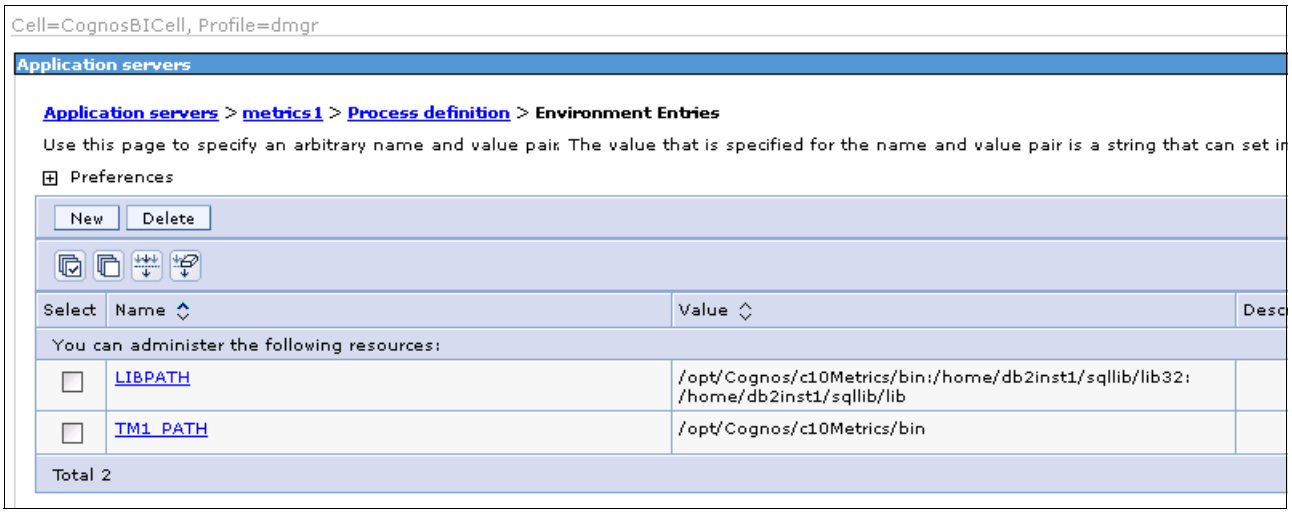

Figure 5-10 Environment variables for the Metrics Server Application Server

Related documents:

- IBM Cognos Business Intelligence Installation and Configuration Guide

  http://pic.dhe.ibm.com/infocenter/cbi/v10r1m0/index.jsp?topic=%2Fcom.ibm.swg.im.cognos.inst_cr_winux.10.1.0.doc%2Finst_cr_winux.html

- IBM Cognos Metric Studio User Guide

  http://pic.dhe.ibm.com/infocenter/cbi/v10r1m0/index.jsp?topic=%2Fcom.ibm.swg.im.cognos.ug_mm.10.1.0.doc%2Fug_mm_id3443sc_metrics.html

- Cognos Metrics Manager Installation Guide

  http://pic.dhe.ibm.com/infocenter/cfpm/v10r1m0/index.jsp?topic=%2Fcom.ibm.swg.im.cognos.inst_cr_winux.10.1.0.doc%2Finst_cr_winux_id6570InstallMetricsManagerCM_single.html

5.3 IBM Cognos TM1 installation considerations

We do not cover the installation of Cognos TM1 9.5.2 on AIX in detail here; it is covered well in other documents, including the IBM Cognos TM1 9.5 information center:

http://publib.boulder.ibm.com/infocenter/ctm1/v9r5m0/topic/com.ibm.swg.im.cognos.tm1_install.9.5.2.doc/tm1_install.html

5.3.1 Primary considerations

The primary considerations are as follows:

- The TM1 server and TM1 Admin Server will run as root.
- The default Cognos TM1 installation process requires access to X Windows.
- If you are migrating an existing TM1 server from Windows, be aware that file names on AIX are case-sensitive, but are on Windows are not. The TM1 server on AIX will attempt to convert all the files to lowercase so you must be sure that this conversion is done correctly. If you have existing Cognos BI reports that will access this TM1 server as a data source, they might require some adjustment after the migration because of the change to all lowercase file names.
- On AIX, Cognos TM1 supports only the Admin Server and TM1 server components. If you need the Cognos TM1 Architect, Contributor, or Web Client they must be installed separately on a Windows system.
- If you do not have name resolution established for your TM1 Admin Server, you must use the HOSTS file on Windows to connect by using the Cognos TM1 Architect, Contributor, or Web Client. These components do not connect to the TM1 Admin Server with only an IP address.
- Cognos TM1 is configured to make use of the maximum amount of memory available to the IBM AIX operating system. If you plan to run other software on the system with Cognos TM1, you might need to use the AIX **`ulimit`** command to make sure that each software package has adequate memory available to it.

5.3.2 Creation of a TM1 data source in Cognos BI

The creation of a data source is covered in detail in the *Cognos BI Administration and Security Guide* in the Data Management chapter:

http://publib.boulder.ibm.com/infocenter/c8bi/v8r4m0/index.jsp?topic=/com.ibm.swg.im.cognos.ug_cra.8.4.0.doc/ug_cra_id25Introduction.html

Other chapters describe the creating of a data source and the specifics of TM1 data sources.

We also found an environment variable that must be set within WebSphere to let the BI server know where to find the TM1 libraries (and SSL information). To set this variable, open the WebSphere Administration console and navigate to the Cognos Application Server (or servers). Select the Cognos application server; then, under Server Infrastructure, open select the **Java and Process Management** twistie to open it, and select **Process definition**. When that interface opens, go to Additional Properties and select **Environment Entries**.

Create an environment variable named TM1_PATH and give it the value of your Cognos BI bin directory. In our case, we use `/opt/cognos/bin64`, as shown in Figure 5-11.

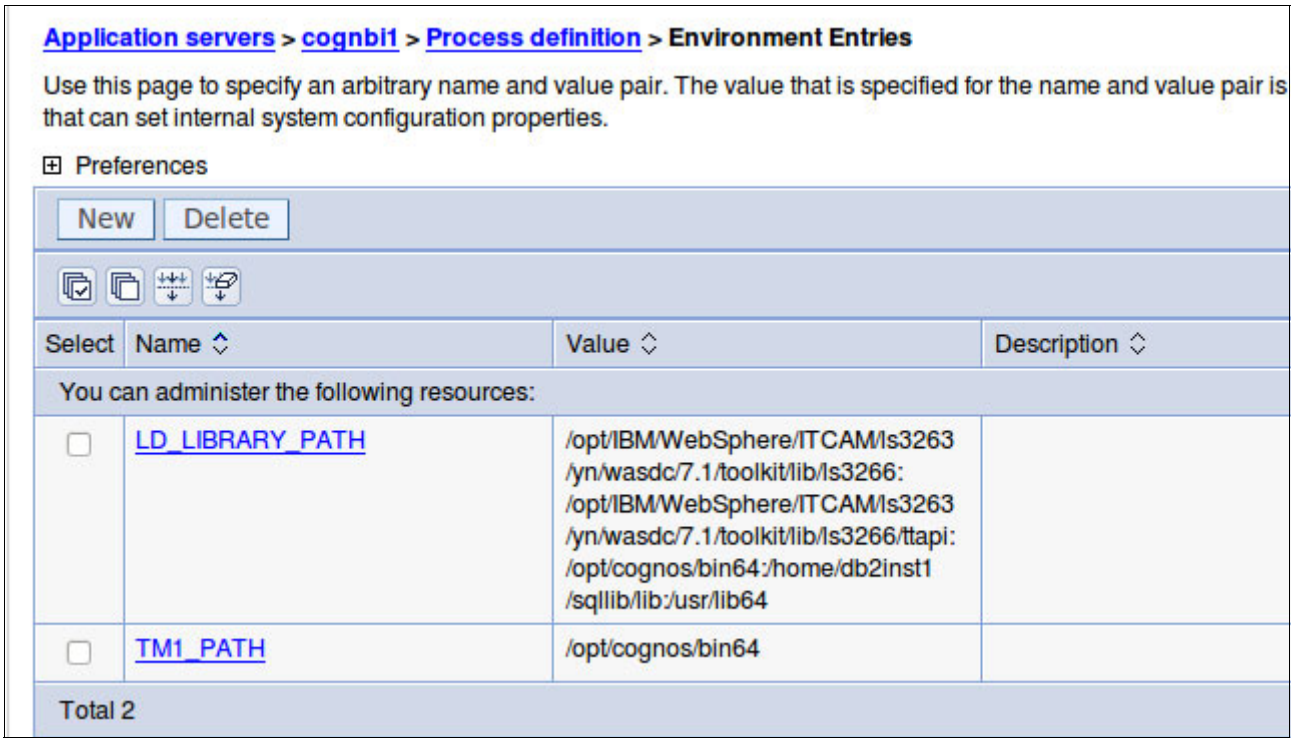

Figure 5-11 WebSphere Application Server environment variables

5.4 SPSS Modeler installation considerations

We do not cover the whole installation process of IBM SPSS Modeler components in detail. These steps are in the IBM SPSS Modeler Installation Guides, for example, these guides:

- Server Installation UNIX
- Client Installation Concurrent Network License

We show the considerations for a deployment infrastructure based upon our zEnterprise hybrid infrastructure, the decisions we face, and specific installation issues.

5.4.1 Architectural considerations and deployment preparation

The fit-for-purpose methodology used in *Using zEnterprise For Smart Analytics: Volume 1 Assessment*, SG24-8007 advises to place IBM SPSS Modeler Server on Linux on System z, and the Modeler Client on an x86-based zBX blade that is running Windows Server 2008 R2.

We also have to consider further parts of the IBM SPSS Modeler product suite:

- Sentinel RMS License Manager
- IBM SPSS Modeler Batch
- Data sources and SPSS Data Access Pack

Sentinel RMS License Manager

Installing SPSS in an distributed infrastructure or at least installing the SPSS Modeler Client on a network or terminal operating system requires you to install Sentinel RMS License Manager.

This License Manager can handle your pool of concurrent user licenses and provide a license to any connecting SPSS Modeler Client (if any are free at this time).

You should *not* install the Network License Manager on the same server as one of the clients because of availability aspects. This way can cause unavailability: if, for any reason, the client computer is not available, the License Manager would not be available also.

Because of this recommendation and the fact that it is not available on Linux for System z, we put the License Manager on our Power 7 blade that is running AIX.

Before you install the License Manager, make sure you are running the latest Version of Java and that the Java Home variable is set.

For detailed Installation Instructions, see *Client Installation Concurrent Network License Guide*, which is located in your SPSS documentation directory.

IBM SPSS Modeler Batch

IBM SPSS Modeler Batch is a component that provides the complete analytic capabilities of the standard SPSS Modeler Client but without access to the regular user interface and without the presence of graphical user interfaces. It provides capabilities to work with SPSS Modeler Server in a distributed and command-line-based batch mode.

In our scenario, we placed the IBM SPSS Modeler Batch directly next to IBM SPSS Modeler Server on Linux on System z. We describe installation in 5.4.2, “SPSS Modeler Client installation and configuration” on page 54 and 5.4.3, “SPSS Modeler server installation and configuration” on page 55. For detailed Installation Instructions, see *IBM SPSS Modeler 14.2 Batch Installation for UNIX*.

`ftp://ftp.software.ibm.com/software/analytics/spss/documentation/modeler/14.2/en/BatchInstallationForUNIX.pdf`

Data Sources and SPSS Data Access Pack

We work with DB2 for z/OS V9 as the major data source and target. In this way, we avoid flat files and benefit from additional performance and data management capabilities that are provided by DB2 for z/OS, and utilities that are related to DB2 (for example, data sharing, existing backup and recovery procedures).

To be able to connect to Data Sources such as DB2, IBM Informix®, and Oracle, enable the IBM SPSS Modeler Server to access them by using ODBC. The required ODBC capabilities are provided by the IBM SPSS Data Access Pack.

An exemplary installation and configuration process of the IBM SPSS Data Access Pack (to connect to a DB2 for z/OS) is discussed in 5.4.4, “SPSS Modeler data access” on page 58.

Architectural overview

Figure 5-12 summarizes the various aspects in the architectural model.

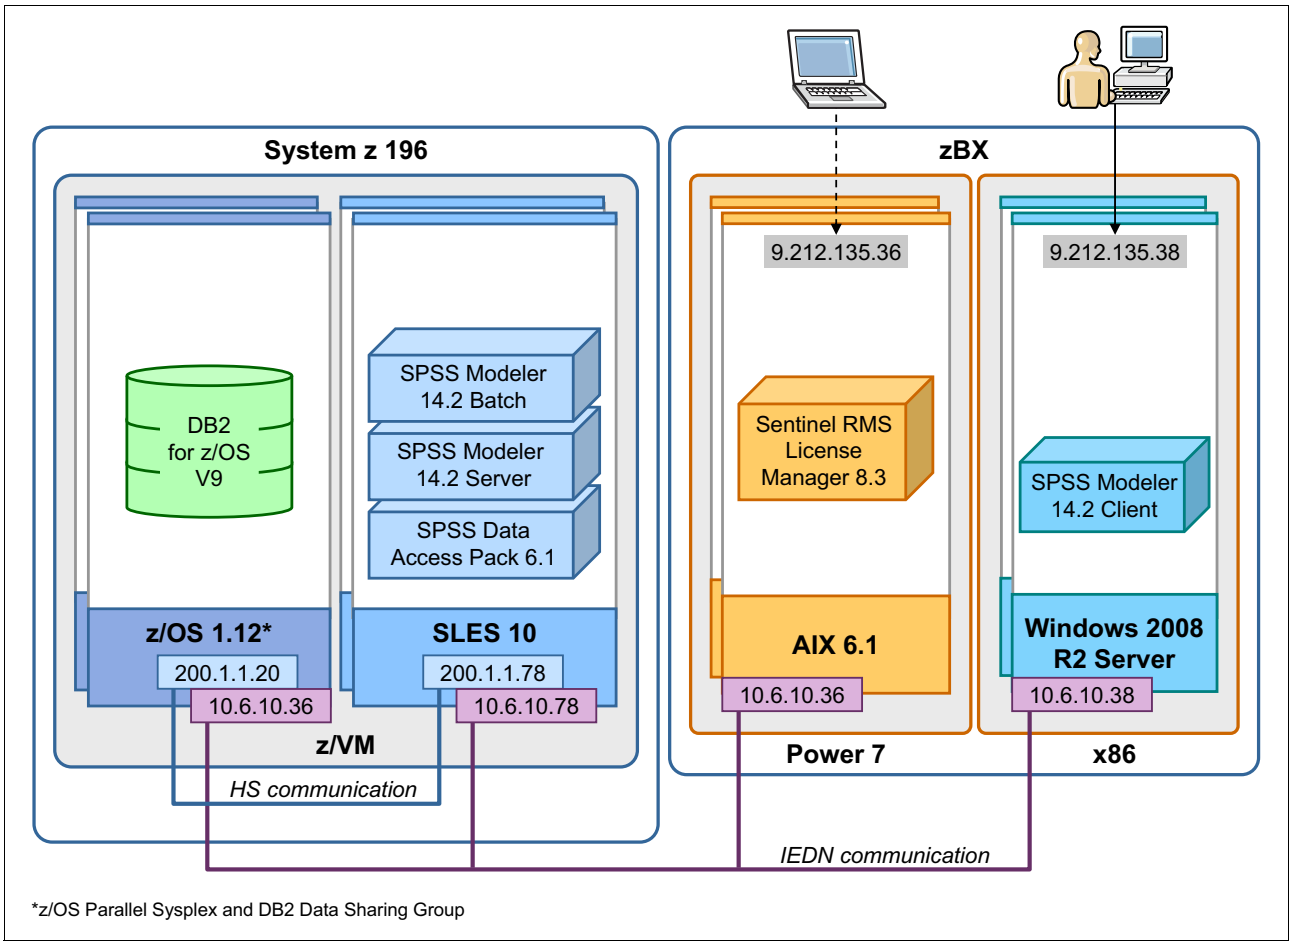

Figure 5-12 SPSS Modeler logical architecture

As shown, we are able to deploy all required SPSS components in a zEnterprise based hybrid infrastructure. The result consists of the following components:

- DB2 for z/OS on z/OS
- IBM SPSS Modeler 14.2 Server (including Batch and Data Access Pack) on Linux on System z (SLES 10)
- Sentinel Network License Manager on AIX 6.1 (on a POWER 7 based zBX blade)
- SPSS Modeler 14.2 Client on Windows Server 2008 R2 (on a x86 based zBX blade)

Figure 5-13 shows how a request flows through the components.

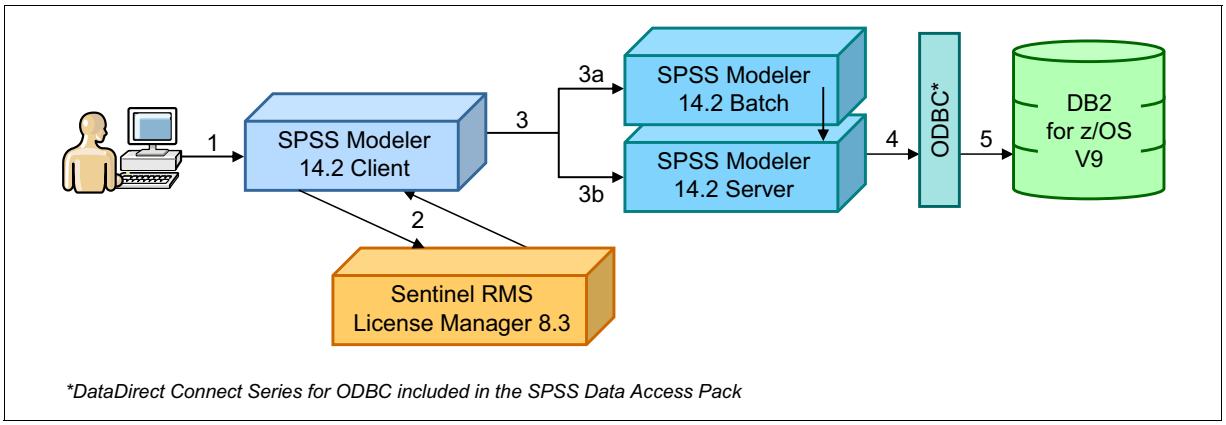

Figure 5-13 SPSS request flow

This architecture allows us to meet the advised fit-for-purpose outcome of *Using zEnterprise For Smart Analytics: Volume 1 Assessment*, SG24-8007 and benefit from zEnterprise possibilities, such as the following examples:

- High availability with z/OS in Parallel Sysplex Mode and DB2 for z/OS Data Sharing Group capabilities
- HiperSockets for fast and secure data transfer between DB2 for z/OS and SPSS Modeler Server components
- IEDN for private and fast data transfer between Modeler Client and Modeler Server components

  Because the License Manager is also available over IEDN, licenses can be provided through this channel also.
- Unique System z capabilities regarding, for example, availability, security, workload management

5.4.2 SPSS Modeler Client installation and configuration

This section lists system requirements, and describes client licensing issues and installation.

System requirements

The requirements are as follows:

- Minimum free disk space: 2 GB of available hard-disk space
- RAM: 1 GB in minimum; 2 GB or more is preferred
- Browser: Mozilla Firefox 3.x or later, or Microsoft Internet Explorer 7 or later, for online help

SPSS Modeler client licensing issues

There are several options to license your IBM SPSS Modeler, depending on the type of usage and the client operating system:

- Single user license
- Site license
- Network license

Be aware that even if you have only one IBM SPSS Modeler Client that will access the SPSS Modeler Server, you must choose Sentinel Network License Manager when you install this client on a terminal operating system such as Windows Server 200x.

Because we have Windows Server 2008 R2 as the operating system for our client, we must use a network license model and Sentinel License Manager.

Concurrent user license: You need concurrent user license for this type of licensing mode.

IBM SPSS Modeler Client installation

The client installation itself is runable in both silent mode and in GUI mode. Because the installation process itself is well documented in the installation guide.

Related documents:

- Client Installation Concurrent Network License Guide

  ftp://ftp.software.ibm.com/software/analytics/spss/documentation/modeler/14.2/en/ClientInstallation_ConcurrentNetworkLicense.pdf
- Network License Administrators Guide

  http://public.dhe.ibm.com/software/analytics/spss/documentation/statistics/20.0/en/client/InstallationDocuments/Linux/Network_License_Administrators_Guide.pdf
- Server and Performance Guide

  ftp://ftp.software.ibm.com/software/analytics/spss/documentation/modeler/14.2/en/ServerandPerformanceGuide.pdf

5.4.3 SPSS Modeler server installation and configuration

This section lists system requirements, and describes installation and batch features.

System requirements

The requirements are as follows:

- Minimum free disk space: 1 GB (recommended)
- RAM: 4 GB (recommended)
- Client software: Must be at the same release level as the IBM SPSS Modeler Server software

IBM SPSS Modeler Server installation

The server installation itself is runable in both silent mode and GUI mode. Because the installation process is well documented in the installation guide, we do not describe the installation in detail. See the appropriate installation guide (for example, named *Server Installation<yourOperatingSystem>*).

During the server installation process, your choice of license type must match your previous choice when you installed the client. In our case, we use Network License and provide the credentials to connect to the previously installed License Manager.

Select a license type (Figure 5-14) and provide an address (Figure 5-15).

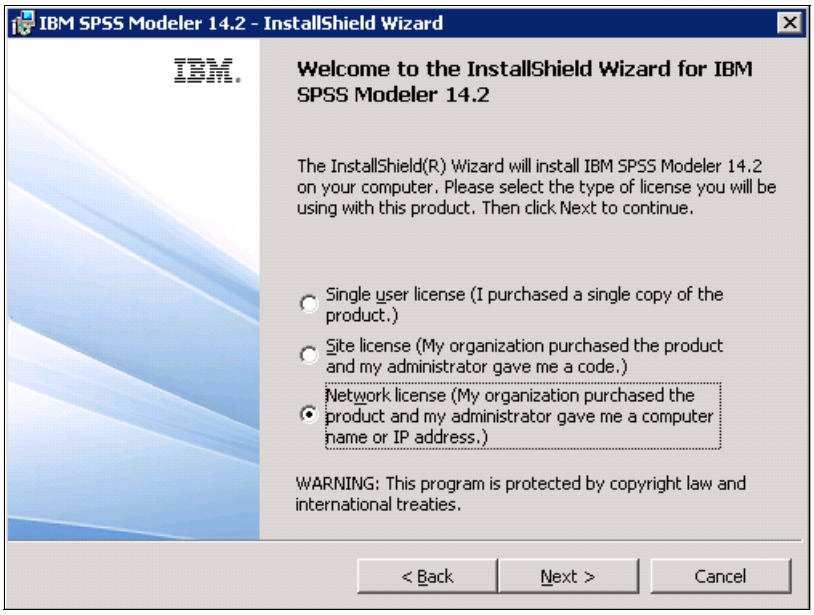

Figure 5-14 IBM SPSS Modeler 14.2 Client: Select a license type

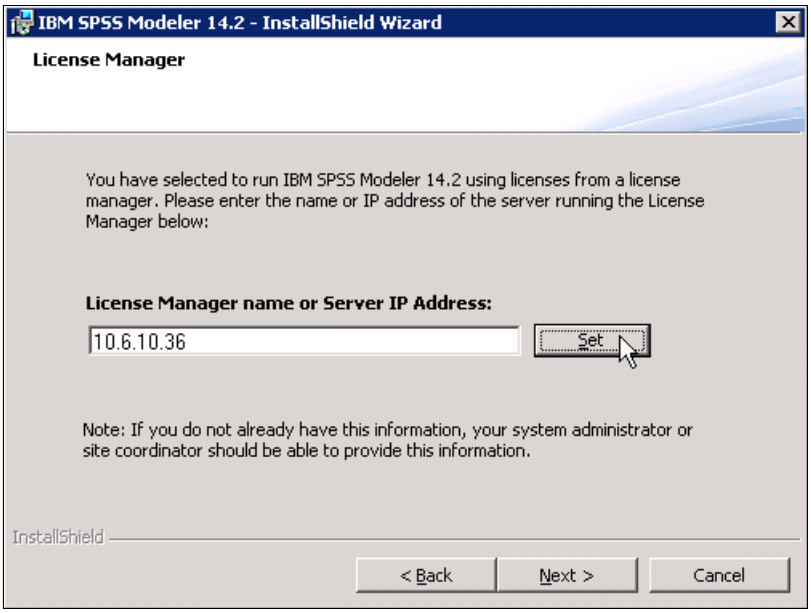

Figure 5-15 IBM SPSS Modeler 14.2 Client: Provide License Manager IP Address

Batch feature

As previously described, IBM SPSS Modeler Batch is a component that you use to process SPSS streams in batch mode. It provides support for the complete analytical capabilities of SPSS Modeler, without access to the regular user interface.

Examples of tasks that are appropriate for batch mode:

- Running a time-consuming modeling exercise in the background
- Running a stream at a scheduled time (for example, overnight, when the resulting load on the computer will not affect day shift production
- Running a data preprocessing stream on a large volume of data (for example, in the background or overnight)
- Running regularly scheduled tasks, such as monthly reports
- Running a stream as an embedded part of another process, such as a scoring engine facility

We do not describe the installation steps. For installation details, see *IBM SPSS Modeler Batch Installation for UNIX*:

ftp://ftp.software.ibm.com/software/analytics/spss/documentation/modeler/14.2/en/BatchInstallationForUNIX.pdf

After the IBM SPSS Batch is installed, you can run it by submitting the statement in Example 5-3.

Example 5-3 Executing SPSS Batch

```
clemb -server -hostname myserver -port 80 -username dminer -password 1234 -stream
report.str -execute
```

For a complete list of command-line arguments, see *IBM SPSS Modeler Batch User's Guide:*

ftp://ftp.software.ibm.com/software/analytics/spss/documentation/modeler/14.2/en/BatchUsersGuide.pdf

Related documents:

- *IBM SPSS Modeler Server for UNIX Installation Instructions*

  ftp://ftp.software.ibm.com/software/analytics/spss/documentation/modeler/14.2/en/ServerInstallationUNIX.pdf
- *IBM SPSS Modeler Batch Installation for UNIX*

  ftp://ftp.software.ibm.com/software/analytics/spss/documentation/modeler/14.2/en/BatchInstallationForUNIX.pdf
- *IBM SPSS Modeler Batch User's Guide*

  ftp://ftp.software.ibm.com/software/analytics/spss/documentation/modeler/14.2/en/BatchUsersGuide.pdf
- *IBM SPSS Modeler Server Administration and Performance Guide*

  ftp://ftp.software.ibm.com/software/analytics/spss/documentation/modeler/14.2/en/ServerandPerformanceGuide.pdf

5.4.4 SPSS Modeler data access

If you want to connect to a database with SPSS Modeler components, you must have an ODBC data source installed and configured for the relevant database. Be sure to grant or get the needed permissions.

The IBM SPSS Data Access Pack includes a set of ODBC drivers that can be used for this purpose. We need these drivers, because we use DB2 for z/OS as data source and target.

Supported ODBC drivers

For the most recent information about which databases and ODBC drivers are supported and tested for use with SPSS Modeler 14.2, see the product compatibility matrixes on the corporate support site:

http://www.ibm.com/support

Installation and configuration

The Installation process of the IBM SPSS Data Access Pack is described in the *IBM SPSS Data Access Pack Installation Instructions* and is therefore not described here.

However, we do mention the most important points, which might be overlooked. After you successfully install the appropriate binary, complete the following tasks:

1. Set ODBC home directory to the current directory. Run the following command within the installation directory of the Data Access Pack:

   `./setodbcpath.sh.`

2. Edit the `modelersrv.sh` file to call `odbc.sh` at start.The `odbc.sh` file sets the environment variables for your ODBC drivers, which is important to be done when the server starts. Add the lines, shown in Example 5-4, to your `modelersrv.sh` file (adjusted to your paths).

   Example 5-4 ODBCINI in modelersrv.sh

   ```
   . /opt/IBM/SPSS/ModelerServer/14.2/odbc.sh
   ODBCINI=/opt/IBM/SPSS/ModelerServer/14.2/odbc.ini
   export ODBCINI
   ```

3. Edit `odbc.ini` file to access the data source. Because we use DB2, modify the following section so that the parameters, listed in Example 5-5, match to your infrastructure:

   `[DB2 Wire Protocol]`

   Example 5-5 odbc.ini modifications

   ```
   IpAddress=200.1.1.20
   Location=RDBNDW00
   LogonID=COGSPSS
   Password=XXXXXX
   ```

4. Stop and then restart the IBM SPSS Modeler Server.

Related document: *IBM SPSS Data Access Pack Installation Instructions*

ftp://public.dhe.ibm.com/software/analytics/spss/documentation/statistics/20.0/en/client/InstallationDocuments/Linux/IBM_SPSS_Data_Access_Pack_Installation_Instructions.pdf

5.4.5 Other considerations

Also, be aware of license information, License Manager, tracing, enabling server logging, and levels for release and fix packs.

Single user-license on terminal server

If, during IBM SPSS Client startup, you see the message shown in Figure 5-16, the client is pointing to the wrong license information.

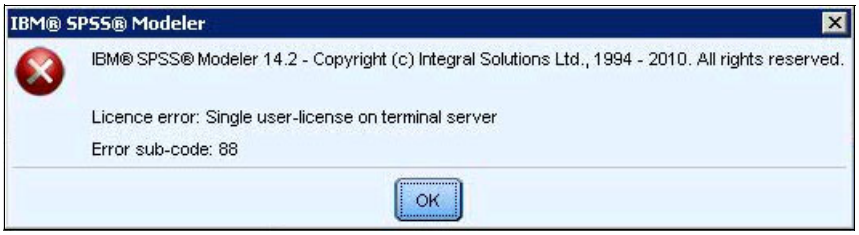

Figure 5-16 Single user-license on terminal server

The following options can help you resolve the issue:

- Make sure you chose the correct license type during your client installation. On a terminal server operating system, you have to select the **Network license** option and provide the appropriate licenses by using Sentinel License Manager.
- Make sure you provided the correct license type in your license manager.
- On the client's local system find the `lservrc` file, open it with a text editor, and delete the content. This file is the first place for the client to check for locale and single user licenses. Entries in this file, because of former evaluation or incorrect license type selection, will avoid the client from connecting to a license server. Restart the Sentinel License Manager Service.

Changing or moving the Network License Manager

If you need to change the referred License Manager in your client configuration, change the referring entry in the `spssprod.inf` file:

1. Change the following entry and point it to your new License Manager location:

   ```
   DaemonHost = 10.6.10.36
   ```

2. Save your changes and restart the IBM SPSS Modeler Client.

Tracing

If you have trouble with your data connection by using the ODBC drivers that are provided by the IBM SPSS Data Access Pack, you can enable tracing. To enable tracing, modify the `[ODBC]` section in the `odbc.ini` file, shown in Example 5-6. Also see Example 5-7 on page 60.

Example 5-6 Tracing section in odbc.ini

```
Trace=[0 | 1]
TraceFile=trace_filename
TraceDll=ODBCHOME/lib/xxtrcyy.zz
ODBCTraceMaxFileSize=file_size
ODBCTraceMaxNumFiles=file_number
```

Example 5-7 General ODBC section with tracing activated

```
[ODBC]
IANAAppCodePage=4
InstallDir=/opt/IBM/SPSS/ModelerServer/14.2
Trace=1
TraceFile=odbctrace.out
TraceDll=/opt/IBM/SPSS/ModelerServer/14.2/lib/XEtrc25.so
```

You can find further detailed information about Progress DataDirect Connect Series for ODBC in the user guide from the manufacturer, which is available at the following location:

http://www.datadirect.com

Enable server logging

IBM SPSS Modeler Server keeps a record of its important actions in a log file named `messages.log`, which is normally located in the following directory:

```
<IBM SPSS Modeler Server installation>/config
```

To enable a several server logging capabilities, complete the following steps:

1. In a text editor, open the `options.cfg` configuration file, which is in the following directory:

   ```
   <IBM SPSS Modeler Server installation>/config
   ```

2. Modify the following entry to set the directory where the log file should be stored:

   ```
   log_directory, "/opt/IBM/SPSS/ModelerServer/14.2/log/"
   ```

3. Modify the following entries to log database access and errors:

   ```
   log_database_access, y
   log_database_errors, y
   ```

4. Modify the following entries to log file access and errors:

   ```
   log_file_access, Y
   log_file_errors, Y
   ```

For further logging options, see the *IBM SPSS Modeler Server Administration and Performance Guide*:

ftp://ftp.software.ibm.com/software/analytics/spss/documentation/modeler/14.2/en/ServerandPerformanceGuide.pdf

Release and fix pack level

The client software must be at the same release level as the IBM SPSS Modeler Server software.

You can download the latest fix pack from the following location

http://www-947.ibm.com/support/entry/portal/Downloads/Software/SPSS/SPSS_Modeler?lnk=mhsd

5.4.6 SPSS Modeler and Cognos TM1 integration considerations

Currently, there is no native or direct data, or process integration of SPSS and TM1. However, this does not mean that there is no way to use SPSS-created data in your TM1 dashboards and reports.

TM1 accessing SPSS data

To access SPSS-created data in your TM1 dashboards and reports, you can use a kind of staging area. You start the SPSS stream processing as you always do, by using IBM SPSS Modeler Client or a batch command. IBM SPSS Modeler Server processes your stream and reads from and also writes to a database, for example. Consider this database as your "staging area." After SPSS finishes stream processing, you can access the calculated and written data of SPSS within this database, with TM1, and display the data in your dashboard and report. See Figure 5-17.

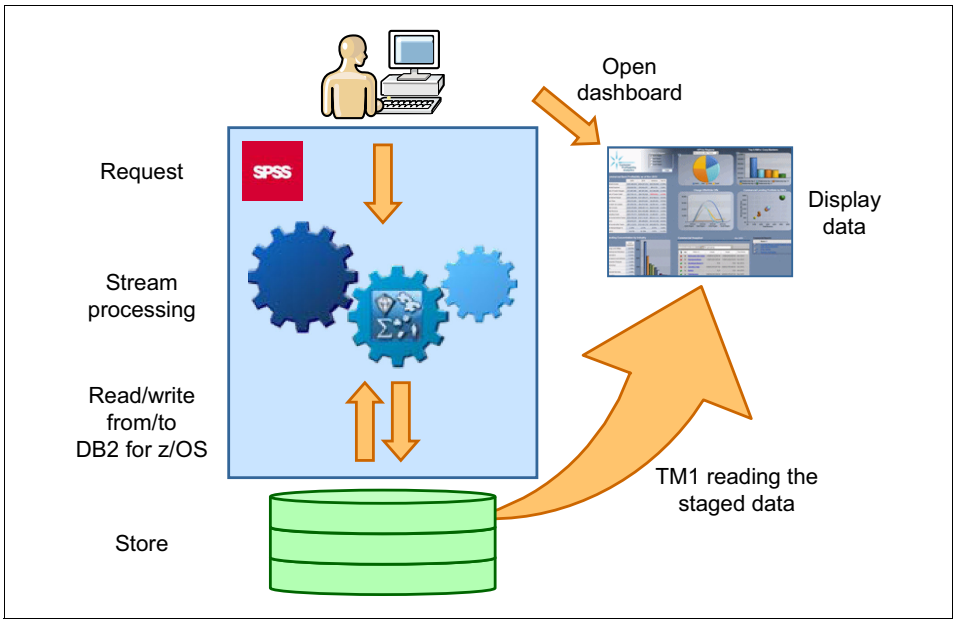

Figure 5-17 TM1 accessing SPSS data

TM1 starting SPSS stream

A TM1 TurboIntegrator (TI) process has a function named `ExecuteCommand`.

This function is valid only in TurboIntegrator processes and executes a command line during a process. It has the following syntax:

```
ExecuteCommand(CommandLine, Wait);
```

The function has the following arguments:

CommandLine Defines the command line you want to execute.

Wait Defines whether the process should wait for the command to complete execution before continuing to the next process statement. It can have the following argument values:

- **0** Proceed without waiting.
- **1** Wait for the command line to successfully execute before proceeding

For details about the TM1 TurboIntegrator functions, see the following location:

http://ibm.co/ovBY5x

Because TM1 is installed on AIX, and the IBM SPSS Modeler Server and IBM SPSS Modeler Batch are installed on Linux on System z, we have to send a command from AIX to Linux on System z.

For this purpose, we can encapsulate an IBM SPSS Modeler Batch command (as described in “Batch feature” on page 57) into an SSH command, which will be started by the TI `ExecuteCommand` function mentioned previously. With the `Wait` parameter, you can be sure to get a feedback when the command line executes successfully.

The logical process is shown in Figure 5-18.

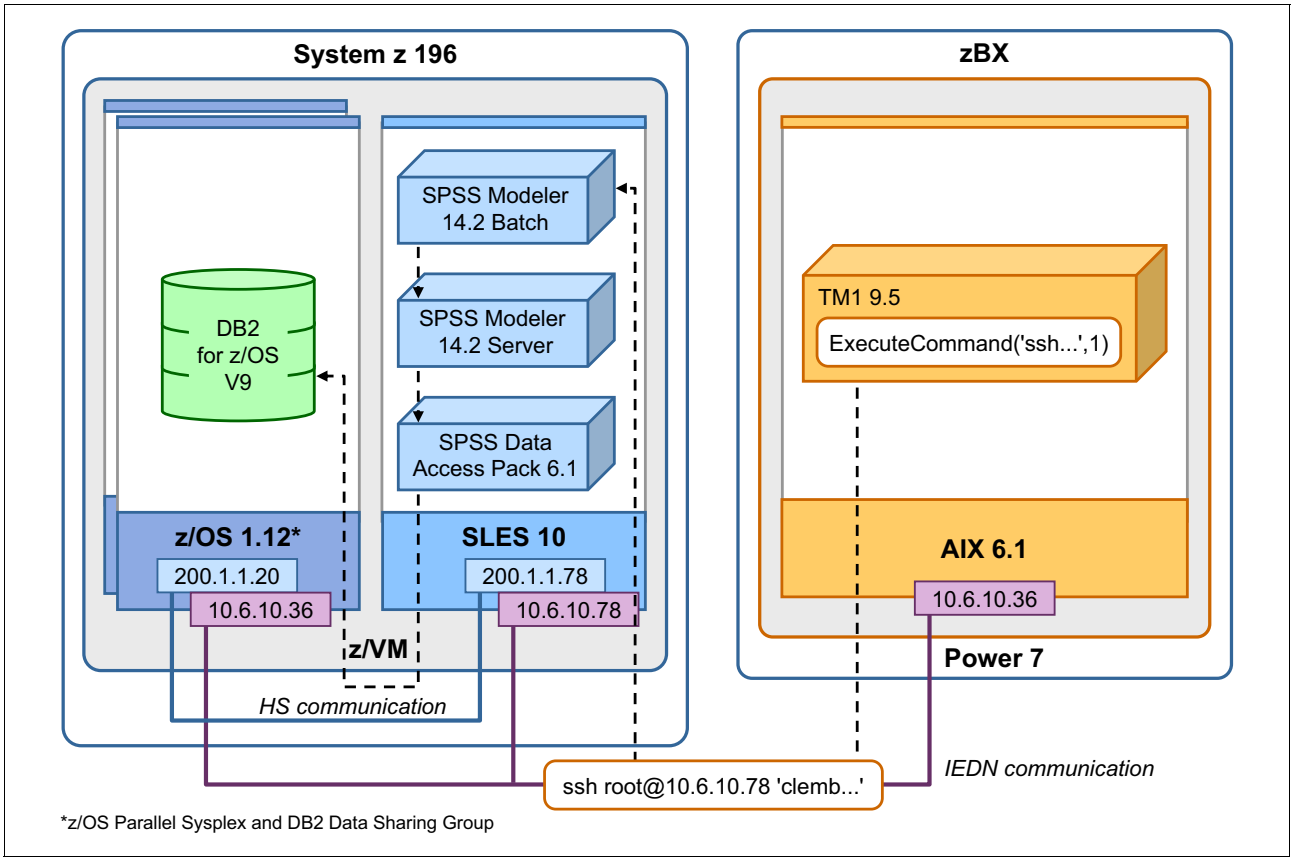

Figure 5-18 TM1 TI ExecuteCommand function calls SPSS Batch command

Example 5-8 shows how this statement might look. In this example, we execute an SSH call from our zBX AIX machine over IEDN (10.6.xx address) to call the **`clemb`** command. We use the fully qualified **`clemb`** command because the `ModelerBatch` directory is not included in our $PATH variable. We still use the standard port and host name of our SPSS Server and tell **`clemb`** to execute the `ProductAssociationByBranch` stream. Wait `1` indicates to TM1 to wait for the successful completion of the command before executing the next step.

Example 5-8 Encapsulated TM1 TI/SSH/SPSS Batch command line

```
ExecuteCommand('ssh root@10.6.10.78 "/opt/IBM/SPSS/ModelerBatch/14.2/clemb -server
-hostname cognspss -port 28051 -username root -password ****** -stream
/spssdata/ProductAssociationByBranch.str -execute"',1);
```

Figure 5-19 shows the result within TM1 Architect, and indicating that completion is successful.

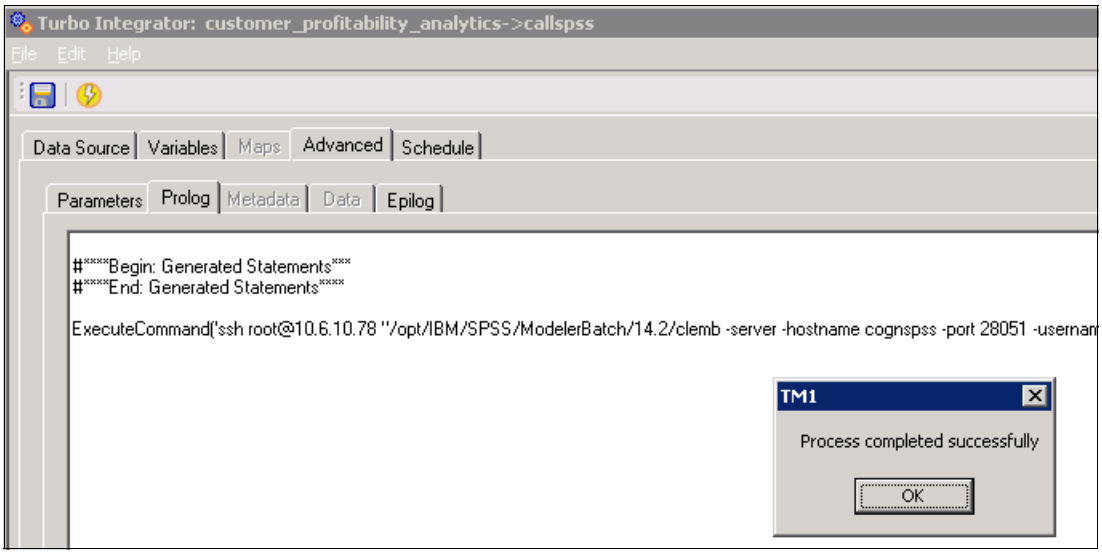

Figure 5-19 TM1 calling SPSS Modeler Batch and completion feedback

Notes:

- If you do not want to type `SSH password` each time you start the TMI process by calling the **`ssh`** command, consider generating and providing appropriate RSA keys.
- If you want to enable logging during **`clemb`** command execution, add the following parameter to the **`clemb`** command:

  `-log <directory/logfile>`

6

Testing the application

To ensure that our new analytics functions would perform well after they were in production, we needed to test a variety of functions to determine whether bottlenecks existed and needed to be resolved.

We chose four scenarios to test all software components of the CPA blueprint by using an existing Rational Performance Test lab:

- Cognos Dashboard scenario
- Metric Studio Scorecard scenario
- TM1 Report scenario
- Cognos Business Insight Report scenario

6.1 Setting up a hybrid test environment

In Chapter 3, “Technical design” on page 11, we detail the environment on which we are implementing the CPA blueprint. This environment includes an existing capability to test applications by using Rational Performance Tester.

Rational Performance Tester is an application for testing of applications, by using repeatable automated scripts. After a testing scenario is planned, Rational Performance Tester can record each scenario and play it back again with a number of variations. The variations can include information input to the application, the number of users, and several other variations. After capturing the relevant scenarios, we see how the system behaves as more users are added by using Rational Performance Tester’s scalable distributed playback.

6.1.1 Capturing a test scenario

Rational Performance Tester has the ability to capture the activities from a browser and record them for later replay. To group various scenarios together, we start by creating a new workspace.

For each scenario, we then choose to create a new performance test from a recording, in this case a Rational Performance Tester HTTP recording. Each scenario is given a unique name to help us track them.

An initializing Recorder window indicates when you are ready to begin, at which point we open a browser. Rational Performance Tester recommends starting by clearing the browser cache and removing any cookies that are associated with the application. We begin the recording by entering the URL of the website from which we are going to start the scenario and proceed through the steps one at a time, making sure to let the pages fully load each during each step. To end the scenario, we simply close the browser.

After the browser is closed, Rational Performance Tester automatically generates the test from what it has recorded, and provides an opportunity to view or edit the new test scenario.

6.1.2 Replaying the scenario

After we finish any editing that might be required to the test scenario, we can run the test to ensure that it completes successfully. A single test can be run only as a single virtual user.

We run the test several times, making sure that we are satisfied with the results and any delays that might be represented.

6.1.3 Scheduling a combined test

To simulate multiple, concurrent users, a schedule is required. By creating a schedule, we can define how many users will be simulated, and how Rational Performance Tester should increase the number and how the users should be grouped. Initially, the schedule contains only users, group, and ramping parameters; however, after the schedule is created, we add the tests that we already captured. Because we want to run all of our scenarios, we add each scenario to the schedule.

We previously selected our scenarios to fit with certain predefined roles. Because a different number of users is represented for each role, we also change the number of users that will run each scenario to better represent how the test will be executed after the application is in production.

The schedule is run by selecting it in Rational Performance Tester and clicking **Run schedule**. While it is running, you can see several reports in real time; these reports represent response statistics and other metrics. After the test is complete, we view the reports to see how well the system performed each scenario and the schedule as a whole.

6.1.4 Hybrid Test environment summary

To simulate a large number of users running various scheduled test scenarios, distributing the tests across multiple Rational Performance Tester test agents might be necessary. Figure 3-3 on page 17 shows the network infrastructure that we used for our implementation, including a BladeCenter that is attached to the administrative network at the top left of the diagram. By running Rational Performance Tester agents on several of the blades in the BladeCenter, we are able to simulate thousands of users that are accessing the CPA test scenarios that follow.

For more information about Rational Performance Tester, see the following resources:

- *Using Rational Performance Tester Version 7*, SG24-7391
- Rational Performance Tester website:

  http://www.ibm.com/software/awdtools/tester/performance/

We decided to capture four scenarios to ensure that all components of the CPA blueprint are tested. We describe each scenario in the following sections.

6.2 Cognos Dashboard scenario

The Cognos Dashboard scenario tests Cognos 10.1 BI, DB2, the TM1 Data Source, and the Cognos Metrics Server. The dashboard, shown in Figure 6-1, provides an integrated view across multiple key dimensions.

The simulated user in this scenario navigates to the user's Cognos home page and opens the dashboard. As the widgets on the dashboard open, data is pulled in from the Metrics Server and from TM1.

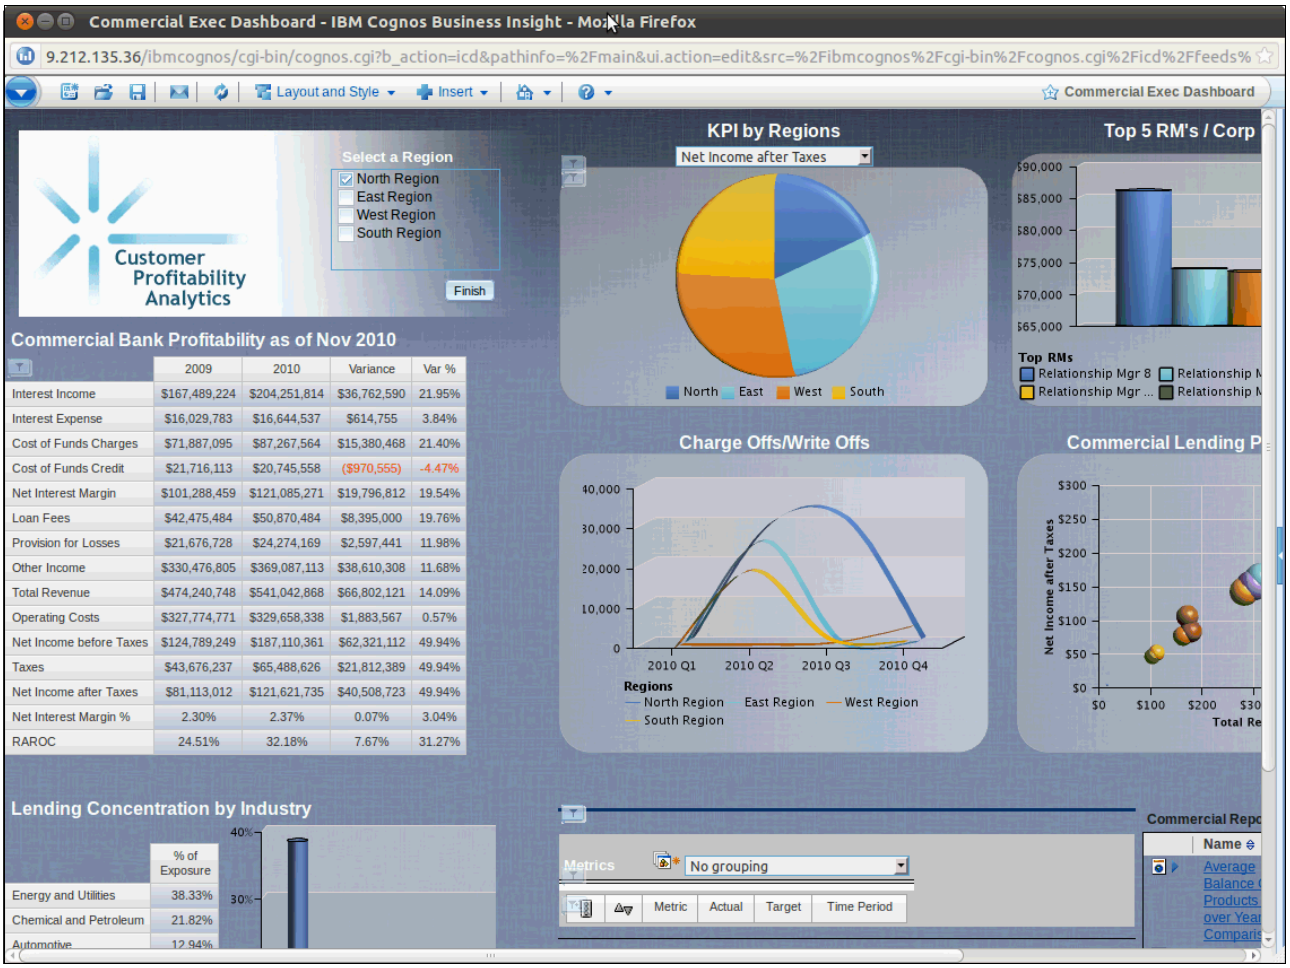

Figure 6-1 CPA dashboard

6.2.1 Test plan

A Cognos BI user opens a browser and navigates to the BI interface. The user opens folders and selects the dashboard.

6.2.2 Running the test

Although this scenario is one of the simpler ones in terms of the number of steps required, this scenario touches almost every major component of the application. The various widgets gather a variety of data in a variety of formats. Overall, this test requires the transfer of the most network data of all the tests. It also puts the heaviest load on the Cognos BI server.

6.3 Metric Studio Scorecard scenario

The Metric Studio Scorecard scenario tests Cognos BI, DB2, and the Cognos Metrics Manager. The scenario, shown in Figure 6-2, displays several key metrics with graphical indicators of their status.

The simulated user in this scenario navigates to the user's Cognos home page and navigates to the Customer Profitability Metrics package. Data is retrieved from Cognos Metrics Manager and DB2.

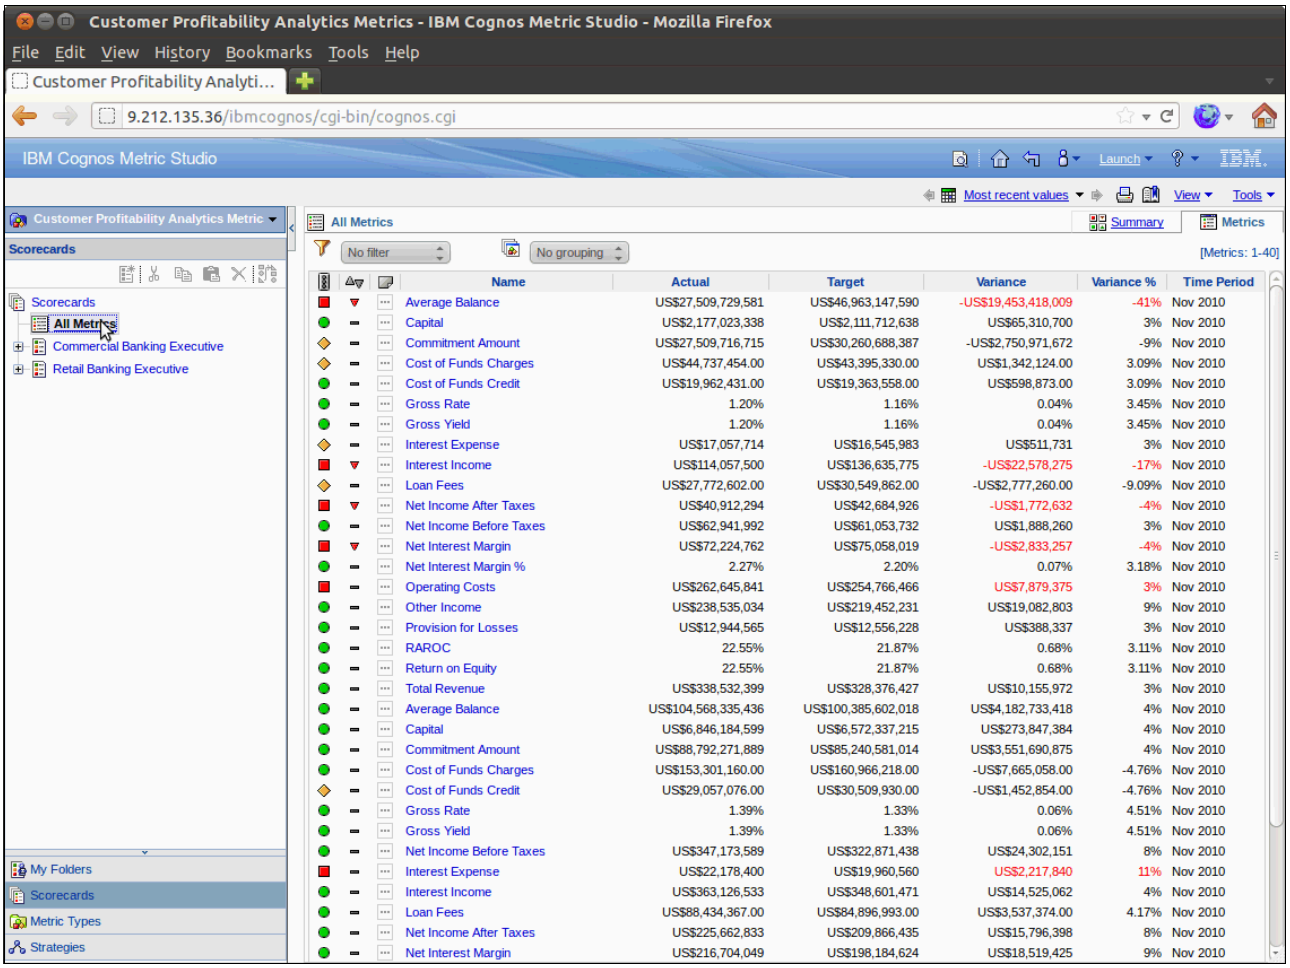

Name	Actual	Target	Variance	Variance %	Time Period
Average Balance	US$27,509,729,581	US$46,963,147,590	-US$19,453,418,009	-41%	Nov 2010
Capital	US$2,177,023,338	US$2,111,712,638	US$65,310,700	3%	Nov 2010
Commitment Amount	US$27,509,716,715	US$30,260,688,387	-US$2,750,971,672	-9%	Nov 2010
Cost of Funds Charges	US$44,737,454.00	US$43,395,330.00	US$1,342,124.00	3.09%	Nov 2010
Cost of Funds Credit	US$19,962,431.00	US$19,363,558.00	US$598,873.00	3.09%	Nov 2010
Gross Rate	1.20%	1.16%	0.04%	3.45%	Nov 2010
Gross Yield	1.20%	1.16%	0.04%	3.45%	Nov 2010
Interest Expense	US$17,057,714	US$16,545,983	US$511,731	3%	Nov 2010
Interest Income	US$114,057,500	US$136,635,775	-US$22,578,275	-17%	Nov 2010
Loan Fees	US$27,772,602.00	US$30,549,862.00	-US$2,777,260.00	-9.09%	Nov 2010
Net Income After Taxes	US$40,912,294	US$42,684,926	-US$1,772,632	-4%	Nov 2010
Net Income Before Taxes	US$62,941,992	US$61,053,732	US$1,888,260	3%	Nov 2010
Net Interest Margin	US$72,224,762	US$75,058,019	-US$2,833,257	-4%	Nov 2010
Net Interest Margin %	2.27%	2.20%	0.07%	3.18%	Nov 2010
Operating Costs	US$262,645,841	US$254,766,466	US$7,879,375	3%	Nov 2010
Other Income	US$238,535,034	US$219,452,231	US$19,082,803	9%	Nov 2010
Provision for Losses	US$12,944,565	US$12,556,228	US$388,337	3%	Nov 2010
RAROC	22.55%	21.87%	0.68%	3.11%	Nov 2010
Return on Equity	22.55%	21.87%	0.68%	3.11%	Nov 2010
Total Revenue	US$338,532,399	US$328,376,427	US$10,155,972	3%	Nov 2010
Average Balance	US$104,568,335,436	US$100,385,602,018	US$4,182,733,418	4%	Nov 2010
Capital	US$6,846,184,599	US$6,572,337,215	US$273,847,384	4%	Nov 2010
Commitment Amount	US$88,792,271,889	US$85,240,581,014	US$3,551,690,875	4%	Nov 2010
Cost of Funds Charges	US$153,301,160.00	US$160,966,218.00	-US$7,665,058.00	-4.76%	Nov 2010
Cost of Funds Credit	US$29,057,076.00	US$30,509,930.00	-US$1,452,854.00	-4.76%	Nov 2010
Gross Rate	1.39%	1.33%	0.06%	4.51%	Nov 2010
Gross Yield	1.39%	1.33%	0.06%	4.51%	Nov 2010
Net Income Before Taxes	US$347,173,589	US$322,871,438	US$24,302,151	8%	Nov 2010
Interest Expense	US$22,178,400	US$19,960,560	US$2,217,840	11%	Nov 2010
Interest Income	US$363,126,533	US$348,601,471	US$14,525,062	4%	Nov 2010
Loan Fees	US$88,434,367.00	US$84,896,993.00	US$3,537,374.00	4.17%	Nov 2010
Net Income After Taxes	US$225,662,833	US$209,866,435	US$15,796,398	8%	Nov 2010
Net Interest Margin	US$216,704,049	US$198,184,624	US$18,519,425	9%	Nov 2010

Figure 6-2 Cognos Metric Studio scenario: Customer Profitability Analytics Metrics

6.3.1 Test plan

The user opens the Cognos BI home page. The user selects the Packages folder and then opens the Customer Profitability Analytics Metrics package.

6.3.2 Running the test

This test requires only the use of Cognos BI, Metrics Manager, and DB2. The amount of rendering required is minimal so less stress is put on the BI server than in the Dashboard scenario. However, much data can be pulled from DB2, which indicates the importance of good bandwidth between the Metrics Server and the DB2 back end.

6.4 TM1 Report scenario

The TM1 Report scenario tests Cognos BI and Cognos TM1. The scenario, shown in Figure 6-3, displays several key metrics with graphical indicators of their status.

The simulated user in this scenario goes to the user's Cognos home page and goes to the selected report. Data is retrieved from Cognos TM1 and a chart is rendered by Cognos BI.

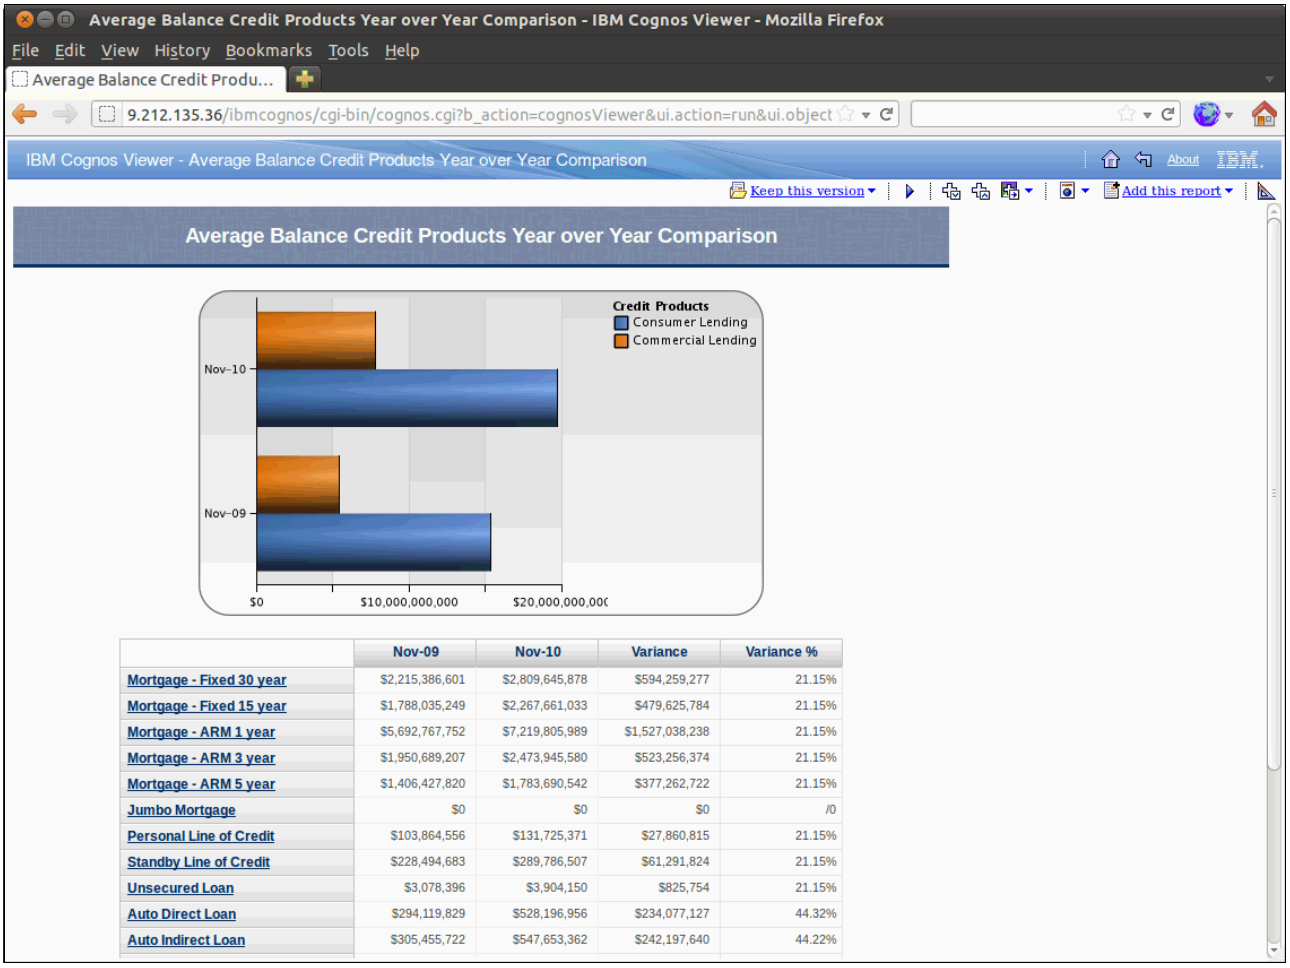

	Nov-09	Nov-10	Variance	Variance %
Mortgage - Fixed 30 year	$2,215,386,601	$2,809,645,878	$594,259,277	21.15%
Mortgage - Fixed 15 year	$1,788,035,249	$2,267,661,033	$479,625,784	21.15%
Mortgage - ARM 1 year	$5,692,767,752	$7,219,805,989	$1,527,038,238	21.15%
Mortgage - ARM 3 year	$1,950,689,207	$2,473,945,580	$523,256,374	21.15%
Mortgage - ARM 5 year	$1,406,427,820	$1,783,690,542	$377,262,722	21.15%
Jumbo Mortgage	$0	$0	$0	/0
Personal Line of Credit	$103,864,556	$131,725,371	$27,860,815	21.15%
Standby Line of Credit	$228,494,683	$289,786,507	$61,291,824	21.15%
Unsecured Loan	$3,078,396	$3,904,150	$825,754	21.15%
Auto Direct Loan	$294,119,829	$528,196,956	$234,077,127	44.32%
Auto Indirect Loan	$305,455,722	$547,653,362	$242,197,640	44.22%

Figure 6-3 TM1 Report scenario

6.4.1 Test plan

The user opens the user's Cognos BI home page, selects the Reports folder and then the Commercial Reports folder, and then selects the Average Balance Credit Products Year over Year Comparison report, which represents a Cognos TM1 cube.

6.4.2 Running the test

Although this test uses only Cognos BI and TM1, it shows how the two can work together to create a meaningful report. In this case, the graph is rendered by using BI's graphics engines, and the data is provided by TM1. Although the test does not stress either system by itself, when combined with other tests, it will compete for BI resources.

6.5 Cognos Business Insight Report scenario

The Business Insight Report scenario tests Cognos BI, DB2, Cognos TM1, and SPSS. The scenario, shown in Figure 6-4, displays predicted performance of several products by region.

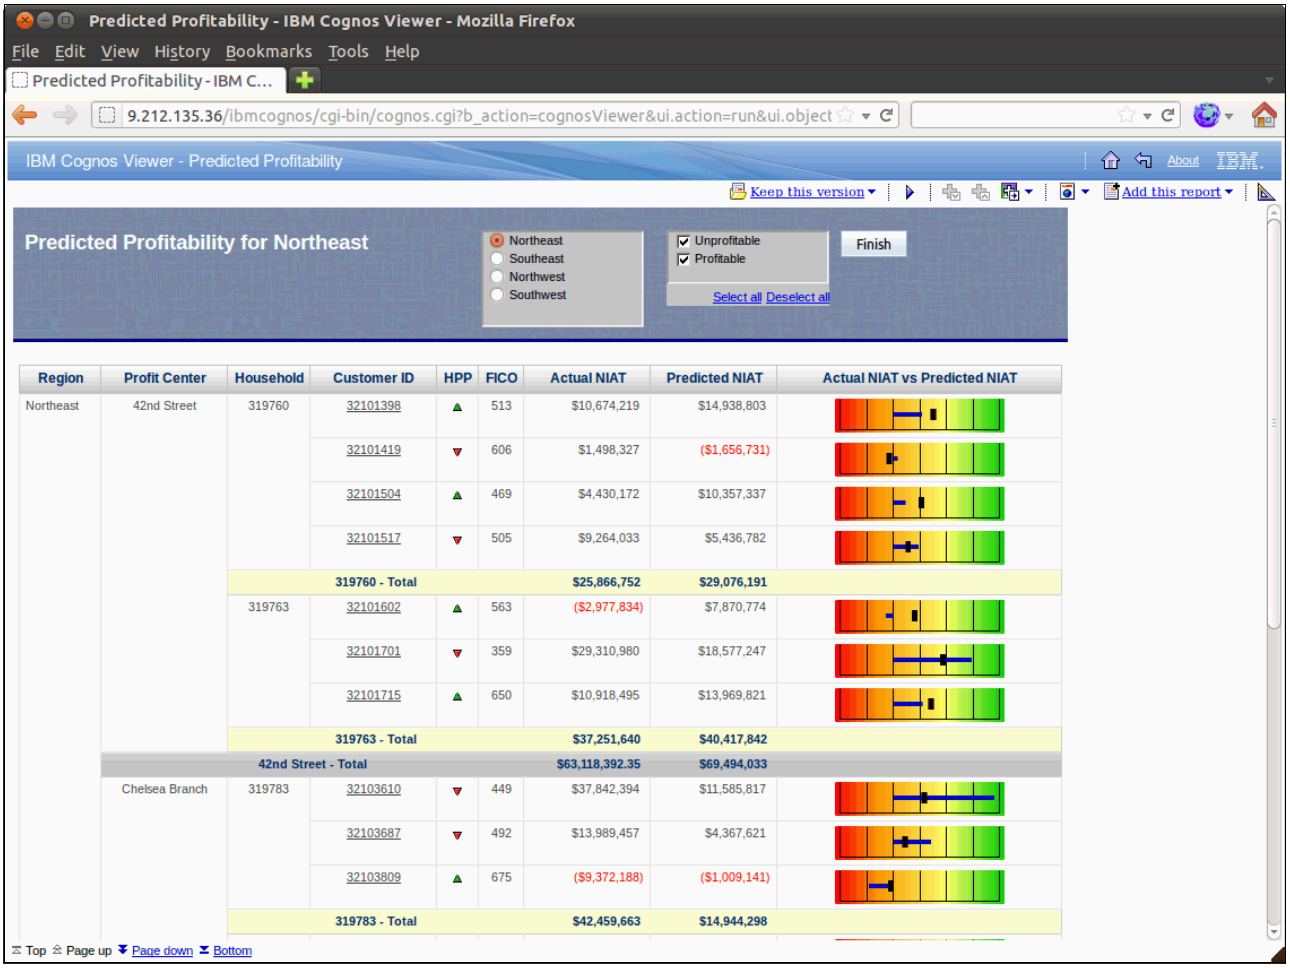

Region	Profit Center	Household	Customer ID	HPP	FICO	Actual NIAT	Predicted NIAT	Actual NIAT vs Predicted NIAT
Northeast	42nd Street	319760	32101398	▲	513	$10,674,219	$14,938,803	
			32101419	▼	606	$1,498,327	($1,656,731)	
			32101504	▲	469	$4,430,172	$10,357,337	
			32101517	▼	505	$9,264,033	$5,436,782	
		319760 - Total				$25,866,752	$29,076,191	
		319763	32101602	▲	563	($2,977,834)	$7,870,774	
			32101701	▼	359	$29,310,980	$18,577,247	
			32101715	▲	650	$10,918,495	$13,969,821	
		319763 - Total				$37,251,640	$40,417,842	
	42nd Street - Total					$63,118,392.35	$69,494,033	
	Chelsea Branch	319783	32103610	▼	449	$37,842,394	$11,585,817	
			32103687	▼	492	$13,989,457	$4,367,621	
			32103809	▲	675	($9,372,188)	($1,009,141)	
		319783 - Total				$42,459,663	$14,944,298	

Figure 6-4 Cognos Business Insight report—predicted profitability

6.5.1 Test plan

The user opens the Cognos BI home page, selects the Reports folder and then the Retail Reports folder. In the folder, the user selects the Predicted Profitability report, which uses data from Cognos TM1 and SPSS.

6.5.2 Running the test

This test pulls together most of the components of the Analytics blueprint. SPSS depends on data from TM1, creates tables in DB2, and depends on BI to render the table. Monitoring during this test demonstrates how each of these components perform, and how each of the underlying networks is used.

6.6 Conclusions

With Rational Performance Tester, we can create individual scenarios that will be used in our production environment and to combine them together to create a full system test. The test exercises several important aspects of the system design and gives us an idea of the area (or areas) that might need improvement.

In our scenarios, Rational Performance Tester uncovered the dependence of each scenario on Cognos BI's graphics engines. To improve performance while running multiple concurrent tests, additional engines were configured, which improved the overall throughput.

7

Review the design points

This chapter has a review of the design for this solution implementation, including solution functionality, compatibility, high availability and scalability, and data proximity.

7.1 Functionality

The IBM Customer Profitability Analytics blueprint that is deployed in our project has three major functions:

- Profitability calculator

  This function provides the bank with a profitability calculator that can efficiently produce customer account profitability statements. Cognos TM1 and Metrics Server provide support for this function.

- Reporting functionality

  This component has a robust reporting capability that shows profitability information to the correct business user at the correct time to enhance business outcomes. the solution includes executive dashboards and scorecards for the bank's major lines of business, and also financial and marketing analyst dashboards. Cognos BI server running on Linux on System z provides powerful reporting, dashboard, and scorecard support for this function.

- Predictive analytics

  This component has the ability to apply predictive analytics to the historical profitability information and customer attribute data to predict what will happen next. SPSS provides support for this function.

7.2 Compatibility

Our project environment has many software products that need to be integrated, including Cognos BI, Cognos Metrics server, Cognos TM1, and SPSS.

7.2.1 Compatibility of Cognos BI and Cognos Metrics Server

We plan to install Cognos BI on Linux and access a DB2 database on z/OS. Meanwhile we want to use scorecarding capabilities in our reports. This function is provided by Cognos Metric Studio.

IBM Cognos Metric Studio (as part of IBM Cognos Business Intelligence) was not yet available in a version for Linux on System z. The Metric Studio Icon/Add-On was available and visible only with a Metrics Server installed in our business intelligence infrastructire.

Therefore, we needed to install Cognos Metrics Server on a supported platform, such as AIX, and so we integrate Cognos BI and Cognos Metrics Server in this way:

1. Install Cognos BI application tier component and Content Manager on Linux. The Content Manager will connect to a content store and other DB2 databases on z/OS.
2. Install Cognos Metrics server gateway, application tier component, and Cognos BI gateway in the same directory in AIX for resource sharing.
3. Cognos Metrics server will share all databases with Cognos BI, so it also will access the Content Manager in Linux.

7.2.2 Compatibility of Cognos BI, Cognos TM1, and SPSS

Cognos BI can use many kinds of data sources for reporting. In this project, we define Cognos TM1 as a data source of Cognos BI. Then, the profitability that is calculated by TM1 can be displayed in Cognos BI blueprint dashboard.

Also, we can define SPSS as a data source of Cognos TM1. Then, Cognos TM1 will import the SPSS predictive results into its cubes for further processing.

7.3 High availability and scalability considerations

High availability is a system design approach and associated service implementation that ensures a prearranged level of operational performance during a contractual measurement period.

Scalability is the ability of a system to adapt to increased processing demands in a predictable way, without becoming too complex, expensive, or unmanageable. As you deploy a system to larger numbers of users, often in separate locations and time zones and with separate language requirements, scalability becomes increasingly important.

7.3.1 High availability and scalability of Cognos BI

In our project environment, we configure each Cognos BI component on one computer simply for proof of concept, ignoring high availability.

All IBM Cognos BI components have built-in failover features to ensure that IBM Cognos BI handles exceptions well.

You can configure components in each of the IBM Cognos BI tiers to enhance availability. As a general rule, make all IBM Cognos BI components available on at least two computers. If the computer on which an IBM Cognos BI component is running fails, another computer takes over.

If, for tuning reasons, you are not running all IBM Cognos BI components on every IBM Cognos BI server, ensure that each component is running on at least two servers. In the event of a computer failure, the remaining component processes requests. Although performance might degrade, there will be service.

Figure 7-1 on page 76 shows a typical Cognos BI high availability architecture. Note these details:

- All Cognos BI components, gateway, dispatcher, and content manger can be installed on two computers to provide high availability.
- Gateways on a web server can balance the workload to all dispatchers that are registered in the active Content Manager. Many kinds of load balancing methods exist.
- Only one Content Manager is active, and others are in standby mode. This failover mechanism works because dispatchers and the active Content Manager routinely communicate with each other. If a dispatcher can no longer reach the active Content Manager, the dispatcher signals a standby Content Manager, which becomes the active Content Manager.
- Content store and other databases reside in a DB2 data sharing group on a z/OS Parallel Sysplex, which provides high availability.

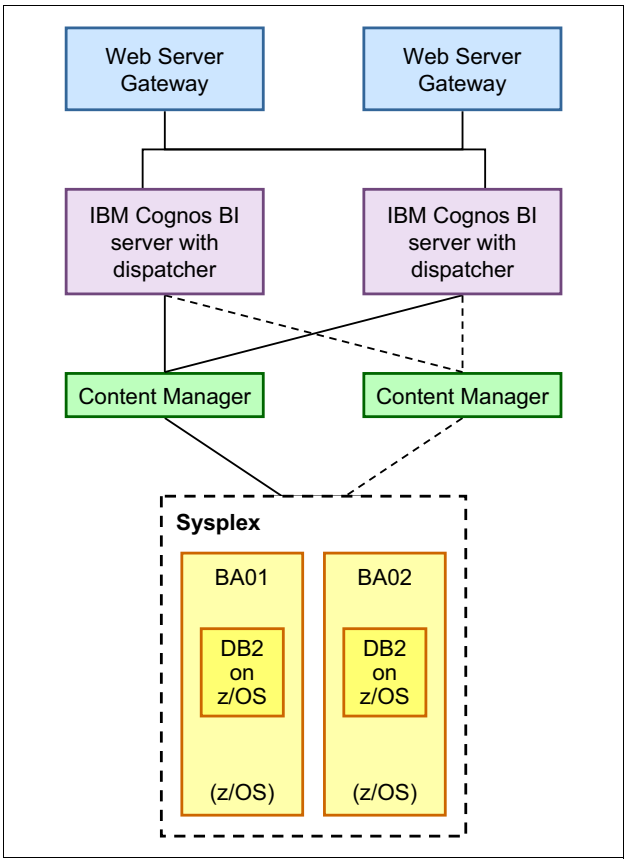

Figure 7-1 Cognos BI high availability architecture

IBM Cognos BI was designed for scalability. It scales in two dimensions:

- Vertical: It scales vertically by using more powerful computers. For example, if you install Cognos BI on Linux, it can scale vertically by using a powerful System z machine.
- Horizontal: You can install and configure Cognos BI components on many computers to get horizontal scalability. The architecture is shown in Figure 7-1.

7.3.2 High availability and scalability of SPSS

From an infrastructure and hardware perspective, System z offers a high level of availability and scalability. In addition, you can benefit from capabilities that SPSS Modeler offers.

When you decide to run a single instance of SPSS Modeler Server, you are dependent on the availability of the physical and logical system that hosts this instance and also on the faultless execution of the instance. Clustering several SPSS Modeler Server instances can improve availability and scalability, and also your load balancing capabilities.

Server cluster: A server cluster is a group of servers that are interchangeable in terms of configuration and resources.

We can think of several clustering scenarios:

- One LPAR directly running Linux on System z with several SPSS Modeler instances
- Several LPARs directly running Linux on System z with several SPSS Modeler instances
- One LPAR running z/VM with 1-n Linux on System z guests, hosting 1-m SPSS Modeler instances.
- Several LPARs running z/VM with 1-n Linux on System z guests, hosting 1-m SPSS Modeler instances.

These various possibilities of running x LPARs with y Linux and z SPSS Modeler Server instances are only a sampling. Having your System z environment running in a Parallel Sysplex can further improve availability.

All of these options require some type of “manager” to control stream executions, load balancing, and instance availability.

With IBM SPSS Collaboration and Deployment Services, a plug-in named the Coordinator of Processes (COP) can be used to manage services in the network. See Figure 7-2. The COP provides these server management capabilities. Each service to be managed by a COP must be registered with the COP upon starting. The service sends updated status messages to the COP periodically. Any necessary configuration files can be stored in the IBM SPSS Collaboration and Deployment Services Repository and retrieved during initialization. The capability of clustering IBM SPSS Modeler Servers has been available in IBM SPSS Collaboration and Deployment Services since version 3.5.

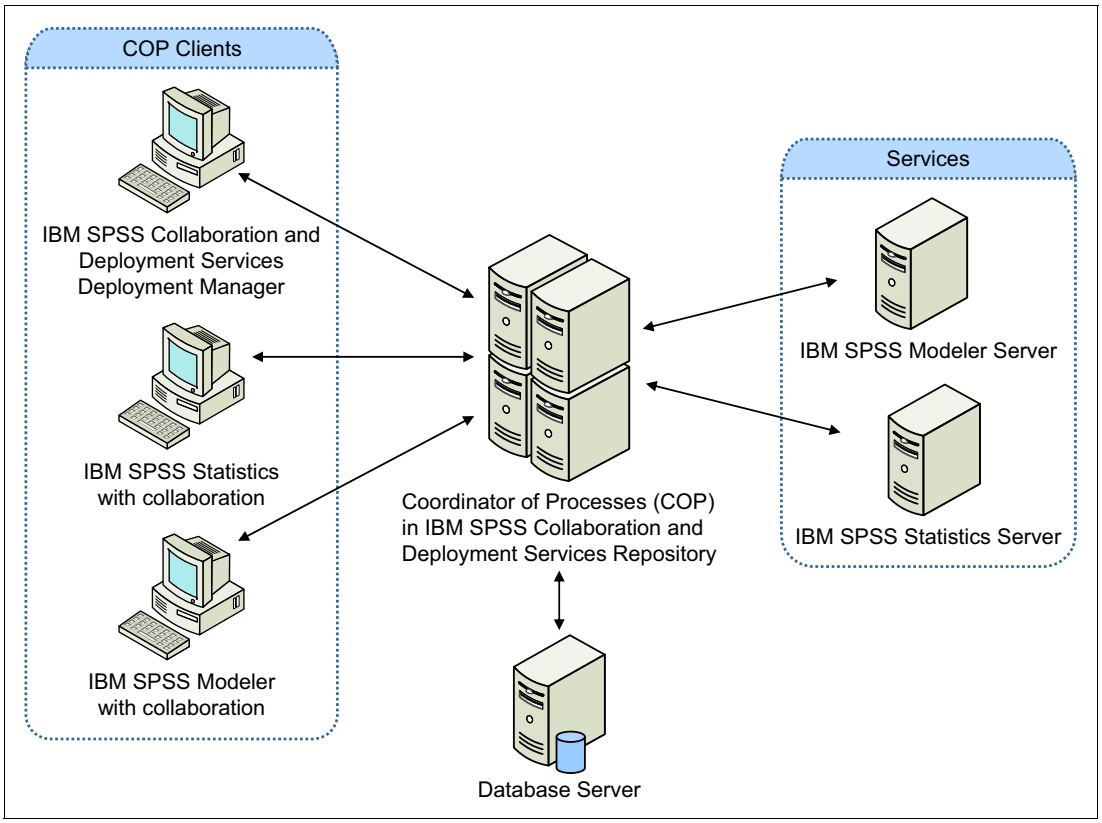

Figure 7-2 The Coordinator of Processes logical architecture

The Coordinator of Processes determines which server is best suited to respond to a processing request, using an algorithm that will balance the load according to several criteria, including the server weights, user priorities, and current processing loads.

For more information, see the *IBM SPSS Collaboration and Deployment Services 5 Coordinator of Processes Service Developer's Guide*:

ftp://public.dhe.ibm.com/software/analytics/spss/documentation/cds/5.0/en/ws/SPSSCopDevGuide.pdf

When you connect to a server or server cluster in SPSS Modeler, you have the option to ask COP to search for a server or cluster for you.

Coordinator of Processes configuration

To register an SPSS Modeler Server instance as a CoP service, you have to modify the configuration settings of the instance. The COP configuration settings panel is shown in Figure 7-3.

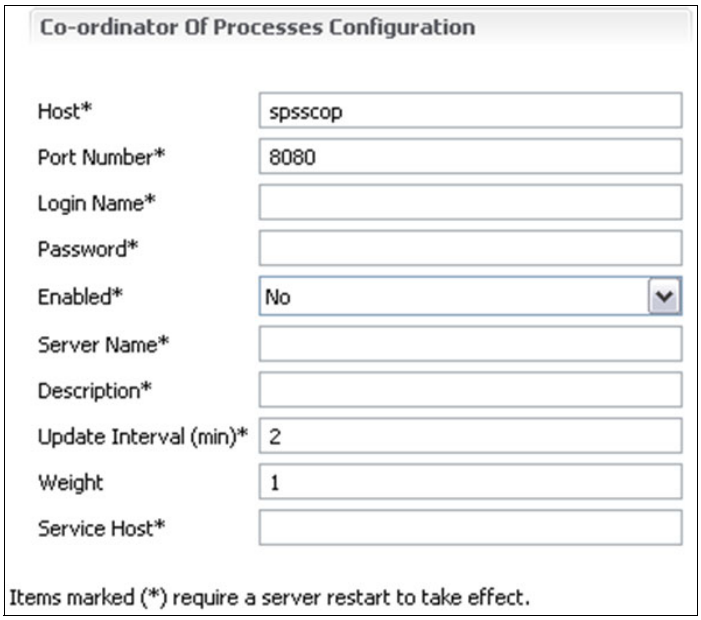

Figure 7-3 The Coordinator of Processes configuration settings

The configuration items are described in Table 7-1.

Table 7-1 Configuration settings

Setting	Value
Host (cop_host)	Host name or IP address of the COP service.
Port Number (cop_port_number)	Port number of the COP service. The default is 8080.
Login Name (cop_user_name)	User name for authentication to the COP service (may include a security-provider prefix).
Password (cop_password)	Encrypted password for authentication to the COP service.

Setting	Value
Enabled (cop_enabled)	Whether the server should attempt to register with the Coordinator of Processes. The default is not to register.
Server Name (cop_service_name)	Name of this server instance.
Description (cop_service_description)	Description of this instance.
Update Interval (min) (cop_update_interval)	Number of minutes between keep-alive messages. The default is 2.
Weight (cop_service_weight)	Weight of this instance. The value is an integer in the range of 1 - 10. A higher weight attracts more connections. The default is 1.
Service Host (cop_service_host)	Fully-qualified host name of the IBM SPSS Modeler Server host. The default is derived automatically.

Related documents:

- *IBM SPSS predictive analytics: Optimizing decisions at the point of impact*, REDP-4710

  http://www.redbooks.ibm.com/abstracts/redp4710.html?Open
- *IBM SPSS Modeler Server Administration and Performance Guide*

  ftp://ftp.software.ibm.com/software/analytics/spss/documentation/modeler/14.2/en/ServerandPerformanceGuide.pdf
- *IBM SPSS Modeler Administration Console User Guide*

  ftp://ftp.software.ibm.com/software/analytics/spss/documentation/modeler/14.2/en/ModelerAdministrationConsole.pdf
- *IBM SPSS Modeler 14.2 Batch User's Guide*

  ftp://ftp.software.ibm.com/software/analytics/spss/documentation/modeler/14.2/en/BatchUsersGuide.pdf

7.4 Data proximity

Data proximity is the relative distance of data placement to the application based on latency. Closer data proximity can reduce latency and improve application response time when accessing data. In our project environment, most of the data is stored in DB2 on z/OS.

- Cognos BI running on Linux accesses DB2 on z/OS through a high-speed HiperSockets network.
- SPSS server running on Linux also can access DB2 on z/OS through the HiperSockets network.

8

Monitoring performance in the zEnterprise hybrid environment

In this chapter, we describe the unique monitoring requirements for zEnterprise hybrid workload performance, how we used a suite of Tivoli monitoring solutions to manage this hybrid environment, and suggestions with additional monitoring tools for future consideration. For this project, we are using the Cognos Business Intelligence and SPSS Modeler as the application workloads to be monitored in a z196 and zBX ensemble environment that is running on the Smarter Banking Showcase System at the IBM Montpellier Data Center.

8.1 Management requirements: What to monitor

In our deployment scenario with the Business Intelligence workload on the zEnterprise environment, to ensure high performance throughput, we identified a set of key performance indicators (KPI) to be monitored at the systems resource level, and at the application and transaction level. The objective is to establish visibility and monitoring into all tiers of Cognos Business Intelligence and SPSS Modeler key components that support the infrastructure with an end-to-end view of the following information:

- Hardware
  - CPC (z196 in this case), zBX, Unified Resource Manager
  - On/Off status, exceptional wait or hung state, capacity statistics
  - LPAR weight definitions versus CPU utilization across zEnterprise Ensemble
  - System z Integrated Information Processor (zIIP) capacity and utilization exceptions that are consumed by the Cognos DB2 database
- Operating systems
  - z/OS, z/VM, Linux on System z, AIX, Windows OS
  - CPU, memory, I/O, paging rate, real and virtual storage utilization, highwater marks, and exception thresholds
  - Workload Manager statistics and service-level indicators
  - z/OS UNIX processes and file I/O for DB2 distributed threads, TCP/IP address space, and BI-related processes
- Networks
  - IEDN, INMN, HiperSockets, TCP/IP connections, Open System Adapter (OSA) performance
  - IEDN, INMN network traffic, OSA interface status, performance statistics within the zEnterprise Ensemble private network, and regular OSA performance for customer data network outside of an ensemble between open systems and z/OS

    Two new OSA types (OSX and OSM) were introduced with the zEnterprise Ensemble. Monitoring the health of the intraensemble data network (IEDN) and the intranode management network (INMN) is critical because the communication and data transmission between CPC and zBX relies on this private networking availability. This new area for monitoring is described in later sections.
  - TCP/IP connection statistics associated by each application, DB2 connect gateway, DB2 subsystems, and Cognos-DB2 distributed threads, exceptional connections rejected or backlogged, TCP/IP buffer utilization thresholds
  - HiperSockets performance statistics between z/VM, Linux on System z, and z/OS, and exceptions

- Databases
  - Cognos Content Store on DB2 z/OS, SPSS on DB2 LUW (Linux on System z), and Cognos Metrics Data Store on DB2 LUW (AIX)
  - DB2 system resource utilization, hit ratio and highwater marks: maximum threads, buffer pools, EDM pool, dynamic SQL statement cache, I/O, locking, logging, and exceptional long running queries
  - DB2 accounting and statistics reports for key transactions and applications
- Middleware
  - WebSphere Application Server on Linux on System z (SLES 10 in our scenario) and AIX, Cognos BI gateway on SLES and AIX, IBM HTTP Server, Cognos Dispatcher, Cognos Content Manager, Cognos Report Server, SPSS Modeler Server, SPSS Modeler Batch, IBM DB2 Connect™ on SLES, Metrics Server and Metrics Server gateway on AIX, Cognos TM1 Architect, and TM1 Client processes on Windows OS
  - JVM monitoring, memory leaks, JDBC performance, exceptional zombie processes, pool threads utilization, throttling, hung threads, run away transactions, end-to-end response time
  - Connectivity between components such as BI gateway, DB2 Connect performance statistics, and abnormal state

In addition to the basic metrics monitoring, three types of usage and objectives are required for ongoing monitoring and tuning efforts:

- System Health Dashboard: Enables real-time monitoring of end-to-end BI infrastructure to detect and identify potential issues before they affect the BI services.
- Problem Determination and Resolution: Highlights exceptional metric data that provides zEnterprise and or BI administrators with hints about what is broken and needs to be fixed within the overall infrastructure.
- Historical Trending: Enables the IT system and BI administrators to easily review the system activity to identify trends versus exceptions and the associated actions to reduce future reoccurrences.

8.2 Current hybrid workload monitoring solution

Given Cognos BI runs across multiple operating systems with multiple web and application tiers, we looked at the current available monitoring solutions that are provided by Cognos and individual domain system monitors. Most provide domain expert monitoring in silo mode without cross-system linkage and automatic discovery for all the components in the zEnterprise and zBX environment. We selected the Tivoli Monitoring Suite because it provides the most coverage with integrated single user interface and top-down monitoring approaches with the end-to-end big picture of all involved OS, databases, and middleware tiers, plus the monitoring capabilities for zEnterprise private networking. The following sections describe our monitoring tips and hints that we learned with the target BI workload on zEnterprise and zBX. In addition, a brief description of our practice with the Unified Resource Manager management functions and Cognos 10 monitoring functions is included.

8.2.1 Systems monitoring deployment approach

In our target test environment, we deployed the following solution suite to meet the basic monitoring requirements for Cognos BI and SPSS hybrid workloads on zEnterprise, within the project schedule:

- IBM Tivoli OMEGAMON® XE on z/OS V4.2.0
- IBM Tivoli OMEGAMON XE for DB2 Performance Expert V4.2.0
- IBM Tivoli OMEGAMON XE for Mainframe Networks V4.2.0
- IBM Tivoli OMEGAMON XE for z/VM
- IBM Tivoli Monitoring Agent for Linux on System z
- IBM Tivoli Monitoring Agent for AIX v6.2
- IBM Tivoli Monitoring Agent for Windows
- IBM Tivoli OMEGAMON Dashboard Edition for z/OS
- IBM Tivoli Composite Application Manager for Application Diagnostics (Monitoring WebSphere Application Server and HTTP Server)

By using the monitoring products, we can collect key performance statistics, track and monitor events in real-time, and send alerts to the designated support staff. The collected metric data and monitored events can feed into the OMEGAMON dashboard, which was customized for the Smarter Banking Showcase. The goal is to use the monitoring tool to automatically detect abnormal situations and take corrective actions or notifications. The following minimum set of key areas were used to set up the thresholds:

- Hardware tier (z196, zBX)
  - CPU utilization (percentage used, percentage free, all servers)
  - Memory utilization (percentage used, percentage free, all servers)
  - Disk utilization (percentage used, percentage free, all servers)
  - Disk I/O (all servers)
- Operating system tier (z/OS, Linux on System z, AIX, Windows)
  - Service availability (Cognos BI, Cognos TM1)
  - Process availability statec(java.exe, HTTP (httpd), cgServer.sh, SPSS Modeler process)
  - Process utilization, DASD and file utilization, I/O rate, paging rate
- Network tier
  - TCP/IP connections, utilization, and highwater mark
  - Open System Adapter (OSA) performance statistics and status
  - Connection backlog rejected, zombies without activities
  - Buffers queued on TCP connections that might cause outages
- Application tiers
  - Cognos
    - HTTP process: Cognos Connection and Cognos gateway (HTTP)
    - Cognos Application Report Server Monitoring and Metrics
    - Cognos Dispatcher and Cognos Report Server (JVM)
  - SPSS Modeler 14.2 on Linux on System z
  - Database
    - Database performance
    - SQL query performance
    - DB2 Connect gateway performance

The monitoring infrastructure flags technical infrastructure events, such as a DB2 database lockout or a network component failure, and sends alert events to the Tivoli Enterprise Portal. Figure 8-1 shows how the Tivoli Enterprise Portal is used to provide a high-level view of the current health of the Cognos Business Intelligence Application running on the zEnterprise Smarter Banking Showcase infrastructure.

Figure 8-1 BI Health Check dashboard

Figure 8-1 displays a business view for BI infrastructure status:

- BI multitier component status
- DB2 Data Warehouse status
- BI transaction arrival rate per second
- BI transaction response time
- Overall CPU usage
- Infrastructure events, for example server failure

8.2.2 Systems monitoring architecture

In this project scope, we deployed IBM Tivoli Monitoring Suite Version 6.2 as the foundation product. The monitoring suite provides a multitier, multi-platform, client/server architecture for maximum flexibility and scalability to correlate and aggregate cross-systems KPIs and systems events.

Tivoli Monitoring V6.2 includes four components:

- Tivoli Enterprise Portal
- Tivoli Enterprise Portal Server
- Tivoli Enterprise Monitoring Server
- Tivoli Enterprise Monitoring Agent

Figure 8-2 shows a high-level view of the architecture showing the OMEGAMON XE and Tivoli Monitoring clients, servers, and agents. The Tivoli Enterprise Monitoring agents for open systems and OMEGMAON agents for z/OS send sampling performance statistic data from the individual subsystems and operating systems to the Tivoli Enterprise Management Server. The Tivoli Enterprise Portal Server then extracts the data from the Tivoli Enterprise Management Serve hub, processes and aggregates the information, and presents it to the Tivoli Enterprise Portal.

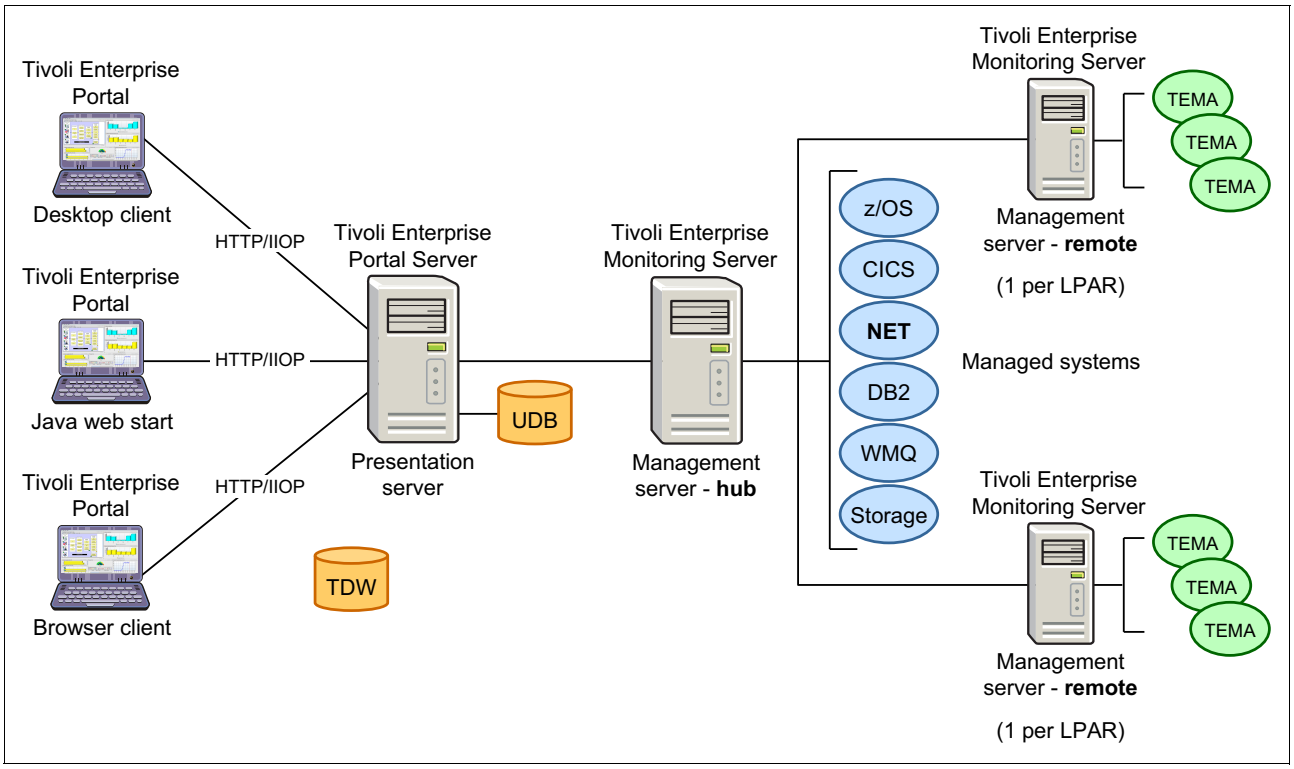

Figure 8-2 Systems monitoring architecture

In our Business Intelligence deployment scenario, we installed corresponding Tivoli monitoring agents for each operating system, DB2 for z/OS, TCP/IP networks, and WebSphere Application Server, to proactively manage the performance of BI infrastructure backbone.

Figure 8-3 shows a mapping of physical component architecture with related monitoring agent.

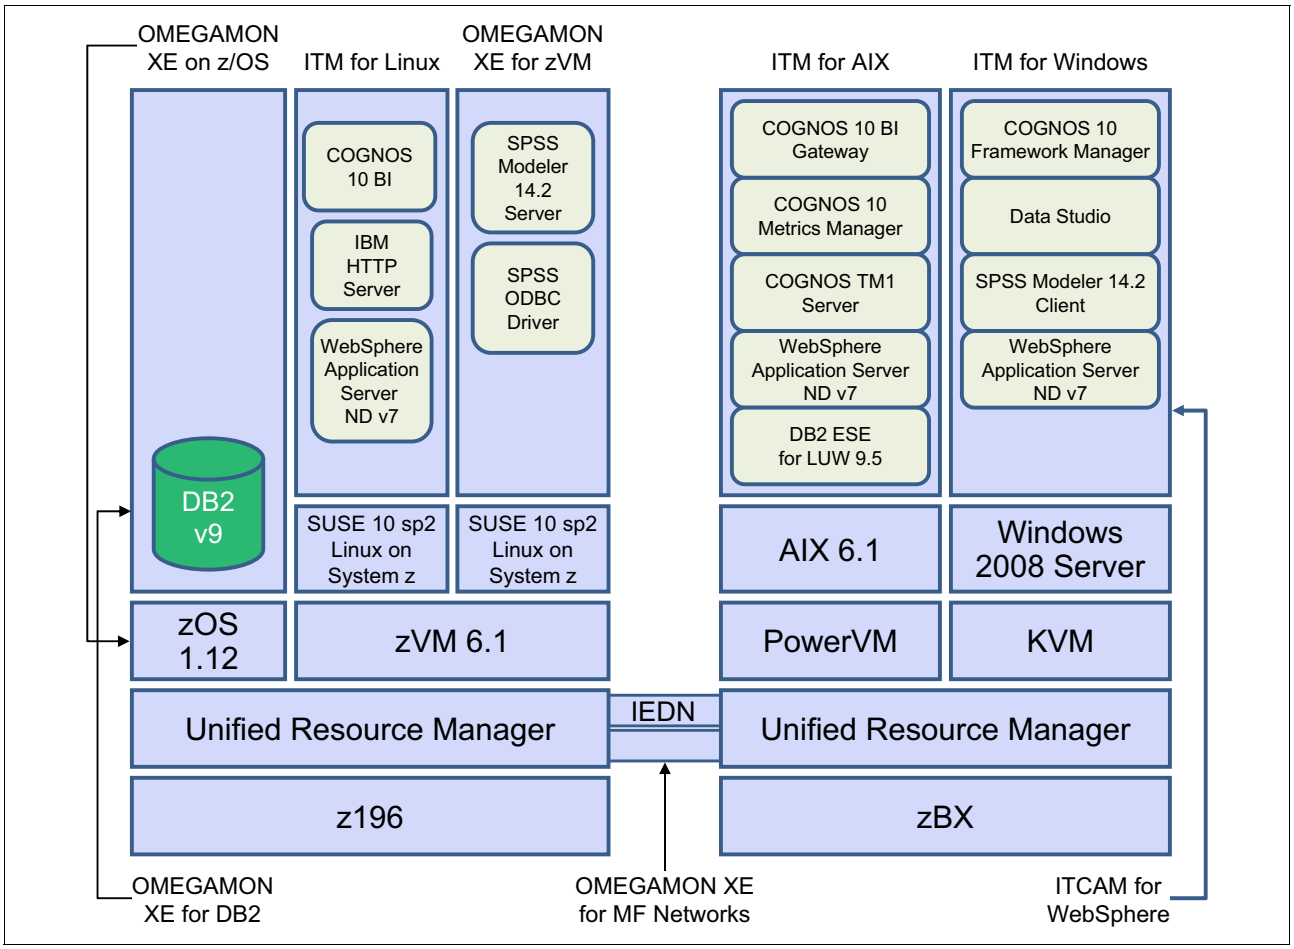

Figure 8-3 Monitoring agents deployed for BI infrastructure

Figure 8-4 illustrates a sample Tivoli Enterprise Portal workspace where it displays the monitoring topology for zBX and graphic view of managed systems. The workspace can be customized and tailored for various physical topology views or business logical presentations.

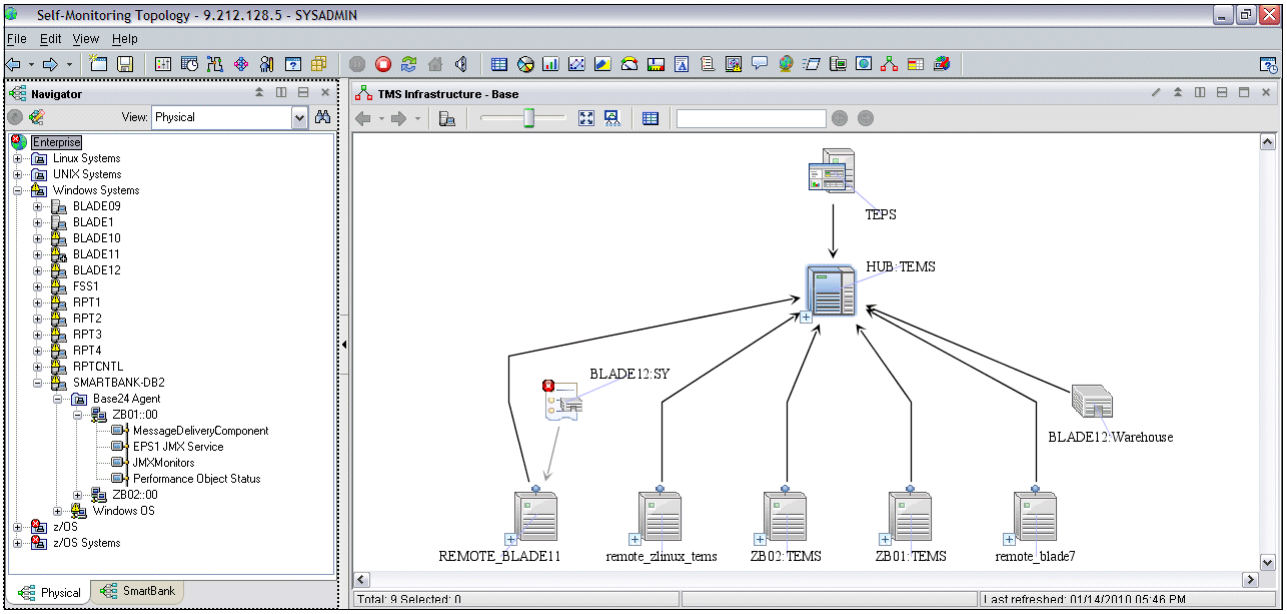

Figure 8-4 Sample zBX monitoring topology

8.2.3 Monitoring operating systems

In this section, we introduce a top-down monitoring approach to managing the zEnterprise hybrid workload performance. We describe key performance indicators under z/OS, z/VM, Linux, AIX, and Windows that can quickly identify potential resource bottlenecks in a zEnterprise ensemble environment.

In general, the performance management functions that are available for use in a zEnterprise runtime environment depend on the suite associated with the Unified Resource Manager. Through a workload performance policy, you can define goals that the Unified Resource Manager uses to manage virtual servers:

- For virtual servers that run z/OS, you can apply these performance goals to the work requests that z/OS and its subsystems support.
- When the z/OS environment includes properly configured guest platform management providers and instrumented IBM middleware products, you can map ensemble workload performance goals to z/OS Workload Manager (WLM) service classes, thus achieving end-to-end goal-based performance management for multitier applications.

In this project scenario, for example, the Cognos BI application workload was defined to Unified Resource Manager as `WKL_SmarterBankingShowcase_Dev_BIWorkload`. Through the top-down logical tuning approach, we look at the Unified Resource Manager performance index first, to see whether the service level for this hybrid workload was achieved. A performance index that is greater than 1 implies that potential slow down or problems were hidden among the underneath components.

We then look at each operating system key performance indicators, including CPU, I/O and demand paging rate, and WLM under each OS. For z/OS, OMEGAMON XE on z/OS can be used to monitor its performance index and detect whether the z/OS WLM service class associated with this workload has achieved the goals. This approach works well, especially

when the workload service policies match between Unified Resource Manager and WLM under z/OS, provided that the z/OS WLM administrator defines new WLM service classes and classification rules for this hybrid work request consistently with the goals defined in the Unified Resource Manager. When the work request enters z/OS, WLM correlates the information about the work request with the work qualifiers in the defined classification rules. Optionally, specifying WLM report classes to monitor the performance of Unified Resource Manager work requests separately from local z/OS work is suggested.

For z/VM, the Unified Resource Manager and the z/VM Resource Manager can manage CPU resources among z/VM virtual servers, but only if the administrator enables CPU management through ensemble-management HMC. CPU management is supported for z/VM hypervisors running on the z196. OMEGAMON XE for z/VM can add value by displaying the CPU and memory resource utilization, and HiperSockets utilizations to help detect potential resource contention or shortage.

For Linux on System z, in addition to Unified Resource Manager management and monitoring features, Tivoli Monitoring Agent for Linux was deployed for this project. The workspace from Tivoli Monitoring for Linux (Figure 8-5) shows multiple processes CPU utilization by Cognos BI.

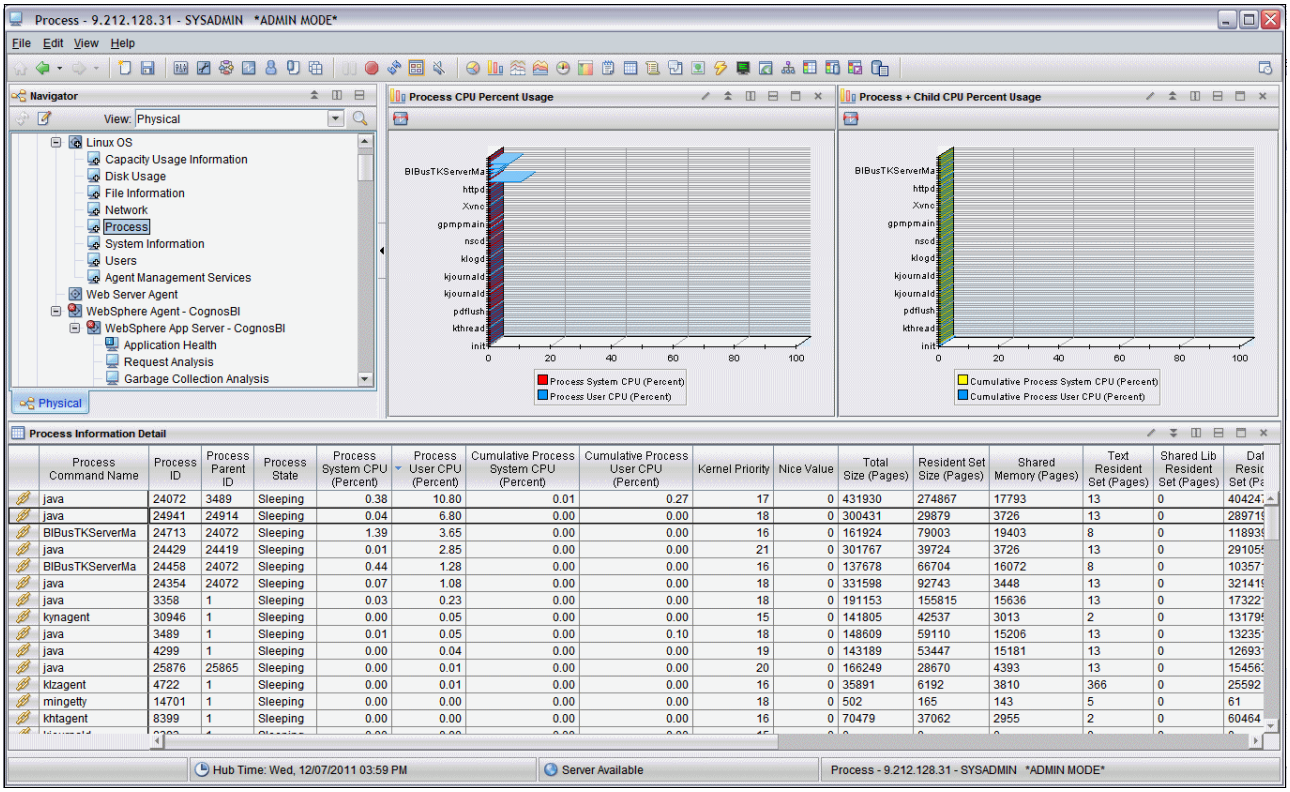

Figure 8-5 Cognos BI Processes CPU utilization under Linux

For AIX and Linux in a zBX environment, IBM PowerVM® Enterprise Edition virtualizes the blades with selective IBM POWER7® and IBM System x blades. PowerVM handles all the access to the hardware resources, providing a virtual I/O server function and the ability to create logical partitions. IBM PowerHA® performs the major functions at the cluster level including manage and monitor OS and HW resources, manage and monitor application processes, manage and monitor network resources, automate application procedures (start, stop, restart, move). Tivoli Monitoring Agent for AIX and Windows operating system added centralized event alerting capabilities and associated application processes and log file monitoring functions with a single monitoring window: the Tivoli Enterprise Portal. This way

helps Unified Resource Manager with more detailed performance data for diagnosing problems faster, and analyzing root cause of problems. Figure 8-6 shows an example of centralized alert management from a customized workspace on Tivoli Enterprise Portal.

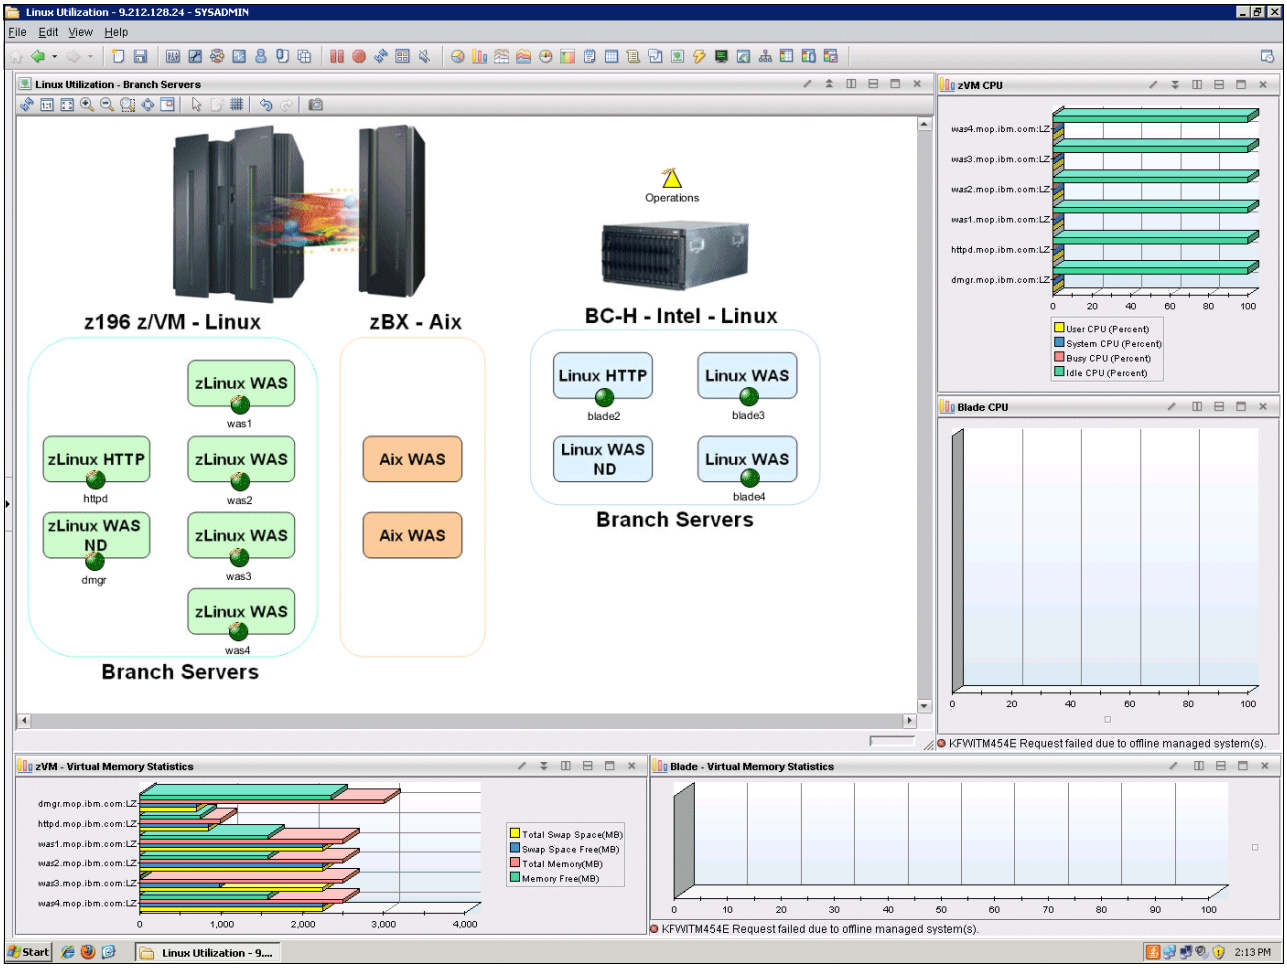

Figure 8-6 Customized dashboard for hybrid operating systems alerts

8.2.4 Monitoring zEnterprise networks

To monitor IEDN performance and whether customer data network affects IEDN or INMN throughput because of network bandwidth or congestions, we set up the situation thresholds through OMEGAMON XE for Mainframe Networks. It monitors the new OSA adapter OSX type and sends alerts if the connection is in an abnormal state. See Figure 8-7 on page 91.

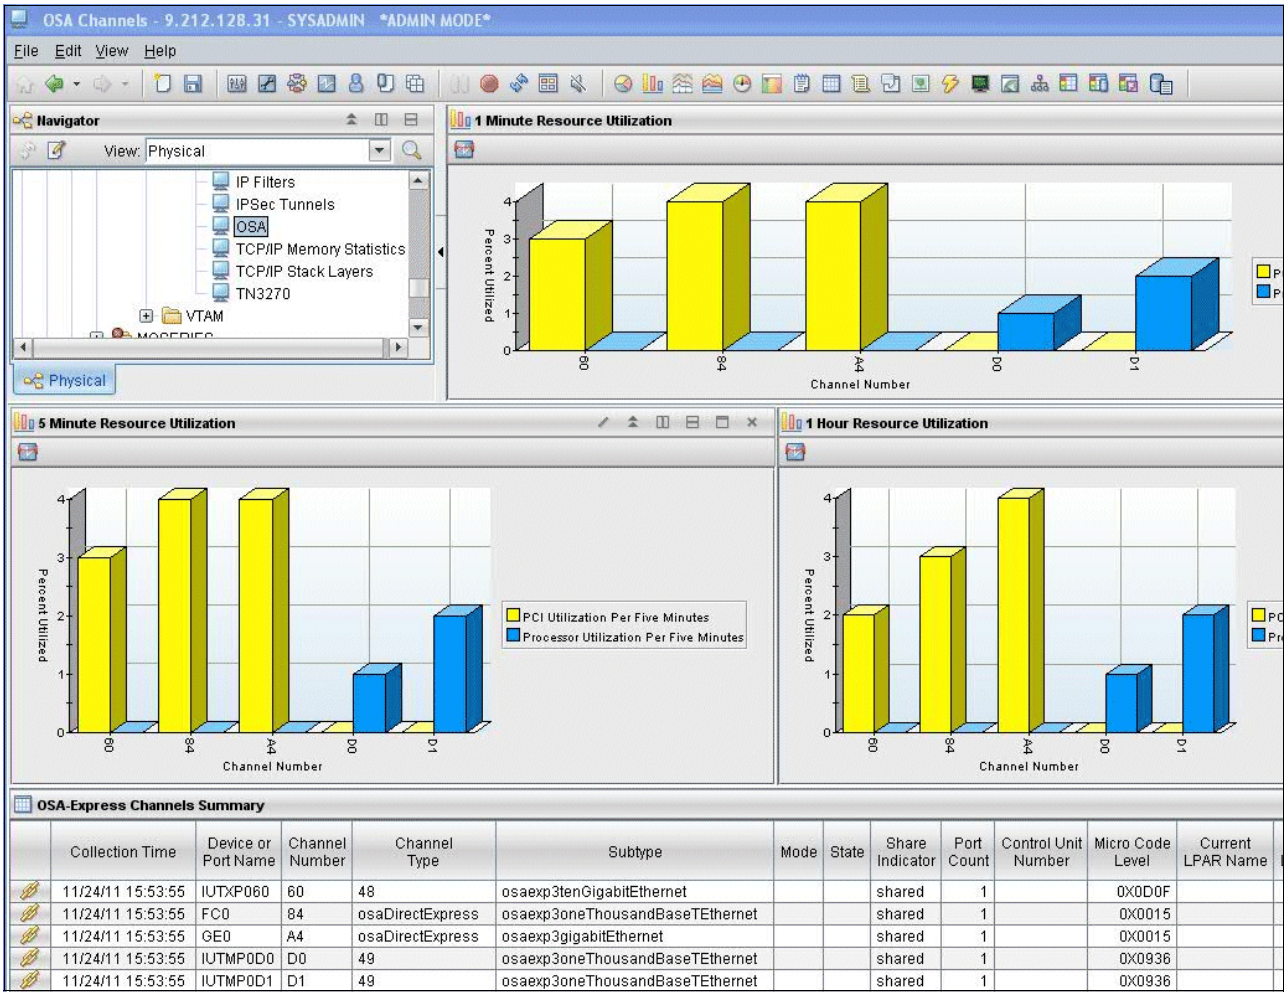

Collection Time	Device or Port Name	Channel Number	Channel Type	Subtype	Mode	State	Share Indicator	Port Count	Control Unit Number	Micro Code Level	Current LPAR Name
11/24/11 15:53:55	IUTXP060	60	48	osaexp3tenGigabitEthernet			shared	1		0X0D0F	
11/24/11 15:53:55	FC0	84	osaDirectExpress	osaexp3oneThousandBaseTEthernet			shared	1		0X0015	
11/24/11 15:53:55	GE0	A4	osaDirectExpress	osaexp3gigabitEthernet			shared	1		0X0015	
11/24/11 15:53:55	IUTMP0D0	D0	49	osaexp3oneThousandBaseTEthernet			shared	1		0X0936	
11/24/11 15:53:55	IUTMP0D1	D1	49	osaexp3oneThousandBaseTEthernet			shared	1		0X0936	

Figure 8-7 Monitoring zEnterprise IEDN and OSX status

You can quickly see how the IEDN and HiperSockets networks are running, and can detect any abnormal situations for Cognos BI and SPSS tasks between zBX and z196, and z/OS and Linux on System z. Figure 8-8 shows the OSA Adapter interface monitoring for the zEnterprise environment.

Description	Interface Type	Physical Address	Current State	Time Entered Current State	MTU Size	Interface Capacity
Loopback	Loopback		Up	11/14/11 16:20:37	65535	0
Loopback	Loopback		Up	11/14/11 16:20:37	65535	0
OSA-Express_QDIO_Ethernet_OSM	OSA_QDIO_ethernet_OSM	0200779...	Up	11/14/11 16:20:39	1500	1000
OSA-Express_QDIO_Ethernet_OSM	OSA_QDIO_ethernet_OSM	0200778...	Up	11/14/11 16:20:39	1500	100
OSA-Express_QDIO_Ethernet_OSD	OSA_QDIO_ethernet_OSD	00145E...	Up	11/14/11 16:20:37	8992	1000
OSA-Express_QDIO_Ethernet_OSD	OSA_QDIO_ethernet_OSD	00145E...	Up	11/14/11 16:20:37	8992	1000
HiperSockets	Hipersocket		Up	11/14/11 16:20:37	8192	0
OSA-Express_QDIO_Ethernet_OSX	OSA_QDIO_ethernet_OSX	02BA3A...	Up	11/14/11 16:20:37	8992	10000
OSA-Express_QDIO_Ethernet_OSX	OSA_QDIO_ethernet_OSX	02BA3A...	Up	11/25/11 12:36:40	8992	10000
Multipath_Channel_Point-to-Point_Samehost	MPC_ptp_samehost		Down	11/14/11 16:20:41	0	0
Multipath_Channel_Point-to-Point_XCF	MPC_ptp_xcf		Up	11/14/11 16:20:41	55296	0
Multipath_Channel_Point-to-Point_XCF	MPC_ptp_xcf		Up	11/14/11 16:26:58	55296	0

Figure 8-8 Monitoring OSA and HiperSockets

In addition to OSA Interface monitoring, OMEGAMON XE for Mainframe Networks can monitor any application that has a TCP/IP connection open on z/OS. This information can be used to determine the volume of data being requested from WebSphere or DB2 distributed threads, for example. When too many concurrent BI requests are flooding into the DB2 z/OS system, if the total number of connections exceeds the maximum allowed, z/OS Communication Server will reject the requests and put them in the backlog queue. The alert notification for excessive backlog rejections can help the BI administrator with early warnings. Figure 8-9 shows an example of the OMEGAMON workspace display and the number of backlog connections that were rejected. If the number gets high consistently over a period of time, the root cause should be investigated to eliminate such bottlenecks.

Listeners with Connections in Backlog or Rejected

System ID	Application Name	Local Port	Connections in Backlog	Backlog Limit	Backlog Connections Rejected	Total Backlog Connections Rejected	Backlog Connections Rejected Time Stamp
MVSB	PORTMAP1	111	0	2	0	6	08/18/11 14:54:40
MVSB	DB2SDIST	5446	5	10	0	0	
MVSC	PORTMAP1	111	0	2	0	6	08/18/11 14:11:03
MVSA	CXEGDSST	1920	0	10	0	7	08/20/11 10:05:35
MVSA	PORTMAP1	111	0	2	0	6	08/18/11 14:35:27

Figure 8-9 Detecting hidden IP connections with backlog rejected

Tips to activate OSA Monitoring: By installation default, OSA data collection is not activated in OMEGAMON XE for Mainframe Networks. Make sure to start the SNMP subagent and configure the TCP/IP profile. For details, see the *IBM Tivoli OMEGAMON XE for Mainframe Networks: Planning Configurations*, SC27-2384 and at the Fix Pack 3 maintenance level or later.

During our project testing phase, we discovered that the OMEGAMON Mainframe Networks did not display OSA adapter information by the default installation. Several steps must be implemented before the data can be displayed in the OMEGAMON OSA workspace:

1. Use the Installation and Configuration Assistance Tool (ICAT) to enable OSA data collection.
2. Verify that the Licensed Internal Code of the OSA adapters meets the requirements that are listed in the OMEGAMON XE for Mainframe Networks Program Directory.
3. Configure and start the SNMP subagent. The SNMP subagent for OSA-Express Direct SNMP subagent is started by using the cataloged procedure, IOBSNMP. A sample IOBSNMP procedure is in the z/OS Communication Server target data set `hlq.SEZAINST`.
4. Configure your TCP/IP profile to match your OSA configuration. If you are using the TCP/IP SNMP subagent for OSA adapter data, specify OSAENABLED and OSASF on the SACONFIG statement in your TCP/IP profile. Reserve a port on the PORT statement, as in the following example:

   ```
   SACONFIG OSAENABLED OSASF 721 COMMUNITY public
   ```

5. Define the OSA adapter to TCP/IP by using either DEVICE and LINK, or INTERFACE statements.
6. If you are using the OSA-Express Direct SNMP subagent for OSA adapter data, install the OSA MIB in the `MIBS.DATA` data set.

For more information about configuring and starting the OSA subagent, see chapter 3 of *IBM Tivoli OMEGAMON XE for Mainframe Networks: Planning Configurations*, SC27-2384.

8.2.5 Monitoring Cognos BI performance on zEnterprise

To ensure optimal application performance across heterogeneous systems on zEnterprise, we used several Tivoli monitoring agents to manage Cognos BI multitier architecture for the targeted BI application:

- Presentation or web tier: IBM HTTP Server and Cognos gateway

  The web tier provides user session connectivity to Cognos BI applications. The IBM Cognos components that fulfill this role are referred to as the IBM Cognos gateway. This component manages all web communications for the IBM Cognos Platform. The workload on the gateway requires minimal processing resources. Depending on the specific application requirements, monitoring and adjusting settings to optimize the HTTP web server performance is important. For high availability or scalability requirements, you can also deploy multiple redundant gateways along with an external HTTP load-balancing router.

  During the testing phase, IBM Tivoli Monitoring for Linux helped monitor the number of connections, and we learned the setting for `maxProcessors` needed to be modified from the default value of 75 to 1000 connections. This step helped improve the performance and scalability. This parameter determines the number of available threads that can be handled by the HTTP process. Also, changing the acceptCount from 100 to 500 helped.

- Application tier (JVM): Cognos Dispatcher, Report Server, and Content Manager components with a collection of Java services

  The application tier for the IBM Cognos Platform is made up of three main components:

  – IBM Cognos Dispatcher
  – IBM Cognos Report Server
  – IBM Cognos Content Manager

  The *Cognos Dispatcher* does the load balancing of requests at the application tier. This dispatcher component is a lightweight Java servlet that manages application services. At startup, each Cognos Dispatcher registers locally available services with the Cognos Content Manager. You can tune the performance of Cognos Platform by defining how the dispatcher handles requests and manages services.

  For this project, we followed a normal configuration for Cognos Dispatcher with two Cognos Report Server processes for each processor, and 8 - 10 threads per processor in four low-affinity threads plus one high-affinity thread. High-affinity connections are used to process absolute high-affinity requests from the report services and resource consumption is optimized for accessing cache when the request is routed back to the report service process that was used to execute the original process. The Tivoli Monitoring Agent can display the performance statistics at the process level. To monitor the process at a detailed thread level, we used Cognos Administration Console to display and manage the number of threads per reporting service process and reconfigure the connection settings when required.

  The main service for application-tier processing is the *report or query service* (often referred to as the *BIBusTKServerMain processes*). The Cognos Dispatcher starts *Report Server* processes dynamically as needed to handle the request load. An administrator can specify the maximum number of processes that these services can start, and also the minimum number of processes that should be running at non-peak times. In general, Cognos BI reporting service performance is closely tied to processor-clock speed and throughput capabilities. The number of processors in a server and their clock rates are the two primary factors to consider when planning for additional Report Server hardware capacity.

When configuring the Cognos application tier, you must set a Java heap size. The Cognos BI reporting and query service is made up of two underlying components:

- Java servlet-based Cognos Dispatcher services
- Report services that are launched using the Java Native Interface (JNI)

According to the Cognos configuration guideline, set the *JVM heap-size allocation* for Cognos Platform so that Java memory is only as large as necessary to accommodate the processing requirements of the Java based services. This setting ensures that as much memory as possible is available to the Cognos Report Service, which is not Java. You can determine the optimal Java heap size by using Java garbage collection statistics. For this project, we used IBM Tivoli Composite Application Manager (ITCAM) for WebSphere monitoring agent to track JVM and Garbage Collection statistics. Note that the JVM memory includes both JVM heap (where Java classes and objects are) and native heap (for non-Java native libraries, such as DB2 libraries or Cognos native-code parts). Understanding this equation when tuning 31-bit/32-bit Cognos is important. The only way to increase the native heap is by decreasing the JVM heap setting. We did some basic experiments with JVM sizes but found the Cognos starting guidelines for JVM heap size to work best for our workloads:

- Content Manager: 512 MB
- Report Servers: 768 MB

We also increased the web threads session pool from 50 (default) to 500, allowing them to exceed the limit when required.

The third key component in the application tier is the *Cognos Content Manager*, which manages the storage of the customer application data including security settings and configurations, server configuration settings, packages, dashboards, metrics, report specifications, and report output. You can use the Cognos Content Manager to publish models, retrieve or store report specifications, handle scheduling information and manage the security name space. The Cognos Content Manager maintains information in the content store database. A minimum of one Cognos Content Manager service is required for each Cognos Platform implementation. Content Manager performance can benefit from the availability of high-speed RAM resources and typically requires one processor for every four processors that are allocated for IBM Cognos Report Server processing.

Query performance is typically bound to the performance of the server that hosts the data source for IBM Cognos BI. Relational data sources that are tuned to meet the access requirements usually perform better. During the testing phase, we also learned to take advantage of IBM Tivoli Monitoring real-time monitor to proactively capture abnormal situations or events by watching the IBM Tivoli Monitoring for Linux workspaces, which displayed warnings or informational messages.

- Data tier: Cognos Content Store, Security Namespace, and Data Source

  When we ran a DB2 content store and multiple instances of WebSphere Application Server for Cognos BI on the Linux, multiple JVMs and Java processes were observed under IBM Tivoli Monitoring for Linux agent “Processes” workspace. During the initial testing phase, we saw numerous warning messages about increasing too many open files on the Tivoli Enterprise Portal. IBM Tivoli Monitoring quickly detected problems and got our attention to determine how to correct the system-hang problem. See Figure 8-10 on page 95.

Log Analysis

Severity	Message ID	Message Text
Informational	N/A	org.apache.commons.httpclient.HttpMethodDirector executeWithRetry I/O exception (java.net.SocketException) caught when processing request: Too many open files
Informational	N/A	org.apache.commons.httpclient.HttpMethodDirector executeWithRetry Retrying request
Informational	N/A	org.apache.commons.httpclient.HttpMethodDirector executeWithRetry I/O exception (java.net.SocketException) caught when processing request: Too many open files
Informational	N/A	org.apache.commons.httpclient.HttpMethodDirector executeWithRetry Retrying request
Warning	N/A	org.apache.commons.httpclient.HttpMethodBase getResponseBody Going to buffer response body of large or unknown size. Using getResponseBodyAsStream in...
Warning	N/A	org.apache.commons.httpclient.HttpMethodBase getResponseBody Going to buffer response body of large or unknown size. Using getResponseBodyAsStream in...
Error	N/A	org.apache.axis.Message writeTo java.io.IOException: java.net.SocketException: Broken pipeat java.net.SocketOutputStream.socketWrite0(Native ...
Error	N/A	org.apache.axis.Message writeTo java.io.IOException: java.net.SocketException: Broken pipeat java.net.SocketOutputStream.socketWrite0(Native ...

Figure 8-10 Monitoring Cognos warning messages

To solve this problem, we made the configuration changes that are shown in Table 8-1.

Table 8-1 Configuration changes

Domain	Type	Item	Value
root	hard	nofile	32000
db2inst1	hard	nofile	21000

a. Set `ulimit` to `65536` for DB2 LUW content store. The setting can be as follows:

```
ulimit -a
```

b. We changed the `nofile` setting in the `/etc/security/limits.conf` file for both db2inst1 and root.

Using Cognos 10 monitoring features

In addition to the Tivoli Monitoring Suite, IBM Cognos 10 provides an Administration Console that displays Cognos internal performance metrics. Key individual metrics are monitored to provide administrators with better insight into the status and overall health of the Cognos BI various components, including the following examples:

- AverageTimeInQueue
- FailedRequestPercent and SuccessfulRequestPercent
- MillisecondsPerSuccessfulRequest
- NumberOfFailedRequests
- NumberOfProcessedRequests
- NumberOfRequests
- NumberOfSessions
- NumberOfSessionsHighWaterMark
- NumberOfSuccessfulRequests
- QueueLengthHighWaterMark
- ResponseTimeHighWaterMark
- ServiceTimeAllRequests
- ServiceTimeFailedRequests
- SuccessfulRequestsPerMinute
- TimeInQueue
- TimeInQueueHighWaterMark

The internal statistics might help the Cognos administrator or application analyst to fine-tune the Cognos-related configuration for better throughput.

Another complementary monitoring feature is to write custom scripts to look at various internal Cognos performance metric data from AIX, Linux on System z, and Windows. For example, some of the middleware components offer additional monitoring APIs that can be used to get a view of the inner workings of that component. For monitoring Cognos TM1, you can also execute the tm1perfmon.exe utility. A variety of metrics are available to help you determine the health of the TM1 server, including internal memory usage, thread creation and failures, and other counters. These can be scripted to provide only the metrics you want, or the `tm1perfmon.exe` program can be run in a loop for interactive monitoring of the TM1 server.

8.2.6 Monitoring DB2 for z/OS performance for Cognos

For our project testing, the data source was deployed on DB2 for z/OS for Cognos BI and SPSS to access. Because the main application tier of Cognos is on the distributed systems, in this case Linux on System z and AIX on Power7 BladeCenter, monitoring the DB2 distributed threads and DDF performance is important. The general best practices for monitoring DB2 distributed workload are similar for monitoring and optimizing Cognos DB2 performance:

- Monitor maximum concurrent threads and maximum distributed threads to meet concurrent number of user requests.
- Allocate sufficient DB2 buffer pool size and dedicate buffer pool for higher data-in-memory hit ratio to minimize the physical I/O delay.
- Optimize the Cognos DB2 physical table space and data set placement on the balanced I/O channel path and faster DASD device type such as DS8800.
- Take advantage of the DB2 10 latest optimizer enhancements for the lowest CPU cost and fastest SQL query response time.
- Store SQL statement in cache to best utilize the memory for faster access.
- Monitor deadlocks or timeouts to minimize locking contention to ensure better throughput.
- Monitor JDBC type-4 connectivity and DB2 Connect gateway performance by using both OMEGAMON XE for DB2 and OMEGAMON XE for Mainframe Networks to ensure the network connections, with no backlog or orphan state associated with Cognos applications.
- For end-to-end Cognos work requests spanning across CPC and zBX, use OMEGAMON XE for Mainframe Networks to monitor IEDN private network and multiple IBM Tivoli Monitoring for DB2 and OS, ITCAM for WebSphere, and OMEGAMON performance monitors to track underlying operating systems (z/OS, Linux on System z, AIX, Windows), JVM (on Linux and AIX), and databases KPIs, all at a glance. This end-to-end capability can help detect a long running SQL query from holding the locks and other DB2 resources while waiting on JVM, for instance, on the distributed system.

An efficient DB2 performance monitor can provide powerful insight into DB2 systems resource utilizations and detect potential application slow-down bottlenecks through real-time sampling technology. Examples include displaying warnings if the number of threads are approaching maximum, whether real storage for data in the cache is sufficient, calculating buffer pool hit ratio, dynamic SQL statement in cache, and monitoring CPU and I/O thresholds.

Figure 8-11 shows the DB2 subsystem key performance indicators view from the overview workspace of OMEGAMON XE for DB2 Performance Expert.

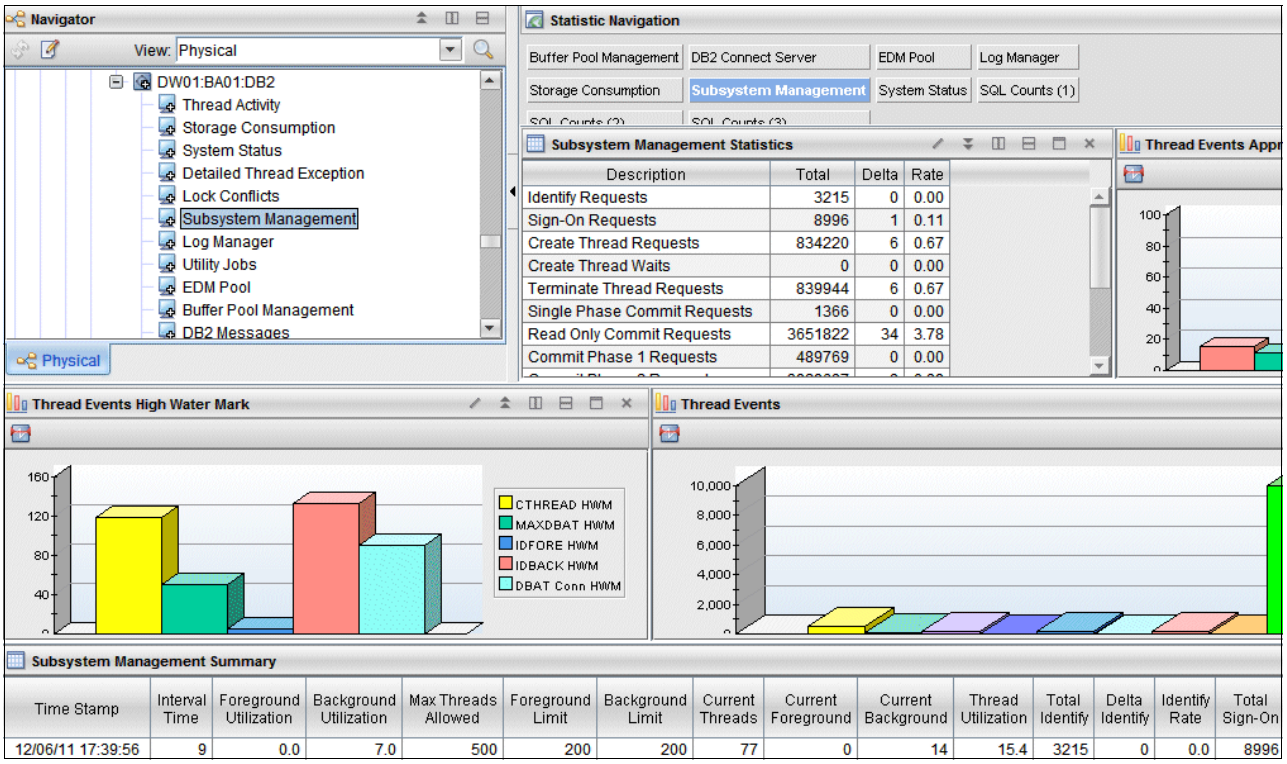

Figure 8-11 DB2 system health overview for Cognos applications

In general, the SQL query performance can be fine-tuned through consistent monitoring and adjusting resource allocations. For large table space scans, DB2 optimizer will dynamically invoke sequential pre-fetch methods for read efficiency, taking advantage of SQL statements in cache for faster reuse, and access big data in real storage. Monitoring DB2 getpages for each read I/O helps detect abnormal high I/O rate. Usually that is a warning sign for either big scans or a shortage of buffers. To monitor DB2 I/O activity, a quick “snapshot” from OMEGAMON DB2 gives an overview of all DB2 volume activity and physical VSAM data set extent numbers. Figure 8-12 shows activity.

Volume Name	Utilization	Service Time	Total I/O	DB2 I/O	Total I/O Rate	DB2 I/O Rate	Number of Data Sets	Number of Extents	Volume Use Percent	DB2 I/O Percent	Extents Per Data Set Ratio
BAW004	0.0	0.0	0	0	0.0	0.0	73	113	0.0	0.0	1.5
BAW010	0.0	0.0	3	3	0.0	0.0	87	182	8.5	100.0	2.0
BAW012	0.0	0.0	3	3	0.0	0.0	77	116	8.5	100.0	1.5
BAW007	0.0	0.0	2	2	0.0	0.0	82	169	5.7	100.0	2.0
BAW005	0.0	0.0	1	1	0.0	0.0	37	81	2.8	100.0	2.1
BAW008	0.0	0.0	6	6	0.0	0.0	46	59	17.1	100.0	1.2
BAW009	0.0	0.0	2	2	0.0	0.0	30	46	5.7	100.0	1.5
BAW003	0.0	0.0	0	0	0.0	0.0	10	15	0.0	0.0	1.5
BAW016	0.0	0.0	11	11	0.0	0.0	63	100	31.4	100.0	1.5
BAW002	0.0	0.0	4	4	0.0	0.0	77	91	11.4	100.0	1.1
BAW013	0.0	0.0	0	0	0.0	0.0	43	72	0.0	0.0	1.6
BAW011	0.0	0.0	0	0	0.0	0.0	11	137	0.0	0.0	12.4

Figure 8-12 DB2 I/O activity on all of DB2 volumes

If the number of extents for a table space gets higher than 100, that might be a warning sign for a reorganization utility to be scheduled to eliminate excessive extents for a faster data fetch. Another usage scenario for monitoring excessive I/O slow-down is if the DB2 I/O number is low and total I/O is high, which indicates other non-DB2 I/O activities occurring on that volume, which might explain why your Cognos request was waiting.

8.2.7 Monitor hybrid workload using Unified Resource Manager

In addition to each component-monitoring and cross-system view from Tivoli Enterprise Portal and Tivoli Monitoring Suite, the zEnterprise architecture provides a built-in Unified Resource Manager to oversee the hybrid workload performance and whether the performance goals are met. The Unified Resource Manager provides basic performance functions and performance-data collection for all workloads for which you have activated service policies. If the goals are not being achieved, detailed performance data from Unified Resource Manager helps identify the source of performance problems. To enable the zEnterprise management functions, the administrator must define the workload and activate a performance service policy for that hybrid workload. The Unified Resource Manager provides the following basic support:

- Monitoring and data collection at the *hypervisor and virtual server* levels, with data provided through these reports: service class report, virtual server report, hypervisor report, resource adjustment reports by workload, virtual server and service class reports.
- Performance management according to the *velocity goals* that you defined for the service classes in the active policy
- CPU management that enables resource allocations to be adjusted dynamically *between virtual servers that have the same hypervisor*.
- Additional customization, which includes installing or starting guest platform management providers (GPMPs) where you can benefit from advanced capabilities and additional performance data:
  - Virtual server topology report, including operating system statistics
  - Hops report for a service class
- Collecting application data from IBM middleware products that have been instrumented with the Open Group Application Response Measurement (ARM) 4.0 standard

  The GPMPs are supported on the following virtual servers:
  - The logical partitions (LPARs) running z/OS V1R12, V1R11, or V1R10.
  - The z/VM guests running z/OS V1R12, V1R11, or V1R10
  - The z/VM guests running Linux
  - The virtual servers running AIX on a POWER blade

In this deployment project, we also tested “report tasks” in the primary HMC to verify workload performance reports for a zEnterprise ensemble. Many types of reports provide performance information from various dimensions. During testing, we examined major Unified Resource Manager reports: workload details and report; and service class, virtual server, resource adjustments, and virtual server topology reports.

Workload Details

Figure 8-13 shows the details of our defined workload. You can select the check boxes of the virtual servers that are included in the workload and the activated performance policy. You also can make changes here.

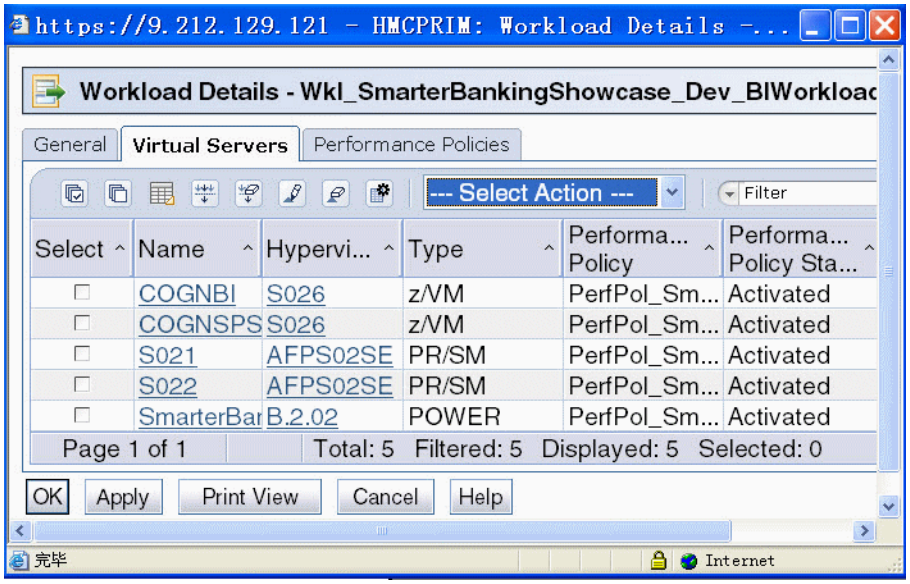

Figure 8-13 Workload Details from Unified Resource Manager

Workloads Report

This report provides high-level performance status and goal achievement information for all the workloads that are defined for the ensemble. The report in Figure 8-14 illustrates our test workload of running 25 concurrent users. We can see that 1.00 is the maximum performance index of our defined BI workload.

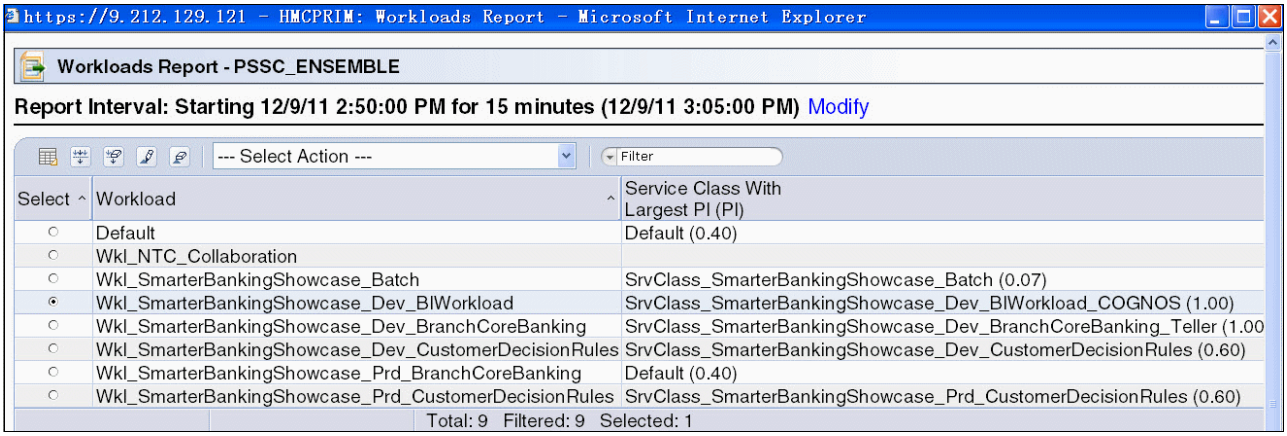

Figure 8-14 Workloads Report

Service Classes Report

This report provides a list of all the service classes that are defined for the workload. For each service class, details include the active performance policy, goal and importance definitions, and actual performance data. If this report indicates performance problems, you can drill down to individual virtual server reports for additional details.

Figure 8-15 shows the service classes report of running a test scenario of 25 concurrent users. It shows the two service classes we defined for Cognos application and SPSS application respectively, and also the performance index of them during a 15-minute interval.

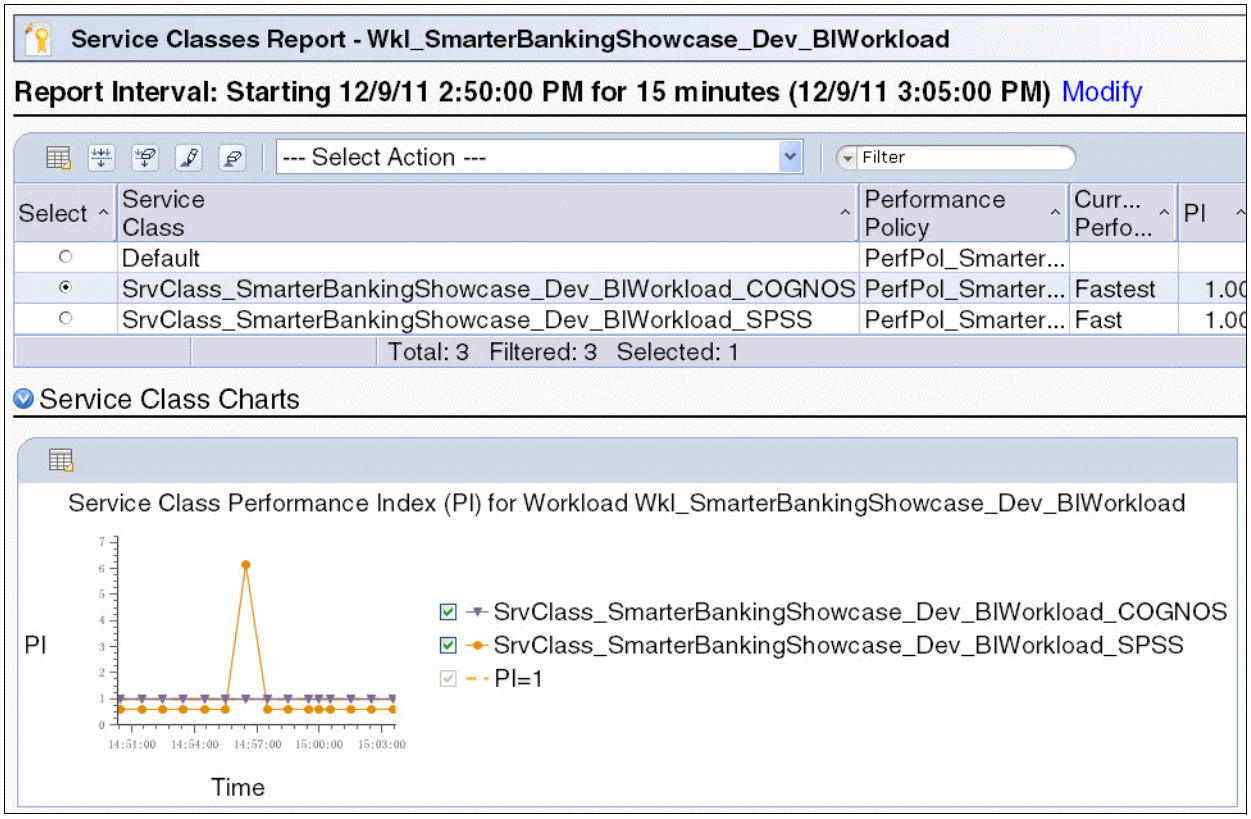

Figure 8-15 Service Classes Report

Virtual Servers Report

This report provides data for virtual servers associated with the workload or a service class. The data includes allocated resources, resource utilization data, and delay statistics.

Figure 8-16 shows the report of running 25 concurrent users. You can see defined virtual processors, current allocated memory, and current CPU utilization of each virtual server.

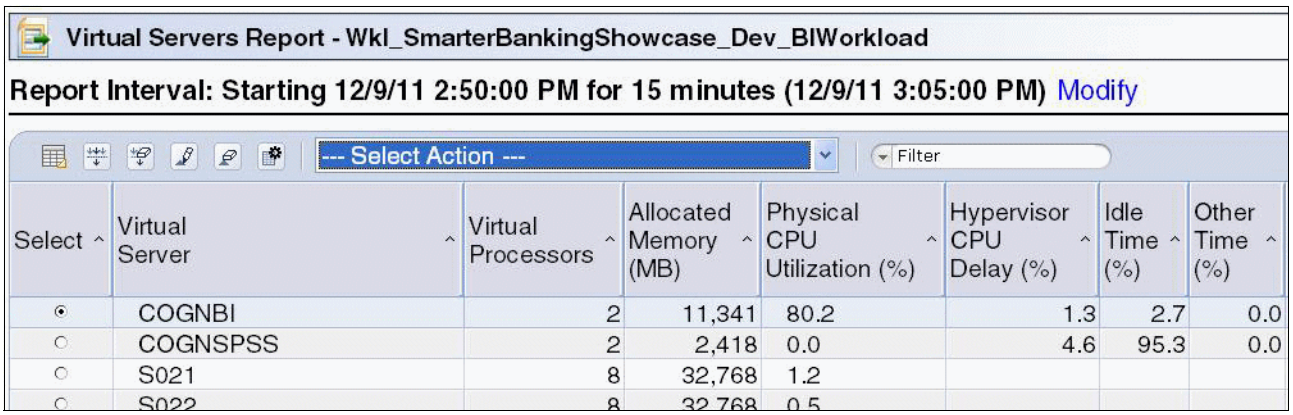

Figure 8-16 Virtual Servers Report

Also you can show a CPU utilization chart during the test period of each virtual server. Figure 8-17 shows the CPU utilization of COGNBI virtual server for the previous test.

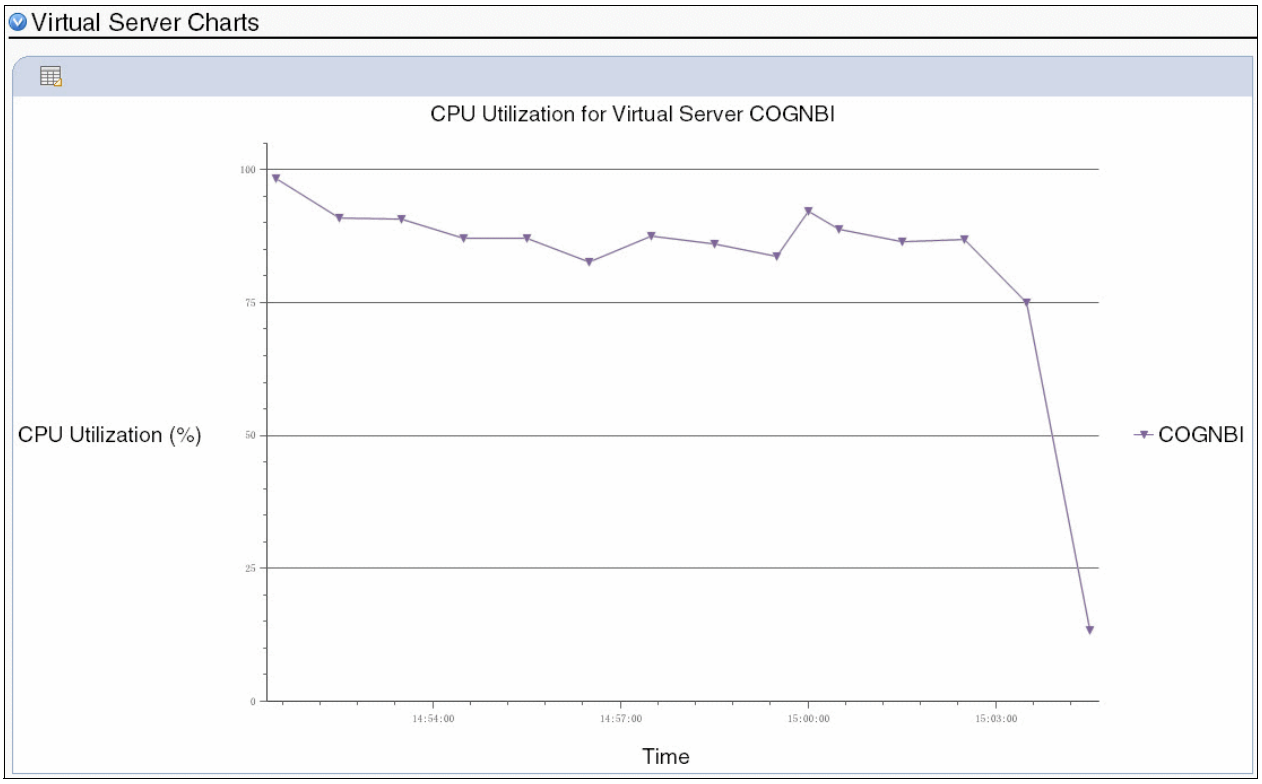

Figure 8-17 CPU Utilization of COGNBI Virtual Server

Resource Adjustments Report

This report lists successful and failed resource adjustments for the selected workload, virtual server, or service class, over a specified interval.

For the BI workload, the AIX virtual server is defined with dedicated processors, so it is not eligible for resource adjustment.

The Linux virtual servers, such as COGNBI, are defined with two processors in shared mode. However, during the testing period, no applications competed for CPU resource, so COGNBI got all CPU cycles that it requested when it was dispatched. From our Service Classes Report, we can see that the performance index of Cognos service class is always one, which means the goal is met, so it does not show any resource adjustment for this virtual server.

Virtual Servers Topology Report

This report provides a graphical view of the relationships between virtual servers that are running the workload and providing the resources to complete the work. This report is available only if a guest platform management provider and ARM-instrumented applications are active on the virtual server.

Within our project time frame, Cognos TM1 and SPSS were not ARM-instrumented software products. Unified Resource Manager could not capture the total workload information and we were not able to display the actual virtual server topology on our HMC.

8.3 Additional monitoring tools for consideration

There are additional tools to monitor and track end-to-end transaction response times with drill-down problem diagnostic capabilities. This section describes tools that were not deployed in this project because of the limited time and scope. However, these tools can help manage total zEnterprise and zBX hybrid workload monitoring with intuitive integrated user interfaces, dynamic end-to-end topology and drill-down problem analysis, and performance trend reporting.

Other advanced management and monitoring functions and tools are recommended for consideration.

8.3.1 End-to-end transaction tracking

A challenge when managing hybrid workload performance is the ability to isolate problems and identify the possible location where the slowing down of work request resides across the multitiered heterogeneous server environment. The Unified Resource Manager provides transaction topology capabilities but requires applications to be ARM-instrumented, as discussed in 8.2.7, “Monitor hybrid workload using Unified Resource Manager” on page 98.

In regard to response time for a given transaction, is there a tool that displays elapsed and CPU times, by segment, across the hops or components? IBM Tivoli Composite Application Manager for Transactions Version 7.3 delivers an “agentless tracking” feature that can show end-to-end transaction flow with network statistics by hops or components, without requiring ARM instrumentation. This function provides a fast and simple way to see an overall view of your hybrid transaction workload performance (Figure 8-18). It displays the network times by components for an end-to-end transaction flow.

All Components

Component	Start Time	Source IP	Latency Time	Average Response Time	Active Connections
IBM HTTP Server	11/09/11 14:30:00	*	0.022	0.337	198
Apache Tomcat	11/09/11 14:30:00	10.*	0.015	0.238	42
RT_STICKER_DL_Abstract_Printer	11/09/11 14:30:00	10.*	0.002	2.390	92
DB/2 Connect	11/09/11 14:30:00	10.*	0.003	0.006	92
WebSphere MQ	11/09/11 14:30:00	*.	0.175	0.054	302
Telnet Server	11/09/11 14:30:00	*.	0.009	2.098	217
CICS Transaction Gateway	11/09/11 14:30:00	10.*.	0.026	16.444	58
FTP Server	11/09/11 14:30:00	10.0.49.1	0.000	0.000	1
IBM Directory Server	11/09/11 14:30:00	10.1.49.152	0.058	0.020	64
Mail Server	11/09/11 14:30:00	10.0.49.65	0.003	0.142	4
Oracle Database Server	11/09/11 14:30:00	10.0.0.38	0.004	0.008	6
SSH Server	11/09/11 13:55:00	10.0.0.83	0.000	0.000	2
WebSphere Application Server	11/09/11 14:30:00	10.0.49.1	0.004	1.274	4

Figure 8-18 ITCAM for transactions tracking by component

In addition to transaction response times by components, the Application Management Console feature of IBM Tivoli Composite Application Manager for Transactions provides a transaction topology view, dynamically with collaborative metrics from data collectors (Internet service monitor, web response time, and robotic scripting tester), and drill-down capabilities to the IBM Tivoli Monitoring or OMEGAMON resource monitors for faster problem diagnostics across zEnterprise virtual servers and networks.

Figure 8-19 illustrates a seamless integration between transaction tracking and in-depth resource bottleneck analysis.

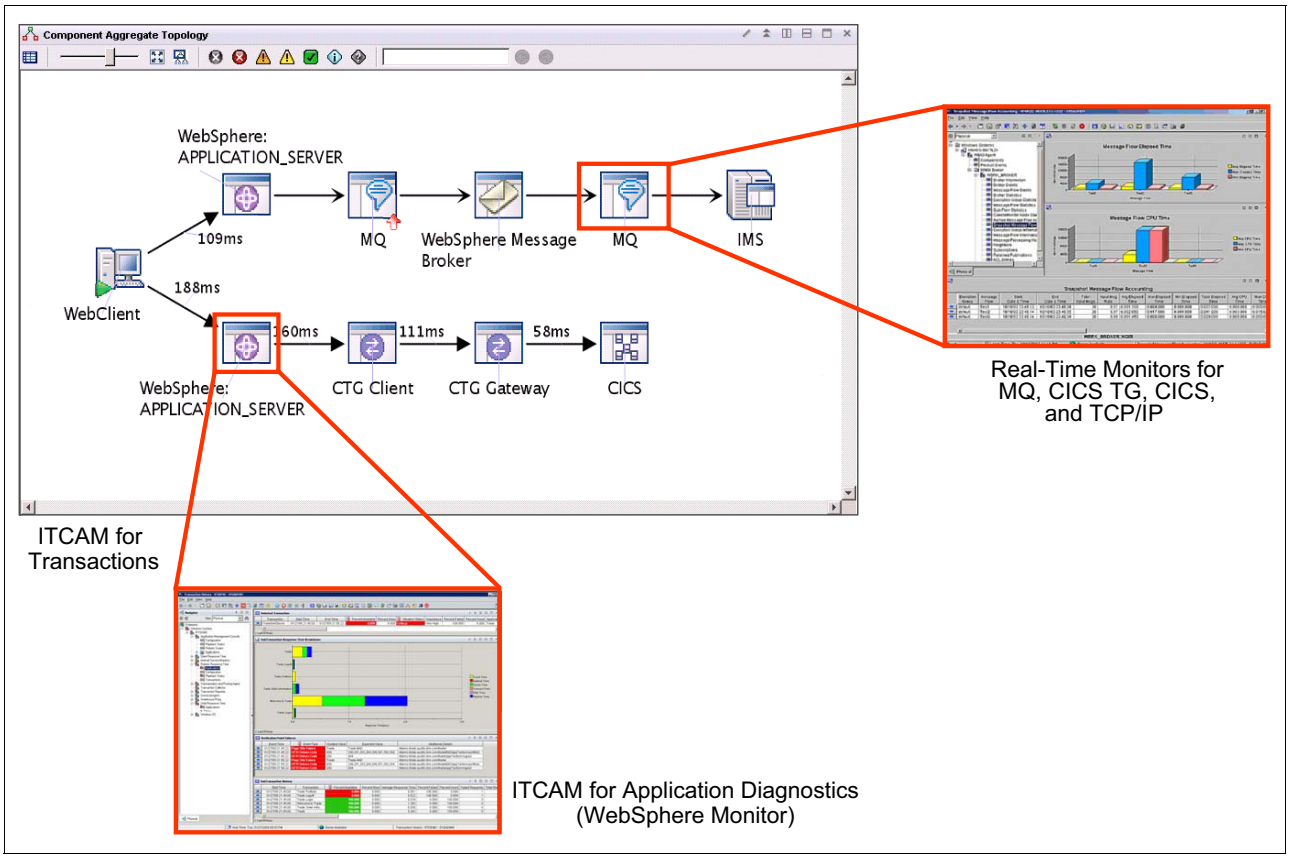

Figure 8-19 End-to-end transaction tracking with drill-down analysis

8.3.2 Integrate BI application monitoring with zEnterprise

In 8.2, “Current hybrid workload monitoring solution” on page 83, we discussed using the current Tivoli Monitoring Suite for OS, WebSphere, and databases to monitor the key processes for Cognos BI web, application and data tiers, and using the Cognos BI Administration console to monitor Cognos internal performance thresholds and abnormal statistics. For longer term capacity-growth planning, you should establish an integrated console for visibility and control with both systems and application internal multitier monitoring to ensure business service goals be achieved.

To take advantage of the Tivoli Monitoring Infrastructure for cross system monitoring and IBM Tivoli Monitoring Agent Builder toolkit, you can customize the Tivoli Enterprise Portal with systems management data feeds supplied by Cognos, and present the monitoring data on Tivoli Enterprise Portal with Cognos services in context. You get in-depth performance statistics at the Cognos application tier levels, such as Cognos gateway, Content Store, and Report Server metric statistics, for real-time problem diagnostics and historical trending analysis. The benefit for the customers is to have an integrated dashboard and drill-down feature for BI problem determination across server components, all from one single management point; Tivoli Enterprise Portal provides proactive management for Cognos cross-system complex tiers as well as underneath physical servers and databases.

Figure 8-20 illustrates a customized dashboard for monitoring Cognos internal components, by using Tivoli Monitoring Agent for Cognos.

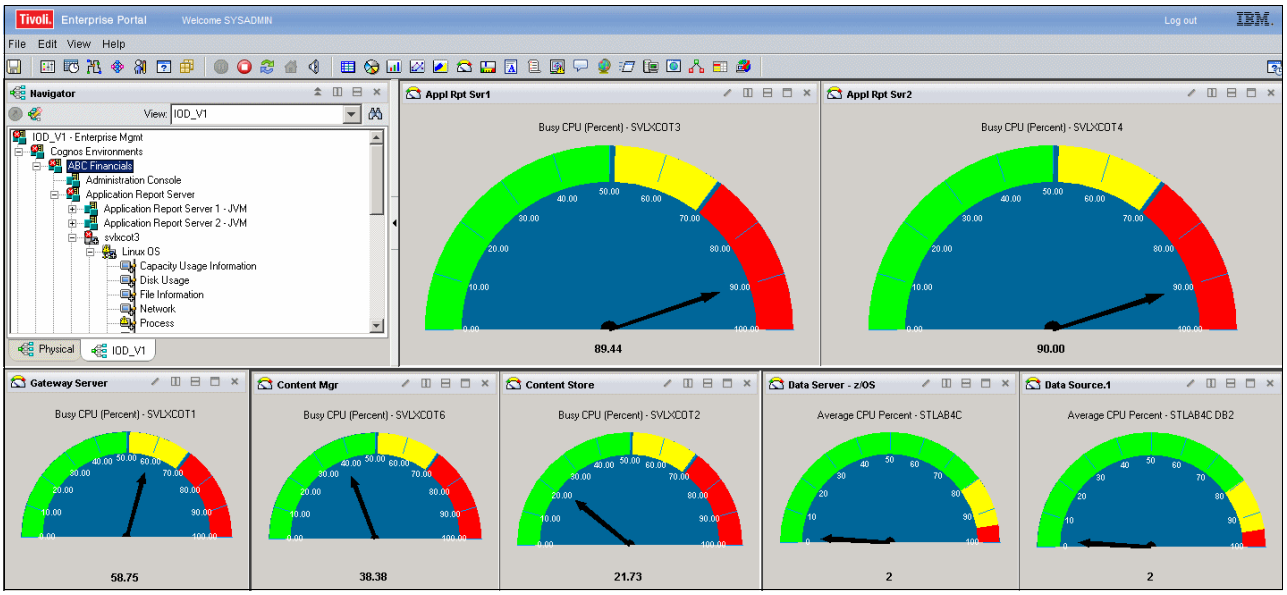

Figure 8-20 Cognos BI application performance dashboard

Using the Tivoli Agent Builder toolkit, you can write scripts to interface with Cognos Systems Management statistic API and display the attribute data on the Tivoli Enterprise Portal. Figure 8-21illustrates that Tivoli Enterprise Portal Navigator on the left panel dynamically reflects Cognos components in context (Administration Console, Application Report Server, Content Managers, Content Store, Data Source, gateway Server). You can define key attributes and thresholds for each tier to be monitored, and Tivoli Enterprise Portal automatically displays alerts and events that are relevant from Cognos web, application, and data tiers.

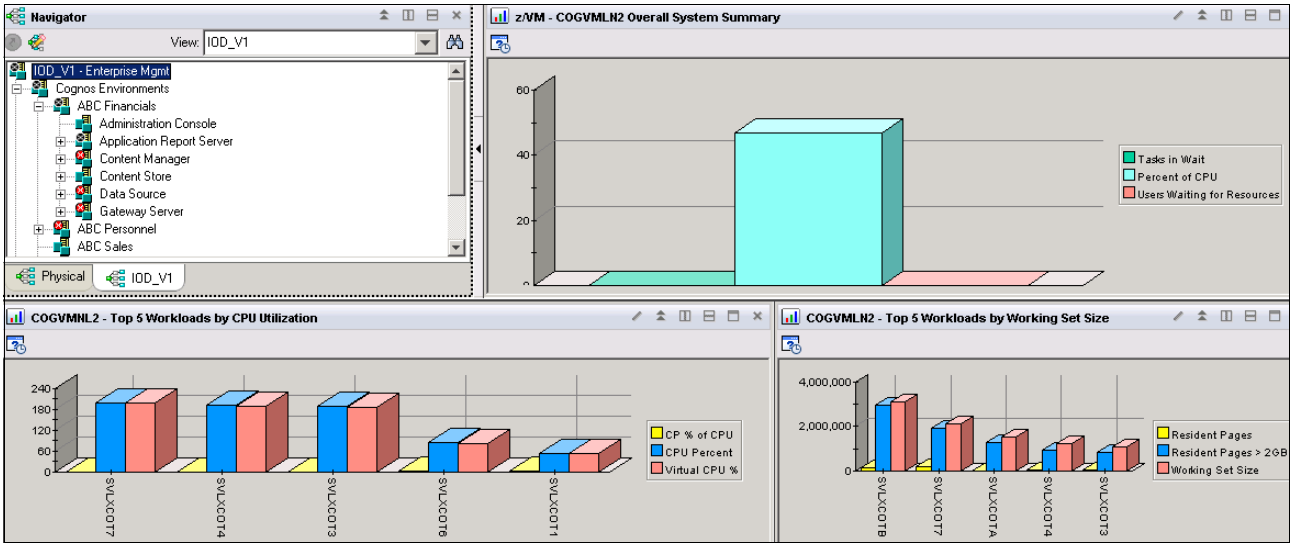

Figure 8-21 Tivoli Monitoring Agent for Cognos BI

You can also install and plug Tivoli UNIX Log Agent into the IBM Tivoli Monitoring infrastructure. This way helps monitor and trap Cognos Application log messages to send early warnings and alerts to the administrator.

8.3.3 Integrate Unified Resource Manager performance API data

IBM Tivoli zEnterprise Monitoring and Discovery Agent, which was developed by using Tivoli Monitoring Infrastructure, exploits the Unified Resource Manager Management API and externalizes the zEnterprise ensemble workload monitoring data on Tivoli Enterprise Portal. This solution helps with the high visibility of total end-to-end overall performance health monitoring and event management of CPC and zBX resource utilization and capacity, with correlated performance statistics from Unified Resource Manager, with its underlying capacity management from IBM PR/SM™, PowerVM and hypervisor.

The Tivoli zEnterprise Monitoring Agent is a monitoring application for the zEnterprise platform, built specifically to discover and monitor the components of the zEnterprise platform, including the physical, virtual, logical layers, workloads, and other related artifacts.

The Tivoli zEnterprise Monitoring Agent collects performance, availability, and configuration information from the hardware level up to the virtual server level. It is based on the IBM Tivoli Monitoring technology and as such, is integrated seamlessly into the monitoring infrastructure that is in place for other monitors such as OMEGAMON XE for z/OS.

The data that is collected by the Tivoli zEnterprise Monitoring Agent helps you to monitor the performance health and availability of the workloads across and within an zEnterprise ensemble. Specifically, this agent monitors the following items:

- The configuration of zEnterprise ensembles and the nodes in those ensembles
- The physical layer of the environment, such resources as central processing centers (CPCs), racks, chassis, and blades
- The virtual layer of the environment, such as virtual servers and virtual networks
- The performance characteristics of the environment, such as workloads, performance policies, service classes, performance indexes, and non-satisfaction indexes

To help you understand the underlying causes of potential problems, the zEnterprise Monitoring Agent provides *workspaces* that are ready for use and that relate zEnterprise workloads to the underlying physical and virtual infrastructure, including networking devices and storage resources. From high-level views, you can click down to more detailed views that answer specific questions about both physical and virtual resources, such as on what platform a virtual server is running and what other programs are running on the same hypervisor. Detailed information about workloads helps you understand service classes, the virtual servers that are associated with a particular service class, and whether these servers are meeting their service class goals. All of this easy-to-access and easy-to-understand information helps you deliver the 24x7 availability that your zEnterprise users expect.

Figure 8-22 shows the new IBM zEnterprise Monitoring and Discovery Agent through the common IBM Tivoli Monitoring infrastructure on Tivoli Enterprise Portal.

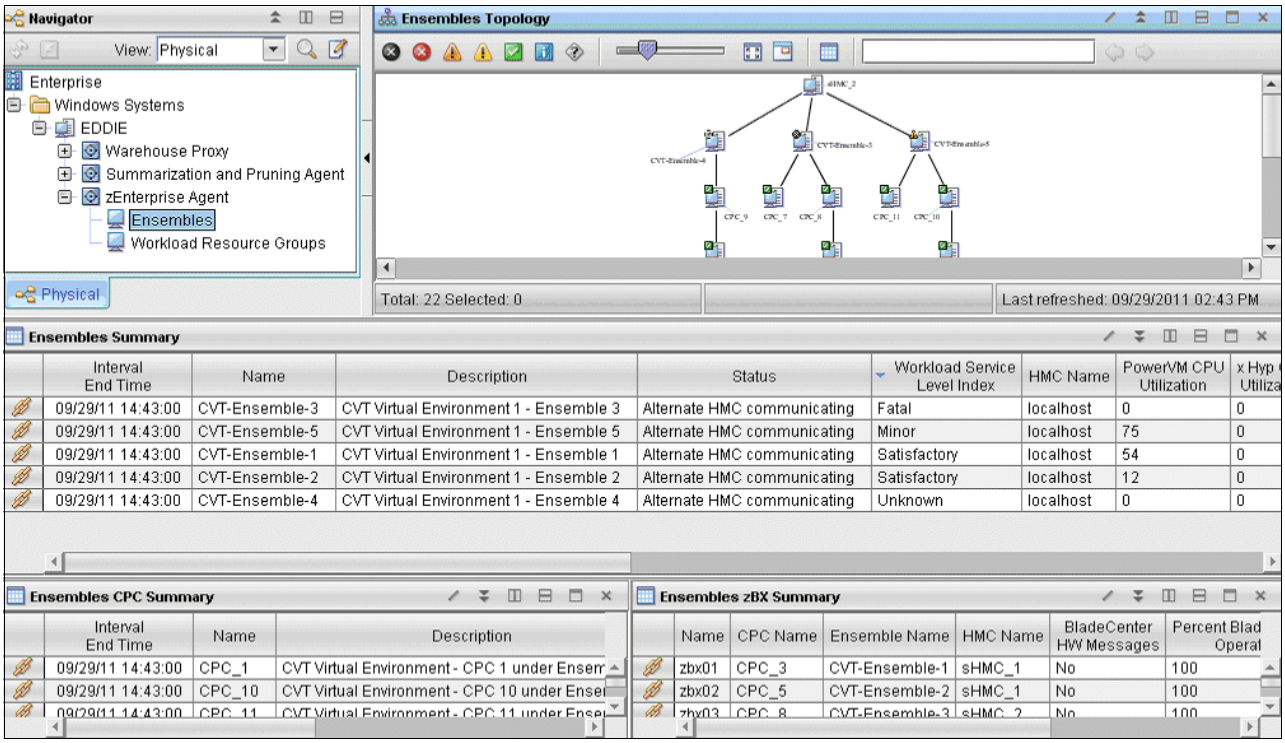

Figure 8-22 Tivoli zEnterprise Monitoring and Discovery Agent

Note this predefined sample workspace is ready for use and displays zEnterprise ensemble topology, its per-platform CPU and memory usage information, and summary information for the central processor complexes (CPCs) and the zEnterprise BladeCenter Extensions (zBXs) that are associated with the ensemble.

The key enabler for this monitoring agent is a component named the Enterprise Common Collector, introduced in the Tivoli Monitoring Infrastructure. It provides a central collection point for zEnterprise data and provides added-value capabilities to complement Unified Resource Manager:

- Manages individual connections to zEnterprise ensembles.
- Tests connections zEnterprise components to verify the correctness of definitions.
- Collects data in parallel from multiple zEnterprise ensembles.
- Enables users to customize the scope and depth of data collection.
- Processes monitoring requests and background data collection activities concurrently.
- Notifies users of inventory, status and property changes and reflects those changes in reports requested through the collector API.

With Tivoli zEnterprise Monitoring Agent, the Enterprise Common Collector extends the existing Tivoli Monitoring Suite with new Unified Resource Manager management statistics, providing total visibility and control for managing the entire zEnterprise ensemble, from the top-level overview down to the source of the problems under the control of a single user interface.

The complementary value proposition is added to the Unified Resource Manager by design to offer the following benefits:

- Enables greater visibility of zEnterprise ensemble's system health on the centralized cross-platform event-automation console in the command center.
- Reduces the processing load on the HMC by providing an alternative source of performance and inventory data.
- Relieves client programs of the burden of message handling that is normally required to keep data current.
- Provides multiple concurrent users with alternative access to the HMC type of information and decreases the amount of traffic to and from the HMC console.
- Adds value to HMC performance information by deriving deltas and generating new metrics based on the sampled data.
- Integrates with existing workload monitors to ensure total performance goals are being met and detects exceptions for Linux on System z in all levels. For example, if the hybrid workload performance index exceeds one, and the service goal is not achieved, the administrator can drill down from the Tivoli Enterprise Portal overview workspace, side-by-side. In this way, the administrator can choose and select virtual server z/OS to look at the WLM service analysis, or look at z/VM monitoring workspace, or look at AIX or Windows for further poor performance indicators to determine the source that caused the bottleneck.

9

Implementation summary

This chapter summarizes the results of the activities that are associated with the configuration of the infrastructure and middleware, deployment of the workload, and the testing activities that are involved in this project.

9.1 Infrastructure and middleware implementation

The assessment phase (described in *Using zEnterprise For Smart Analytics: Volume 1 Assessment*, SG24-8007) identified the optimal product configuration across a range of operating systems within zEnterprise. However, we discovered some unsupported functionality that required a deviation in terms of overall architecture.

Unsupported functionality: This issue is common within IT departments. Although the architect defines the ideal configuration, the specialists must work with the products that are available to them in the currently supported versions.

We contacted the development labs, to request that these unsupported features be evaluated for possible inclusion in future releases of the products.

In the ideal world, the following situations would be in available:

- Cognos TM1 would communicate directly with SPSS Modeler to execute predictive models providing result sets or database updates that could be used dynamically within TM1.
- Cognos Metric Studio would support DB2 for z/OS as the common store.
- Cognos Metric Studio would support WebSphere Application Server 64-bit.

Within IBM Montpellier, there are two principal network backbones, one for IBM internal usage and one for external usage. Because access to both networks was required during this project, a number of additional VLAN and security configurations were required. This requirement caused some additional effort.

With regards to the monitoring of this workload, although a number of reports and services can be used to display specific issues, establishing an enterprise-wide monitoring solution that encompasses the infrastructure, middleware, and user activity for this analytics workload is highly preferred.

The IBM Tivoli suite of monitoring products, including IBM Tivoli Monitoring, OMEGAMON, and ITCAM components are currently deployed with the Smarter Banking Showcase. These products were easily tailored to provide dedicated screens that capture and report on the availability and performance of the analytical services.

9.2 Customer Profitability Analytics blueprint deployment

This package was developed to portray the capabilities of a number of IBM analytical products. The sample data and configuration was built as a functional demonstration running in a Windows environment. However, many customers questioned how these products and components would perform within a high-volume, multi-platform, production-like environment.

After the infrastructure and middleware components were configured, the package was deployed by using the provided installation scripts. Most of the products are platform independent and therefore proved straightforward to install, although several prepackaged elements, such as TM1 objects, were created under Windows, which is case-insensitive. Deploying the same objects under AIX resulted in a mismatch between objects that were stored within libraries and classes. We contacted the labs for advice and were informed that a conversion toolkit was planned. In the meantime, we redeployed the TM1 package to execute within the Windows Server 2008 environment running on a zBX x86 blade.

9.3 Testing

IBM Rational Performance Tester was used to capture and replay a number of the user reporting requests. The display of dashboards, scorecards, and reports involved simulated users, accessing the Cognos BI gateway and executing a report request or selecting one or more of the display options.

The test executions included elements of wait times based on typical web usage and random report selections per execution group.

Within our development environment, we were constrained by network bandwidth and simultaneous development activities, so our simulation testing extended only to 25 users before network time-outs prevented workload from completing.

The comparison of results showed the following information:

- Linear increments of CPU and memory usage based on limited user volumes
- Extended response times when processing a greater proportion of dashboard requests, which use a significant number of HTML graphic objects

  Although this information can be addressed by performance tuning of the Cognos dispatcher service classes, this was outside of the scope and timescale of the project.
- No specific network or data access bottleneck

We plan to refine the performance requirements of this solution and test again with greater volumes. This solution will then be deployed within a production environment and made available for IBM demonstrations.

9.4 Conclusion

System z is traditionally the cornerstone for high-volume, batch-oriented transaction processing, and analytics often involves variable processing demands and the use of graphical business user tools to support a variety of reporting needs. This combination provided a number of conflicts when we considered only the z/OS or z/Linux environments. However, with the introduction of the zBladeCenter Extension that supports blades and operating systems such as AIX, Linux and Windows, analytical processing can exploit a rich set of capabilities within a closely-coupled runtime environment.

The key benefits of deploying the analytics products and a business package onto this environment are as follows:

- A common data store for all products and access to warehouse tables within a scalable database environment, DB2 10 for z/OS
- High performance network connectivity services:
  - HiperSockets between z/OS and Linux on System z
  - IEDN between the zBX blades and the z196 environment
- A range of high availability and scalability options to support growing demands for data, users and analytical content
- The benefit of using enterprise-wide tools to monitor the entire environment, warning about both hardware and software issues, and supporting automated responses to rapidly resolve these issues before they impact the business
- The use of advanced workload management capabilities through the Unified Resource Manager to establish performance policies and goals to meet the demands of running workloads across multiple components

Although this is a complicated project involving a number of components and personnel skills, the benefits of establishing a common framework for the analytical components within this environment can lead to many future deployments of packages to support the growing business requirements.

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

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- *Building an Ensemble Using IBM zEnterprise Unified Resource Manager*, SG24-7921
- *HiperSockets Implementation Guide*, SG24-6816
- *IBM SPSS predictive analytics: Optimizing decisions at the point of impact*, REDP-4710
- *IBM zEnterprise 196 Technical Guide*, SG24-7833
- *Using Rational Performance Tester Version 7*, SG24-7391
- *Using zEnterprise For Smart Analytics: Volume 1 Assessment*, SG24-8007
- *z/VM and Linux on IBM System z: The Virtualization Cookbook for SLES 10 SP2*, SG24-7493

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Other publications

These publications are also relevant as further information sources:

- *IBM Cognos Business Intelligence Installation and Configuration Guide*, GC19-2964
- *OMEGAMON XE for Mainframe Networks Planning Configurations*, SC27-2384
- *zEnterprise System Ensemble Performance Management Guide,* GC27-2607

Online resources

These websites are also relevant as further information sources:

- Cognos Business Intelligence Installation and Configuration Guide

  http://pic.dhe.ibm.com/infocenter/cbi/v10r1m0/index.jsp?topic=%2Fcom.ibm.swg.im.cognos.inst_cr_winux.10.1.0.doc%2Finst_cr_winux.html
- IBM DB2 Analytics Accelerator for z/OS

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

- IBM Cognos Metrics Manager Getting Started Installation Guide

  http://publib.boulder.ibm.com/infocenter/c8bi/v8r4m0/index.jsp?topic=/com.ibm.swg.im.cognos.qrc_mm_inst.8.4.1.doc/qrc_mm_inst.html

- *IBM SPSS Collaboration and Deployment Services 5 Coordinator of Processes Service Developer's Guide*

  ftp://public.dhe.ibm.com/software/analytics/spss/documentation/cds/5.0/en/ws/SPSSCopDevGuide.pdf

- Metric store databases

  https://www.ibm.com/support/docview.wss?uid=swg27019126#metric_stores

- Create metric store database

  http://ibm.co/RPEAT6

- SPSS

  http://www.ibm.com/software/analytics/spss/

- zBX, IBM zEnterprise System

  http://www.ibm.com/systems/z/hardware/zenterprise/

- Rational Performance Tester

  http://www.ibm.com/software/awdtools/tester/performance/

- IBM Cognos Metric Studio User Guide

  http://pic.dhe.ibm.com/infocenter/cbi/v10r1m0/index.jsp?topic=%2Fcom.ibm.swg.im.cognos.ug_mm.10.1.0.doc%2Fug_mm_id3443sc_metrics.html

- Cognos Metrics Manager Installation Guide

  http://pic.dhe.ibm.com/infocenter/cfpm/v10r1m0/index.jsp?topic=%2Fcom.ibm.swg.im.cognos.inst_cr_winux.10.1.0.doc%2Finst_cr_winux_id6570InstallMetricsManagerCM_single.html

- TM1 Installation

  http://publib.boulder.ibm.com/infocenter/ctm1/v9r5m0/topic/com.ibm.swg.im.cognos.tm1_install.9.5.2.doc/tm1_install.html

- Cognos BI Administration and Security Guide (see Data Management chapter)

  http://publib.boulder.ibm.com/infocenter/c8bi/v8r4m0/index.jsp?topic=/com.ibm.swg.im.cognos.ug_cra.8.4.0.doc/ug_cra_id25Introduction.html

- *IBM SPSS Modeler Client Installation (Network License) Guide* (concurrent)

  ftp://ftp.software.ibm.com/software/analytics/spss/documentation/modeler/14.2/en/ClientInstallation_ConcurrentNetworkLicense.pdf

- *Network License Administrators Guide*

  http://public.dhe.ibm.com/software/analytics/spss/documentation/statistics/20.0/en/client/InstallationDocuments/Linux/Network_License_Administrators_Guide.pdf

- *IBM SPSS Modeler Server for UNIX Installation Instructions*

  ftp://ftp.software.ibm.com/software/analytics/spss/documentation/modeler/14.2/en/ServerInstallationUNIX.pdf

- *IBM SPSS Modeler Batch Installation for UNIX*

  ftp://ftp.software.ibm.com/software/analytics/spss/documentation/modeler/14.2/en/BatchInstallationForUNIX.pdf

- *IBM SPSS Modeler Batch User's Guide*

  ftp://ftp.software.ibm.com/software/analytics/spss/documentation/modeler/14.2/en/BatchUsersGuide.pdf

- *IBM SPSS Modeler Server Administration and Performance Guide*

  ftp://ftp.software.ibm.com/software/analytics/spss/documentation/modeler/14.2/en/ServerandPerformanceGuide.pdf

- Supported ODBC drivers

  http://www.ibm.com/support

- *IBM SPSS Data Access Pack Installation Instructions*

  ftp://public.dhe.ibm.com/software/analytics/spss/documentation/statistics/20.0/en/client/InstallationDocuments/Linux/IBM_SPSS_Data_Access_Pack_Installation_Instructions.pdf

- Progress DataDirect Connect Series for ODBC

  http://www.datadirect.com

- TM1 TurboIntegrator functions

  http://ibm.co/ovBY5x

- *IBM SPSS Modeler Administration Console User Guide*

  ftp://ftp.software.ibm.com/software/analytics/spss/documentation/modeler/14.2/en/ModelerAdministrationConsole.pdf

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

IBM Redbooks

Using zEnterprise for Smarter Analytics: Volume 2 Implementation

(0.2”spine)
0.17”<->0.473”
90<->249 pages

Using zEnterprise for Smarter Analytics
Volume 2 Implementation

Explore the use of COGNOS, SPSS, and DB2

Learn how to set up analytics tools on System z

Determine how to test the workload

This IBM Redbooks publication series explains the assessment and implementation of a workload, integrated within IBM Smarter Banking Showcase, and hosted at IBM Montpellier, France. Intended for decision-makers, consultants, architects, administrators, and specialists, this book is the second volume in a series of two:

- Assessment: Volume 1 (SG24-8007) describes how to evaluate the requirements of a new Smarter Analytics workload, addressing the user, system resources, and data processing profiles to identify the most optimal configuration by using IBM methodologies, such as fit-for-purpose. Given that the existing showcase is based on the IBM zEnterprise System, deployment options include IBM z/OS, Linux on IBM System z, IBM AIX running on IBM POWER processor-based blades within the zEnterprise BladeCenter Extension (zBX), and Windows Server 2008 running on System x and BladeCenter blades also within zBX.
- Implementation: Volume 2 (SG24-8008), which you are reading, describes the setups that are involved in deploying the Smarter Analytics workload within the showcase. With multiple components, including IBM Cognos BI, IBM Cognos TM1, Cognos Metric Studio, IBM DB2 for z/OS, and a number of application design tools, the workload spans multiple operating environments. The use of application clustering, setting up performance policies by using Unified Resource Manager, and simulation test execution results are included.

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

SG24-8008-00 ISBN 0738437522