October 2014

IBM Cognos Business Intelligence 10.2.1 for Linux on System z - Performance and z/VM Resource Management

Linux end-to-end Performance team:
Thomas Weber, Dr. Juergen Doelle

IBM System z Cognos BI Performance team:
Thuy V. Nguyen, Tim Lighter, Michael Schapira

Table of Contents

About this publication

This white paper presents information about the setup, performance and z/VM® resource management of a Cognos® Business Intelligence (BI) server for Linux® on System z®, running under z/VM (distributed installation). The tuning of some important Cognos BI server components and a performance study using the z/VM Resource Manager (VMRM) to manage constrained CPU resources is also described in this paper.

Remarks

IBM Knowledge Center web references pointing to specific product documentation chapters are usually very long. Therefore, in this paper, a navigator indicates the path to the specific topic for the product documentation.

For example:
IBM Knowledge Center z/VM welcome page:

z/VM 6.3.0 > System overview > z/VM: General Information > What is new or changed in z/VM V6.3

Note: The web-links referred in this paper are up-to-date as of July, 2014.

Introduction

Read this introduction for general information about business intelligence and business analytics, for a discussion of the objectives of the project described in this paper, and for an introduction of the z/VM Resource Manager used in the performance tests. A table describing the notation conventions is also included in this introduction.

Business intelligence (BI) and business analytics

Business intelligence (BI) and business analytics have become widely accepted solutions for more than just business analysts and chief executives. Large enterprises need to be able to access BI and analytics on a comprehensive and expanding scale. To support customers in taking advantage of business intelligence, this white paper includes information about the setup, performance, and z/VM resource management of a distributed Cognos Business Intelligence (BI) server for Linux on System z, running under z/VM.

Some organizations have various BI capabilities in place but want to add more. Others want to use their capabilities more effectively and efficiently. Imagine you can choose an analytics solution that works for your organization today, but can expand to meet your needs in the future. You could order several options, such as personal discovery tools, and expand to a fuller option later. No matter what you order, the later addition of other services would be seamless and easy.

From business intelligence to financial performance and from strategy management to analytics applications, Cognos software can provide whatever your organization needs to become top-performing and analytics-driven.

Cognos Enterprise combines methods of reporting, analysis, modeling, planning, and collaboration for smarter decision-making and better business outcomes. It equips users with what they need to explore information, analyze key facts, quickly collaborate to gain alignment with key stakeholders and act with confidence to drive better business outcomes.

These premium features are available on Linux on System z under z/VM for a reduced total cost of ownership and an easier centralized environment, which results in an increase of the return on investment.

Objectives

From the maintenance and administration view, BI-workloads have a certain drawback. They need a high amount of system resources in terms of memory and CPU to analyze a large amount of data with short response times. Probably these resources are only needed for a limited amount of time each day. The remaining time, these resources are possibly idling. When running in a virtualized environment, they could be used for other workloads running on the same hypervisor. However, this bears the risk, that the BI-workloads, when running in parallel, do not get sufficient resources and the transaction response times increase. The ideal solution would be a management tool which guarantees that a selected workload gets a certain amount of resources regardless of the load level of the system.

z/VM Resource Manager (VMRM)

With VMRM, z/VM offers a feature that provides a workload management feature for the resource CPU. This study analyzes how VMRM can manage the CPU resource for two different workload types that run in parallel:

- ***Cognos BI:***
  A CPU intensive workload with relative long-running transactions was used. The Cognos BI server components were implemented with several z/VM guests (distributed installation).
- ***Online stock trading system:***
  The DayTrader benchmark application was used, a transactional workload with response times in the millisecond range. DayTrader is a WebSphere Application Server (WAS) application, implemented from three components (IHS, WAS, DB2®). Three instances of this workload were used in parallel.

z/VM was equipped with sufficient memory to ensure that there is no constraint for this resource.

To manage the available CPU resources for the workloads, two strategies were considered:

- ***One-workload-gets-all:*** *VMRM was set up in a way that one workload gets all of the CPU resources that it requests. This ensures the best possible response times, but bears the risk that the other workloads suffer, if the remaining CPU resources run short.*
- ***A-chance-for-others:*** *The idea is that always some CPU resources are available for the less important workloads. They may perform slower, but time-outs should be avoided. This could be implemented by applying more CPU resources than the preferred workload uses at its maximum load. Or by choosing a balanced VMRM resource management, where the preferred workload gets the largest portion of the resources, but not all.*

The system under test was implemented in a way that various load levels of the two different workloads were adjusted, and a Cognos BI workload level below the CPU resource-full condition was chosen. The DayTrader workloads then fill the gap until all CPU resources are almost completely used or even overused. Both strategies and their corresponding VMRM CPU resource goals were tested for various workload levels and combinations.

Notational conventions

Table 1 shows the notation conventions used in this paper, in accordance to IEC 60027-2 Amendment 2.

Table 1. Notation conventions

Symbol	Full name	Derivation
KiB	kibibyte	2 ** 10 byte == 1024 byte
MiB	mebibyte	2 ** 20 byte == 1048576 byte
GiB	gibibyte	2 ** 30 byte == 1073741824 byte
KiB/s	kibibyte per second	2 ** 10 byte / second
MiB/s	mebibyte per second	2 ** 20 byte / second
GiB/s	gibibyte per second	2 ** 30 byte / second

Summary

This paper describes how IBM Cognos BI was deployed and tuned in a virtualized environment on Linux on System z, consisting of an IBM zEnterprise® 196 (z196) model 2817-M66 using 10 CPUs and 70 GiB central storage, connected to an IBM System Storage® DS8800 Model 951.

This study analyzes how a workload mix in a virtualized environment can be managed to ensure a certain response time for a particular workload - in this study, a Cognos BI workload - independent of the load level of other competing workloads in a CPU constrained environment. It is obvious that this can only be reached at the expense of the less important workloads. The level of this expense is an important parameter which needs to be considered when defining any resource management rules.

The most important parameter to influence the assigned CPU resources to a z/VM guest is the relative CPU share. Two methods of resource management are considered.

1. **Defining the relative shares manually (static method)**
   This should provide a base line and show what happens when competing workloads run with no prioritization. For the setup, each guest gets all the virtual CPUs the workload would require. Then, each virtual CPU contributes with a value of 100 to the relative share of the guest. This method is called the fair-share setup in the following context.

   z/VM guests with the largest amount of virtual CPUs will get the highest relative share in sum.

   Obviously the big winner was a larger DayTrader application guest with six virtual CPUs and a high CPU load level. The Cognos BI server showed significant throughput degradation and a moderate degradation in response times, because the individual guests end up with lower relative shares in sum for this method. z/VM guests with a larger amount of virtual CPUs will receive the highest relative share.

2. **Using VMRM to manage CPU resources (dynamic method)**

   - First the relevant VMRM management parameters were determined. As z/VM monitoring is the base for VMRM resource management, the CP monitor sample interval determines the frequency; how often VMRM receives updated CP monitor sample records in order to adjust management parameters accordingly. The default interval of 60 seconds can be too large, if a faster adjustment to changed load levels and hence a faster adjustment of the management parameters is desired. The recommendation is to reduce the CP monitor sample interval to 30 seconds or 15 seconds in that case. Six seconds are found to be too short, and the overhead increased.
   - The most important VMRM parameter for this method is the CPU goal parameter.
     - With the most extreme CPU goal settings of 100 for the Cognos BI workload and 1 for the competing DayTrader workloads (One-workload-gets-all), it is possible to keep the load level and response times of the Cognos BI workload almost stable, independently of a constant parallel DayTrader workload with different load levels or an injected workload peak. However this configuration bears the risk that the competing workloads receive a severe reduction of assigned CPU resources. Even application timeouts can occur in the worst case when the system does not have sufficient CPU resources.
     - A more moderate configuration was a CPU goal setting of 70 for the Cognos BI workload, and 25 for the competing DayTrader workloads. This ensures a certain capacity for these competing workloads, once the system CPU capacity is reached. This takes place at the expense of the Cognos BI workload, but as long as the Cognos BI response times are still in an acceptable range that might be a tolerable situation. Besides that, this configuration is consistent against a further increase of the load level for the DayTrader workloads.

- In summary, the recommended approach is to adjust the VMRM goals for the workload groups in an iterative way, for example, starting with 70/25 and increasing it to 80/20 until the desired service times for the preferred workload are reached.
- VMRM starts regulating the relative shares with the initial value defined in a z/VM user directory entry for the guest. In order to reduce the VMRM regulating intervals for the relative share target according to the CPU velocity goals, it might be useful to start with already lower shares in the user directory for less prioritized workloads to shorten the time for reaching the target shares. After that, VMRM does not change the relative shares, when the system becomes unconstrained again.

In an environment with groups of guests that have cooperating components, or run clustered, there are additional advantages of using VMRM instead of using a static relative share setup:

- Additional guests are just added or removed to the workload group without the need to manually review the relative shares. For example, adding a new low prioritized guest in a non-VMRM environment managed with relative shares only, the relative shares of all guests have to be recalculated, otherwise the CPU portion for all low prioritized guests increase.
- When only some of the high prioritized guests are active, but a high amount of the low prioritized guests, they might get a larger portion of the CPU capacity than intended, because the sum of their relative share is higher.

When using the z/VM Resource Manager, there are some important things to consider:

- The CPU goal definitions are in regard to wait times and not to CPU usage. That means, a setting with a CPU goal of 80 does not indicate that the workload gets 80% of CPU capacity. It might be less.
- VMRM is compatible with the ILMT (IBM License Metric Tool).
- VMRM is NOT compatible with CPU pooling.
- VMRM is only effective when the system is CPU constrained. As long as free CPU capacity is available, all guests can get the CPU resources they require.
- This study is restricted to the VMRM capability to manage the CPU resources. Other VMRM management capabilities are not considered.

VMRM can be a very useful tool to manage CPU resources in an environment with two or more competing workloads, where a single workload has a higher priority than the others. It can ensure constant throughput levels and response times independent of the load level of any competing workloads. It can also simplify the z/VM setup. Especially, the capability to manage environment changes, either in guest setup or with workloads level, makes it very attractive.

System under test - overview

Read this introduction and overview of an IBM Cognos Business Intelligence (BI) server solution for Linux on System z, and how it was implemented as a system under test (SUT) for this project.

The SUT includes basic IBM Cognos BI core components like a Content Manager, application tier components like Report Servers, and a Gateway.

The architecture model consists of logical tiers distributed across three independent systems:

- Client components on local workstations (tier one)
- Cognos BI server components on remote servers (tier two)
- Back-end database server on a remote server (tier three)

Tier one is responsible for the presentation and user interaction with tier two processes. The client components enable the user to interact in a secure manner, by using Cognos BI APIs to access the tier two servers. Tier one client processes never access any tier three services on the database servers directly. Tier one could either be a web browser, a Cognos Mobile iPad application, or a custom client using the Cognos SDK to generate transactions against Cognos BI.

Tier two processes are also known as application logic layer. These processes manage the business logic of the application and provide access to tier three services. It is the layer where most of the data processing occurs. The Cognos BI Gateway and Report Servers represent the application logic layer for this setup. The Report Servers run under the control of a WebSphere Application Server (WAS), while the Gateway extends an IBM HTTP server.

Tier three services are protected from direct access by any tier one components. Usually they are residing in a secure network (inside the System z box for example). Interaction is only possible through the tier two services. The Cognos meta data repository, called the Cognos BI Content Store, and reporting databases (DB2 LUW or z/OS®, SQL server) which store the content to be analyzed, are considered as tier three services in this case.

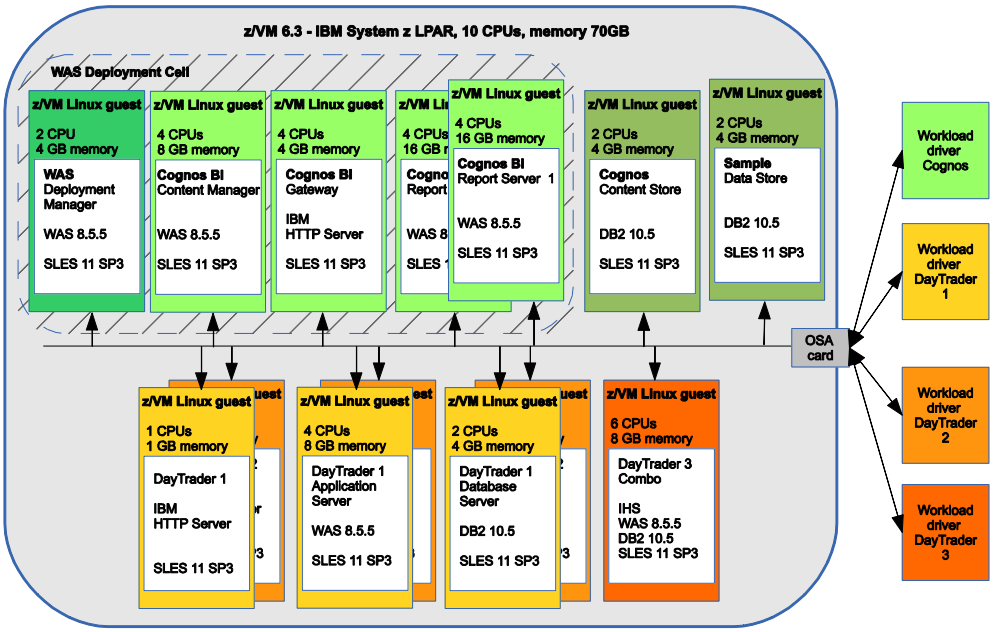

Figure 1: Overview of the system under test (SUT)

The system is implemented in a single IBM System z LPAR in a virtualized z/VM environment. The IBM Cognos BI core components are distributed over four virtual machines and are combined in a WAS Network Deployment cell. The application servers and the HTTP server in the cell can be managed with the WAS Deployment Manager.

The Content Manager database and the BI sample data database run in two virtual machines with IBM DB2 software.

The competing workloads are set up in two virtual machine triplets and one single DayTrader combo system. The virtual machine triplets form a complete DayTrader software stack including an IBM HTTP server, a WAS server running the DayTrader application, and a DB2 database server. The DayTrader combo virtual machine includes all application components in a single instance.

In total, this setup involves:

- 22 virtual CPUs on behalf of the Cognos application
- 20 virtual CPUs on behalf of the DayTrader application

which results in 42 virtual CPUs on 10 real CPUs.

For guest-to-guest communication, the z/VM guests are connected over a Virtual Switch (VSWITCH), which also provides LAN access via a 10Gbit OSA-E card.

An IBM Flex System® blade is used as a Cognos BI and DayTrader workload generator. The IBM Flex System blade is connected over the 10GbE network to the IBM System z mainframe.

Hardware and software configurations

Read this topic for a description of the software and hardware configurations for the SUT. The scenario represents a complete customer-like environment with a Cognos BI Server and three online stock trading systems (DayTrader benchmark application).

The system is implemented in a single System z LPAR running a z/VM version 6.3 hypervisor on a z196 mainframe. The storage subsystem is a IBM DS8800 connected via FICON Express8.

IBM System z Enterprise 196 z/VM LPAR setup

The mainframe hardware for the described scenario consists of a single LPAR for z/VM version 6.3 on an z196 model 2817-M66.

The z196 system is equipped with:

- 10 central processors (CPs) supporting speeds up to 5.2 GHz;
  Integrated Facilities for Linux (IFLs) could have been used as well.
- 70 GiB central storage
- 1x 10GbE OSA-Express2 card
- 8x FICON Express8 S LX (short wave) features supporting ECKD™ for DASD device access

Note:

- The term CPU is used in the following for both kinds of processor types (IFL or CP).
- The term CP is used in the following for the z/VM Control Program.

IBM storage subsystem setup

An IBM System Storage DS8800 Model 951 is used as the storage subsystem.

DASD disk devices are used in the following configuration for the z/VM Linux guest operating systems and data disks:

- DASD Model 27, Model 54 and Model 128 directly attached to the guests
- selection from two ECKD storage pools (storage pool striping enabled), taken from both internal storage subsystem servers
- HyperPAV alias devices used

The details of the disk setup for each guest are listed in the pertaining topics later in this paper.

Note: FCP attached SCSI devices could also be used as disk storage.

Network Setup

This topic describes the network configuration used in the Cognos setup.

- z/VM VSWITCH for all z/VM virtual machines for guest-to-guest communication
- z/VM uses a VSWITCH and a 10 GbE OSA-Express3 card attached to the VSWITCH for external connectivity
- no network encryption

Cognos BI Server virtual machines (distributed installation)

This topic lists and describes all of the components that make up the complete distributed installation.

The settings for the z/VM relative CPU shares of the guests is described in Control virtual machine CPU capacity.

Cognos BI Content Manager

The Content Manager is the component that ensures integrity of the content store database (meta data repository). Transactions against this database are only done by this component.

The role of the Content Manager is a corner stone in the Cognos architecture, and only one can be active at any point in time. However, for continuous availability more than one can be installed and configured. Failover mechanisms will occur if the current one becomes unavailable.

Cognos meta data includes report definitions, scheduling information, security rules, data models, packages and so on.

The z/VM virtual machine resources for the Cognos BI Content Manager comprise the listed equipment:

- 4 virtual CPUs
- 8 GiB memory
- 10GbE LAN via VSWITCH
- 1x DASD model 27 for Linux operating system and Cognos BI software (20 GiB)
- 4x HyperPAV aliases available
- DASD devices and HyperPAV aliases are directly attached to the virtual machine

The Cognos BI Content Manager uses the software listed in Table 2:

Table 2. Cognos BI Content Manager software.

Software	Service level
Novell SUSE Linux Enterprise Server 11 for System z (64-bit)	Service Pack 3 (SP3)
IBM WebSphere Application Server Network Deployment v8.5.5 for Linux on System z (64-bit)	Fix pack 1 (8.5.5.1)
IBM WebSphere® SDK for Java™ Technology Edition 7	Version 7.0.4.1
IBM Cognos BI Server Content Manager	Version 10.2.1

Cognos BI Gateway

Web communication in Cognos BI typically works through gateways, which reside on one or more web servers. A Cognos BI Gateway is an extension of a web server, for example, the IBM HTTP server, through a CGI script or an Apache module.

The Cognos BI Gateway receives client requests and passes them to the report server dispatcher.

The z/VM virtual machines for the Cognos BI Gateway comprise the listed equipment:

- 4 virtual CPUs
- 4 GiB memory
- 10 GbE LAN via VSWITCH
- 1x DASD model 27 for Linux operating system and software (20 GiB)
- 4x HyperPAV aliases available
- DASD devices and HyperPAV aliases are directly attached to the virtual machine

The Cognos BI Gateway uses the software listed in Table 3:

Table 3. Cognos BI Gateway software

Software	Service level
Novell SUSE Linux Enterprise Server 11 for System z (64-bit)	Service Pack 3 (SP3)
IBM HTTP Server	Version 8.5.5.1
IBM Cognos BI Server Gateway	Version 10.2.1

Cognos BI Report Servers (Application tier)

A Cognos BI applications tier can contain one or more Cognos BI Report Servers. A Cognos BI Report Server runs requests, such as reports, analysis, and queries that are forwarded by a gateway.

The z/VM virtual machines for the Cognos BI Report Servers comprise the listed equipment:

- 4 virtual CPUs
- 16 GiB memory
- 10 GbE LAN via VSWITCH
- 1x DASD model 27 for Linux operating system and software (20 GiB)
- 4x HyperPAV aliases available
- DASD devices and HyperPAV aliases are directly attached to the virtual machine

The Cognos BI Report Servers use the software listed in Table 4:

Table 4. Cognos BI Report Servers software

Software	Service level
Novell SUSE Linux Enterprise Server 11 for System z (64-bit)	Service Pack 3 (SP3)
IBM WebSphere Application Server Network Deployment for Linux on System z (64-bit) v8.5.5	Fix pack 1 (8.5.5.1)
IBM WebSphere SDK for Java Technology Edition 7	Version 7.0.4.1
IBM Cognos BI Application Tier	Version 10.2.1

Cognos BI Content Store

The Cognos BI Content Store is a relational database that contains data that the Cognos BI Server needs to operate, such as report specifications, published models, and the packages that contain them, connection information for data sources, information about the external name space and the Cognos name space itself, and information about scheduling and bursting reports.

The z/VM virtual machines for the Cognos BI Content Store database server comprise the listed equipment:

- 2 virtual CPUs
- 4 GiB memory
- 10 GbE LAN via VSWITCH
- 1x DASD model 27 for Linux operating system and software (20 GiB)
- 1x DASD model 236 for backup purposes and database log files (180 GiB)
- 2x DASD model 128 as database container disks (approx. 200 GB for both)
- 20x HyperPAV aliases available
- DASD devices and HyperPAV aliases are directly attached to the virtual machine

The Cognos BI Content Store database server uses the software listed in Table 5:

Table 5. Cognos BI Content Store software

Software	Service level
Novell SUSE Linux Enterprise Server 11 for System z (64-bit)	Service Pack 3 (SP3)
IBM DB2 Enterprise Server Edition V10.5	Fix pack 2 (10.5.0.2)

Sample Database Server for Cognos BI queries

The sample database is a relational database that contains sample data for the Cognos BI queries triggered by the workload driver.

The z/VM virtual machines for the for the DB2 sample database server comprise the listed equipment:

- *2* virtual CPUs
- 4 GiB memory
- 10 GbE LAN via VSWITCH
- 1x DASD model 27 for Linux operating system and software (20 GiB)
- 1x DASD model 128 for backup purposes and database log files (100 GiB)
- 2x DASD model 128 for database container disks (approx. 200 GB for both)
- 20x HyperPAV aliases available
- DASD devices and HyperPAV aliases are directly attached to the virtual machine

The DB2 sample database server uses the software listed in Table 6

Table 6. DB2 Sample Database Server software

Software	Service level
Novell SUSE Linux Enterprise Server 11 for System z (64-bit)	Service Pack 3 (SP3)
IBM DB2 Enterprise Server Edition V10.5	Fix pack 2 (10.5.0.2)

DayTrader benchmark application server virtual machines

The tables in this topic list the virtual machines resources and software versions for all components of the distributed DayTrader server: Gateway Server, Application Server, and Database Server.

DayTrader installations with multiple guests (DayTrader triplets)

The setup with the DayTrader core components installed in multiple guests is used to investigate the behavior of these guests playing together in the system under test. When one of the guests does not get its required CPU resources, the transaction response time for DayTrader will increase. Additionally, the idea is to increase the number of guests competing with the Cognos BI guests for CPU resources.

The subsequent lists and tables provide information about the z/VM virtual machine resources and software versions for all types of DayTrader servers.

DayTrader Gateway Server virtual machine

- 1 virtual CPU
- 1 GiB memory
- 10GbE LAN via VSWITCH
- 1x DASD model 27 for Linux operating system and Cognos BI software (20 GiB)
- 4x HyperPAV aliases available
- DASD devices and HyperPAV aliases are directly *attached to the virtual machine*

The DayTrader Gateway Server uses the software listed in Table 7:

Table 7. DayTrader Gateway Server software

Software	Service level
Novell SUSE Linux Enterprise Server 11 for System z (64-bit)	Service Pack 3 (SP3)
IBM HTTP Server for WebSphere Application Server	Version 8.5.5.1

DayTrader Application Server virtual machine

- *4* virtual CPUs
- 8 GiB memory
- 10GbE LAN via VSWITCH
- 1x DASD model 27 for Linux operating system and software (20 GiB)
- 4x HyperPAV aliases available
- DASD devices and HyperPAV aliases are directly attached to the virtual machine

Table 8. DayTrader Application Server software

Software	Service level
Novell SUSE Linux Enterprise Server 11 for System z (64-bit)	Service Pack 3 (SP3)
IBM WebSphere Application Server Network Deployment (standalone)	Fix pack 1 (8.5.5.1)

DayTrader Database Server virtual machine

- 2 virtual CPUs
- 4 GiB memory
- 10GbE LAN via VSWITCH
- 1x DASD model 27 for Linux operating system and software (20 GiB)
- 4x HyperPAV aliases available
- DASD devices and HyperPAV aliases are directly attached to the virtual machine

Table 9. DayTrader Database Server software

Software	Service level
Novell SUSE Linux Enterprise Server 11 for System z (64-bit)	Service Pack 3 (SP3)
IBM WebSphere Application Server Network Deployment (standalone)	Fix pack 2 (10.5.0.2)

Single system DayTrader installation (DayTrader combo)

The DayTrader combo system combines all DayTrader core components in a single larger guest which competes with the Cognos BI guests for CPU resources.

The z/VM virtual machine resources for the DayTrader combo server comprise the listed equipment:

- 6 virtual CPUs
- 8 GiB memory
- 10GbE LAN via VSWITCH
- 1x DASD model 27 for Linux operating system and software (20 GiB)
- 4x HyperPAV aliases available
- DASD devices and HyperPAV aliases are directly attached to the virtual machine

The DayTrader combo server uses the software listed in Table 10:

Table 10. DayTrader combo server software

Software	Service level
Novell SUSE Linux Enterprise Server 11 for System z (64-bit)	Service Pack 3 (SP3)
IBM HTTP Server for WebSphere Application Server	Version 8.5.5.1
IBM WebSphere Application Server Network Deployment (standalone)	Fix pack 1 (8.5.5.1)
IBM DB2 Enterprise Server Edition V10.5	Fix pack 2 (10.5.0.2)

Client machine used as workload driver

This machine acts as the workload generator for the SUT. It drives custom workloads for the Cognos BI Server and the DayTrader benchmark application.

The hardware of the workload driver machine is listed in Table 11:

Table 11. Workload driver machine hardware

Client hardware	Setup
IBM Flex System	16 CPUs (2.9 GHz) 250 GiB memory 10 GbE network card for application server connectivity

The workload driver machine software versions are listed in Table 12:

Table 12. Workload driver machine software

Software	Service level
Red Hat Enterprise Linux Server Release 6	Service Level 2 (Santiago)
WebSphere Studio Workload Simulator	
Cognos BI Server Workload Simulator	

System under test - setup and tuning

Setup and tuning modifications applied to individual components of the system under test (SUT) are described in this topic. Changes were done for z/VM, Linux on System z, and Cognos BI Server components.

z/VM setup

All z/VM virtual machines defined for the various servers are outlined in detail in topic *Hardware and software configurations*. Besides the z/VM resource allocations, the following additional configuration and tuning was applied.

Virtual Networking

The z/VM guests use a Virtual Switch (VSWITCH) configured LAN for guest-to-guest communication. A 10 GbE OSA-Express3 is connected to the VSWITCH and the guests attached to the VSWITCH reside in the same LAN as the 10GbE OSA Express. This allows the configuration of guest IP addresses from the same network that is also used by the OSA Express card. As a virtual networking solution, the VSWITCH provides good performance and is the recommended method for internal and external z/VM network connectivity.

For more information about z/VM virtual networking and the VSWITCH refer to:
http://www.vm.ibm.com/virtualnetwork/

HyperPAV

The I/O throughput for DASDs can be improved by using parallel access volumes (PAV) or HyperPAV (see also Linux operation system setup).

For example, database servers are candidates for doing a lot of disk I/O. z/VM is set up to assign HyperPAV aliases directly to the virtual machines.

Control virtual machine CPU capacity

z/VM has the concept of shares to control the percentage of CPU power that a virtual machines receives. The CPU shares can be set with the SET SHARE command and the SHARE directory statement. It comes in two flavors as ABSOLUTE shares and RELATIVE shares.

ABSOLUTE shares allocate an absolute percentage of all available system processors to a virtual machine, and guarantee the availability of a certain percentage of processing time.

RELATIVE shares allocate a portion of the total system processors to a virtual machine, minus those processors allocated to virtual machines with an ABSOLUTE share. The range for relative shares is an integer number from 1 to 10000 (the larger the number, the higher the share). In general, you should assign a relative rather than an absolute share to typical users.

Definition of the fair share setup used for non-VMRM measurements

The default relative share for a virtual machine is 100. A common way to specify relative shares is to align them with the number of assigned virtual CPUs to the virtual machine.

For example, if virtual machines get a relative share of 100 per virtual CPU, a first virtual machine (VM1) gets a relative share of 200 for its two virtual CPUs, and a second virtual machine (VM2) gets a relative share of 400 for its four virtual CPUs. As a result, VM2 receives twice as much processing time for its four virtual CPUs as VM1. This ensures that each virtual CPU has the same weight within all virtual machines.

On the other hand, if all guests have the same relative share in a CPU constraint scenario, they get exactly the same amount of CPU capacity, independently of the amount of virtual CPUs for the guest.

QUICKDSP option
QUICKDSP is intended for selective use on virtual machines with critical response time requirements. The scheduler always moves a QUICKDSP user immediately into the dispatch list whenever it is ready to run, regardless of resource requirements and current system load. The virtual machine never waits in the eligible list, if there is one. This option is only recommended for mission critical servers.

None of the virtual machines in the SUT had QUICKDSP enabled.

z/VM Resource Manager (VMRM)

The z/VM Resource Manager (VMRM) provides functions to dynamically tune the z/VM system. Groups of virtual machines can be defined to be part of a workload. The workloads can then be managed by VMRM according to goals that are also defined. Virtual machines can be grouped into different workloads, each managed according to different goals. VMRM automatically adjusts performance parameters when there is contention between virtual machines for a resource. Therefore, VMRM has no effect in environments where resources are unconstrained and CPU resource management is not required.

VMRM uses CP monitor data to obtain regular measurements of virtual machine resource consumption. Based on the supplied definition of workloads and workload goals from a configuration file, VMRM will adjust virtual machine tuning parameters to achieve those goals.

Note:
- Only the CPU management capability of VMRM was used for this project.
- VMRM Cooperative Memory Management and velocity targets for DASDs were not considered.
- VMRM is compatible with ILMT (IBM License Metric Tool).
- VMRM is NOT compatible with CPU pooling.

Sample VMRM configuration file
VMRMSVM is a predefined multi-configuration virtual machine that can be enabled to run the VMRM code. This user ID also maintains the VMRM configuration file on its A-disk.

```
ADMIN MSUSER VMRMADMN
GOAL MAXCPU VELOCITY CPU 70
GOAL MAXCPU VELOCITY CPU 25
WORKLOAD DAYTRADER USER TRDUS*
MANAGE DAYTRADER GOAL MINCPU IMPORTANCE 1
WORKLOAD COGNOS USER COGNUS*
MANAGE COGNOS GOAL MAXCPU IMPORTANCE 10
```

ADMIN MSGUSER
defines the user ID that will receive messages from VMRMSVM.

GOAL CPU
defines a goal that is comprised of velocity targets for CPU (1-100).

Specifies the percentage of time that the workload should receive CPU resources when it is ready to consume them. This is computed by taking the time the users are running, and divided through the sum of time the users are running or waiting for CPU.

WORKLOAD
defines a workload comprised of one or more virtual machines. Identified by a user name, account ID, or ACI group name.

MANAGE
associates a workload to a goal with the importance (1-10) of achieving this goal.

A low and a high velocity target for CPU is defined for the Cognos BI and the DayTrader workloads. The virtual machines (COGNUS*) belong to the Cognos BI Server and are grouped together with the WORKLOAD statement. Also all DayTrader virtual machines (TRDUS*) are grouped. If no ACI group is used, one can use similar user IDs with a common part and shorten the differing part with an asterisk.

The MANAGE statement associates a workload with its CPU velocity goal and goal importance.

The shown sample VMRM configuration describes the following workload constraints:

- Manage the Cognos BI workload with a CPU velocity target of 70 and the importance 10 (preferred workload in a CPU constraint system)
- Manage the DayTrader workloads with a CPU velocity target of 25 and the importance 1 (will receive less share in a CPU constraint system)

Note: The percentage of the CPU goal parameter is calculated in regard to wait times and **not** CPU load. When managing workloads in a prioritized way, a big advantage of using VMRM is that the CPU goals can be defined for workload groups and not only for individual guests like for the standard relative share definition. That means adding or removing guests to a workload group does not change the CPU velocity goals for the workload. In case of a manually managed share setup, where the relative shares are individually set for each system, a chance for the number of guests would require a change of the relative shares for all guests.

Interaction with the CP monitor
VMRM requires CP MONITOR SAMPLE data in order to monitor virtual machine performance. VMRM starts sample monitoring with an interval of one minute if not active. If it is already active, then VMRM does not start sample monitoring and uses whatever interval is already set.

Performance Toolkit
There is a VMRM screen in the Performance Toolkit (FCX241) that gives some information about the currently active VMRM configuration.

Refer to the z/VM V6 R3.0 Information Center:
http://pic.dhe.ibm.com/infocenter/zvm/v6r3/index.jsp

Customization and Tuning > z/VM: Performance > Tuning z/VM Performance > VMRM Tuning Parameters

Linux operation system setup

The virtual machine resource definitions for the Cognos BI and DayTrader servers are outlined in detail in topic *Hardware and software configurations*. Besides that, additional operating system tuning and setup changes were applied as described in this topic.

Networking

The Cognos BI Server components communicate over a virtual guest-to-guest network using a z/VM Virtual Switch (VSWITCH). The virtual network devices appear as `GuestLAN QDIO` type in Linux when the network devices are listed with the *lsqeth* command.

Sample command: `lsqeth` for a virtual network device

```
# lsqeth eth1
Device name                    : eth1
-------------------------------------------
        card_type              : GuestLAN QDIO
        cdev0                  : 0.0.1000
        cdev1                  : 0.0.1001
        cdev2                  : 0.0.1002
        chpid                  : 02
        online                 : 1
        portname               : osaport
        portno                 : 0
        route4                 : no
        route6                 : no
        checksumming           : sw checksumming
        state                  : UP (LAN ONLINE)
        priority_queueing      : always queue 2
        fake_broadcast         : 0
        buffer_count           : 128
        layer2                 : 0
        large_send             : no
        isolation              : none
        sniffer                : 0
```

The buffer count for the network devices was increased for all virtual machines:

- *Increased the number of buffers for inbound traffic to **128** (default is 64) at* `/etc/udev/rules.d/51-qeth-0.0.1000.rules`

  where `1000` *is first* `cdev` *of the virtual network device in this case.*

- *The following line was added to the* `udev` *rule definition to enhance the buffer count:*
- `ACTION=="add", SUBSYSTEM=="ccwgroup", KERNEL=="0.0.1000", ATTR{buffer_count}="128"`

See also **Linux on System z - Tuning hints & tips**:
http://www.ibm.com/developerworks/linux/linux390/perf/

Disk I/O on ECKD DASD

The I/O throughput for an ECKD DASD disk can be improved by using Parallel Access Volumes (PAV) or HyperPAV. This feature can be important for DASDs used with the Cognos BI Content Manager and Content Store where more disk I/O is expected.

The Linux DASD device driver can use this IBM System Storage feature to perform multiple concurrent data transfer operations to or from the same DASD instead of single data transfers. To use HyperPAV, there must be base and alias devices available, which require System z Input/Output Configurations Data Set (IOCDS) definitions. For HyperPAV on an IBM System Storage subsystem, the alias devices are not exclusively referenced to a certain base device, but they are eligible for all base devices in the same logical control unit (LCU). Linux handles HyperPAV alias devices in the same way as normal DASD base devices by using the `chccwdev` command or defining the appropriate `udev` rules for them. When listing the DASD devices with the `lsdasd` command, HyperPAV aliases can be identified by the `alias` status tag. The usage of the HyperPAV aliases is completely handled by the Linux kernel and is not noticeable by the users.

The sample command `lsdasd` shows four DASD devices and 20 HyperPAV aliases for the Content Store database server:

```
# lsdasd
Bus-ID     Status      Name      Device  Type  BlkSz  Size      Blocks
==============================================================================
0.0.6980   alias                       ECKD
0.0.6981   alias                       ECKD
0.0.6982   alias                       ECKD
0.0.6983   alias                       ECKD
0.0.6e80   alias                       ECKD
0.0.6e81   alias                       ECKD
0.0.6e82   alias                       ECKD
0.0.6e83   alias                       ECKD
0.0.6e84   alias                       ECKD
0.0.6e85   alias                       ECKD
0.0.6e86   alias                       ECKD
0.0.6e87   alias                       ECKD
0.0.6f80   alias                       ECKD
0.0.6f81   alias                       ECKD
0.0.6f82   alias                       ECKD
0.0.6f83   alias                       ECKD
0.0.6f84   alias                       ECKD
0.0.6f85   alias                       ECKD
0.0.6f86   alias                       ECKD
0.0.6f87   alias                       ECKD
0.0.6926   active      dasda     94:0    ECKD  4096   21129MB   5409180
0.0.6e01   active      dasdb     94:4    ECKD  4096   184688MB  47280240
0.0.6e07   active      dasdc     94:8    ECKD  4096   100170MB  25643520
0.0.6f07   active      dasdd     94:12   ECKD  4096   100170MB  25643520
```

For details about IOCDS configurations, see the *Input/Output Configuration Program User's Guide*, SB10-7037-11 at the Resource Link® (with a required registration):
http://www.ibm.com/servers/resourcelink

A detailed study addressing HyperPAV and databases can be found in the IBM white paper:
Oracle Database on Linux on System z - Disk I/O Connectivity Study

Other Linux operating system adaptions

For the Cognos BI Server Linux guests, especially the components that run under control of a WebSphere Application Server, an increased limit for open files and number of processes was set.

Content of `/etc/security/limits.conf`**:** setting limits for number of open files and processes

```
# cat /etc/security/limits.conf

# /etc/security/limits.conf
#
#Each line describes a limit for a user in the form:
#
#<domain>       <type>  <item>          <value>
...

*               -       nofile          65512
*               soft    nproc           32768
*               hard    nproc           32768
```

The limits can be permanently set in the Linux configuration file `/etc/security/limits.conf`.

By using the `ulimit` command, the values can be changed in a running Linux system. For example, you can set the number of open files:

```
ulimit  -n 16384
```

Cognos BI Gateway

The IBM Cognos BI Web server tier can contain one or more IBM Cognos BI Gateways that handle the web communication. The SUT includes a single Cognos BI Gateway deployed in its own virtual machine.

The Cognos BI Gateway is an extension of a web server (for example, IBM HTTP Server) and knows the location of a dispatcher service running in a Cognos BI application tier (for example, a report server). The dispatcher is a multithreaded service that uses one or more threads per request.

When an IBM Cognos BI Gateway receives a request, it

- encrypts passwords to ensure security
- extracts information needed to submit the request to an IBM Cognos BI server
- attaches environment variables for the web server
- adds a default name space to the request to ensure that the server authenticates the user in the correct name space
- passes requests to an IBM Cognos BI dispatcher for processing

For more information on Cognos BI Gateways and the Dispatcher, refer to the IBM Knowledge Center for Cognos Business Intelligence 10.2.1:

- Cognos Business Intelligence 10.2.1 > Install > Business Intelligence Architecture and Deployment Guide 10.2.1 > The Multi-tiered Architecture > Tier 1. Web Server: IBM Cognos BI Gateways
- Cognos Business Intelligence 10.2.1 > Install > Business Intelligence Architecture and Deployment Guide 10.2.1 > The Multi-tiered Architecture > Tier 2. Applications: IBM Cognos BI Servers > Dispatcher

Gateway tuning

Apache modules on IBM HTTP Server (IHS)

For enhanced performance and throughput, the default CGI gateway is replaced by an Apache module (*apache_mod)* gateway that can be used with the IBM HTTP Server.

This change applies to the IHS configuration file (`httpd.conf`).

The Apache module is added at the end of the IHS load module list:

```
LoadModule cognos_module <cognos10_location>/cgi-bin/mod2_2_cognos.so
```

The complete procedure is described under the following IBM Knowledge Center navigator:
Cognos Business Intelligence 10.2.1 > Install > Business Intelligence Installation and Configuration Guide 10.2.1 > Installing and Configuring Product Components one One Computer > Configuring the web server > Use compiled gateways for production systems > Use Apache modules on Apache Server or IBM HTTP Server

Provide enough virtual CPUs

The Cognos BI Gateway performance depends on the number of available CPUs. Figure 2 shows that the different Cognos BI loads (increasing number of users) perform better with four virtual CPUs instead of only two. The gateway benefits from the higher degree of parallelism with four virtual CPUs.

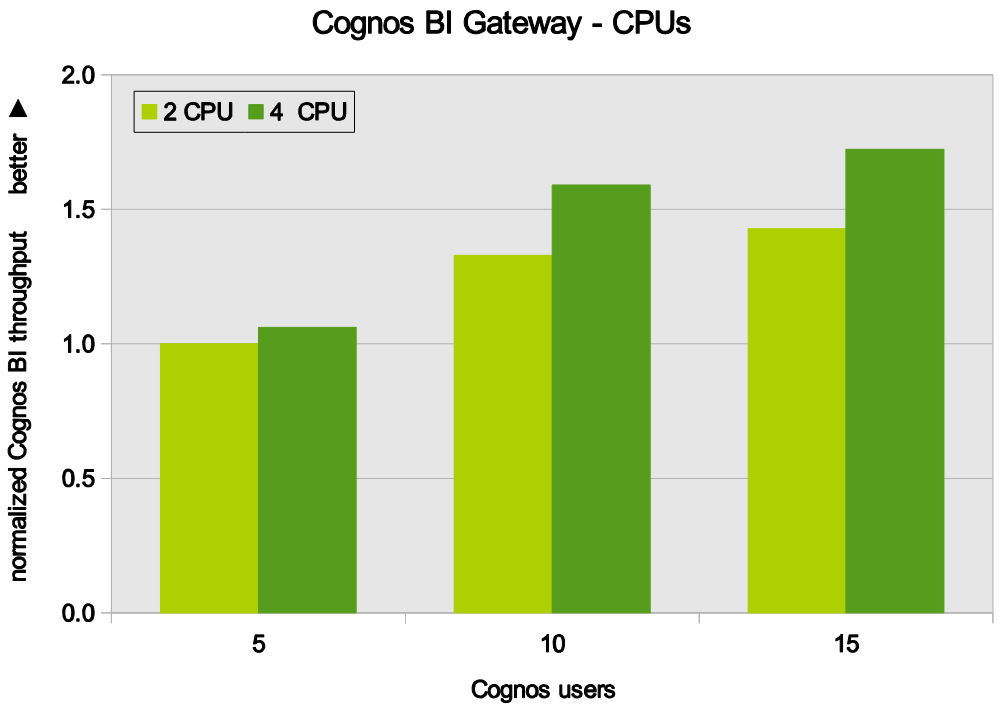

Figure 2. Cognos BI transaction throughput - two compared to four gateway CPUs

If the Cognos BI Gateway does not have enough CPUs, it quickly becomes a bottleneck. This is obvious when looking at the CPU utilization for the Cognos BI load with 10 users in Figure 3, for example.

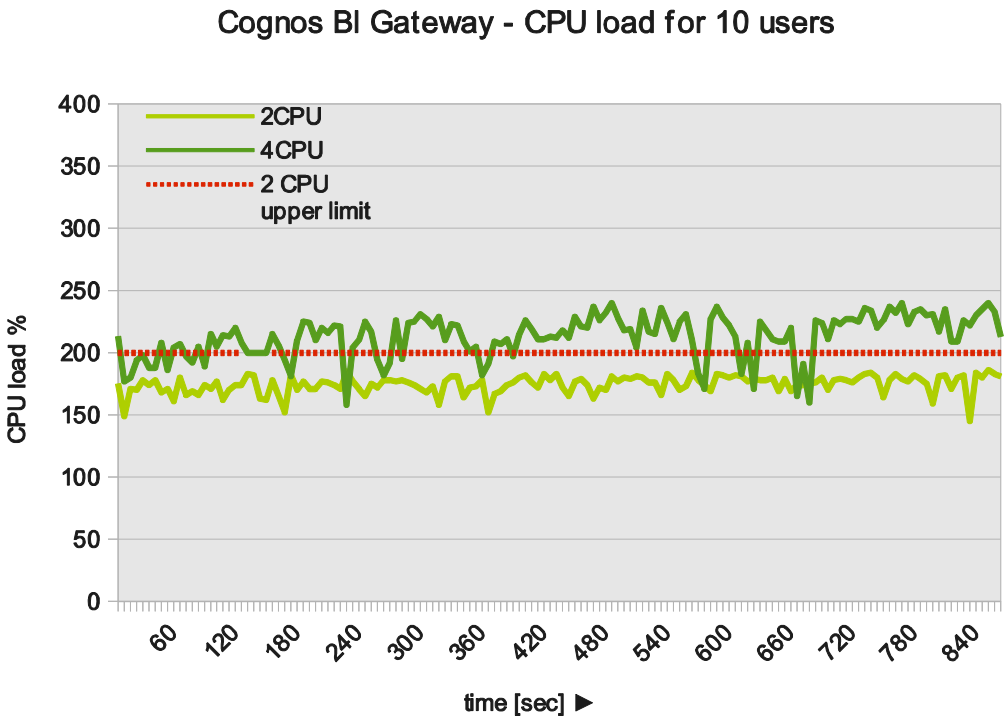

Figure 3. Cognos Gateway - CPU load 2 versus 4 Gateway CPUs

The workload level can be covered with two virtual CPUs at first sight. The CPU load is high but not yet constrained. When the same workload runs with four virtual gateway CPUs, the transaction throughput and the CPU load increase at the same workload level.

Cognos BI application tier

The Cognos BI application tier can contain one or more servers (for example, report servers). Such a server runs requests, for example, reports, analysis, and queries, that are forwarded by a gateway.

The SUT has two Cognos BI report servers as application tier, each in its own virtual machine. Both virtual machines also run a dispatcher service that starts and controls the report service in its instance.

The Cognos BI Gateway was configured to send any requests to the first available dispatcher service in its list of known dispatchers. When a dispatcher service starts, it registers itself with the Content Manager. As a result, each dispatcher is aware of the other dispatchers and can route requests to other present dispatchers. This enables the capability of workload balancing between the dispatchers and hence also between both report servers. The Cognos BI Gateway itself does no request routing between the Cognos BI report servers.

For more information on the Cognos BI Application tier, refer to the IBM Knowledge Center for Cognos Business Intelligence 10.2.1:

- Cognos Business Intelligence 10.2.1 > Install > Business Intelligence Architecture and Deployment Guide 10.2.1 > The Multi-tiered Architecture > Tier 2. Applications: IBM Cognos BI Servers

Also refer to the following technical article on IBM developerWorks®:
IBM Business Analytics Proven Practices: IBM Cognos BI Dispatcher Routing Explained

Report Server tuning

WAS Java virtual machine heap size

The WAS report server initial and maximum Java virtual machine (JVM) heap size was set to 1024 MiB and 2048 MiB respectively. Increasing the JVM heap may often improve performance.

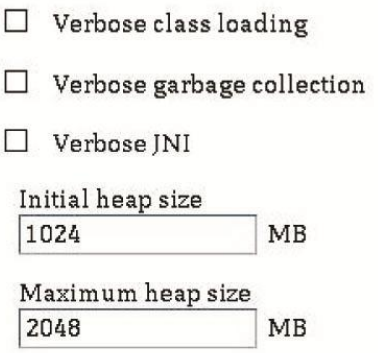

WAS administration console path:
Servers -> Application Servers -> server1 -> Java and process management -> Process Definition -> Java Virtual Machine

WepSphere WebContainer thread pool

The WAS Thread Pool WebContainer maximum was increased from default 50 to 500. Increasing the maximum will improve scalability for the Cognos BI report server.

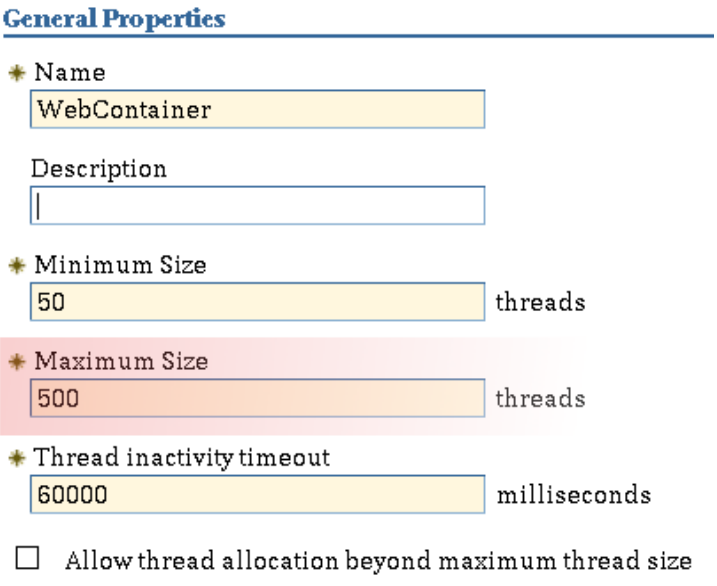

WAS administration console path:
`Servers -> Application Servers -> Server1 -> Thread Pools -> WebContainer`

Report Server session cache
The default processing behavior for the Cognos BI Report Server can be changed by modifying entries in the `rsvpproperties` file. Depending on your specific IBM Cognos BI application and requirements, changing settings in the `rsvpproperties.xml` file may improve the performance.

A `rsvpproperties.xml.sample` file is located in the `<cognos10_location>/configuration` directory.

If required, just copy the sample file to the file `rsvpproperties.xml`.

One parameter which is important for the performance, is the maximum size of the session cache maintained by the report server. It was increased from 20 to 100.

```
..
<property>SessionCacheSize</property
   <value type="long">100</value>
...
```

Note: Increased caches might have unintended side effects if the system can not satisfy higher memory requirements.

Number of processes for the report service
The report service, batch report service, and report data service have several settings that you can configure to optimize the use of system resources.

There are a number of processes associated with the report service and the batch report service. When these services receive requests from the dispatcher, they start processes to handle the requests (also known as BiBus processes). You can specify the maximum number of processes that these services can start any time.

The report server execution mode was explicitly set to 64-bit for the SUT.

For more information on the report server execution mode see:

- *Differences between running a 32 and 64 bit Report Server with IBM Cognos 10.1 RP1 (10.1.1) and above*
- *IBM Knowledge Center for Cognos Business Intelligence 10.2.1:*
  Cognos Business Intelligence 10.2.1 > Install > Installation and Configuration 10.2.1 > Installing and Configuring Server Components on Different Computers > Installing and Configuring Application Tier Components > Enable the 64-bit version of report server

The number of processes should be configured based on the amount of available capacity provided by the Cognos report server. In general, report processing is a CPU-bound process. As a result, the number of CPUs for the server are the main variables to keep in mind when adjusting this setting from the default value of 2.

Figure 4 shows the transaction throughput and response times when scaling the number of report service processes from two to eight processes.

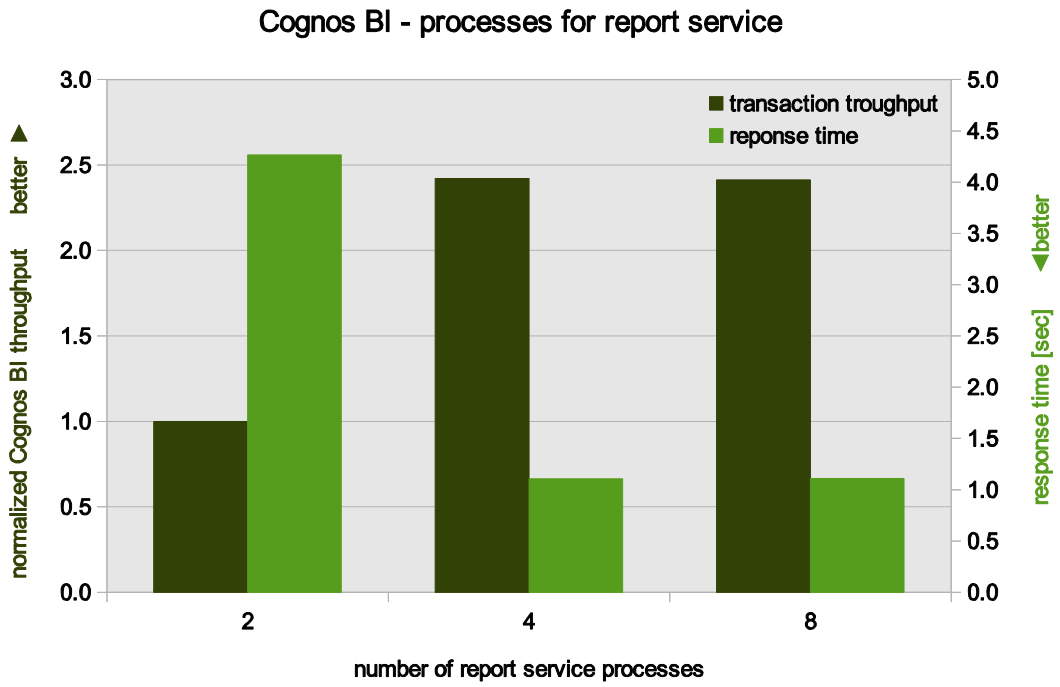

Figure 4: Scaling the number of report service processes

The two report servers in the SUT are defined with four virtual CPUs each. The default value of two system-wide report service processes for a single report server gives lower results in this case.

When setting the number of processes equal or higher than the number of available CPUs, the report server is much better utilized and delivers a higher transaction throughput and lower response times.

Setting the number of processes for the report service for peak and non-peak periods:

☐	Tuning	Number of high affinity connections for the report service during peak period	2	Yes
☐	Tuning	Number of low affinity connections for the report service during peak period	4	No
☐	Tuning	Maximum number of processes for the report service during peak period	4	No

Cognos Connection path:
Configuration -> Dispatchers and Service -> Set properties -> Report Service ->Tuning -> Settings

Query Service - Java virtual heap size

The initial and maximum heap size for the Query Service was increased to 4096 MiB.

☐	Tuning	Initial JVM heap size for the query service (MB) (Requires QueryService restart)	4096	No
☐	Tuning	JVM heap size limit for the query service (MB) (Requires QueryService restart)	4096	No

Cognos Connection path:
Configuration -> Dispatchers and Service -> Set properties -> Query Service -> Settings

Remarks on the Report Server memory footprint
All tuning mentioned above indicates that the memory footprint for the report server can get very big. The following memory-consuming components need to be considered together:

- WAS JVM heap size (2 GiB)
- Query service JVM heap size (4 GiB)
- number of processes for the report service (can grow to more than 1 GiB memory per process)

Figure 5 shows the memory consumption for four report service processes over three hours at a continuous workload level.

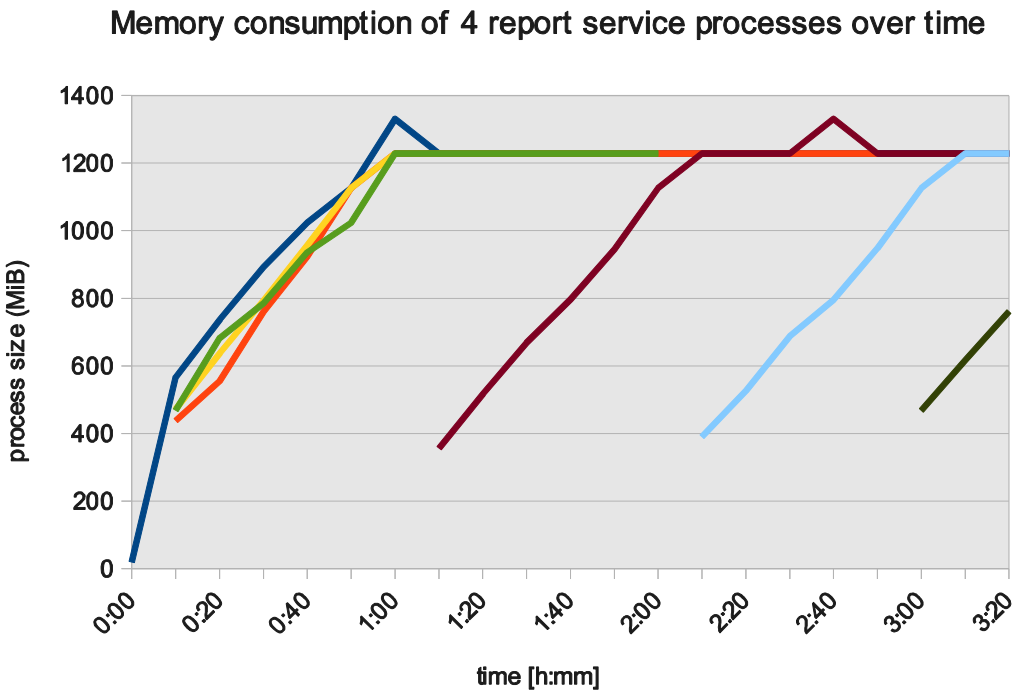

Figure 5: Report service process memory consumption over time

All four processes grow to 1.2 GiB memory in parallel. From time to time, a single process becomes idle and is stopped by the report service. Due to the ongoing workload, a new process is created that again grows over time.

Conclusion: There is a recommendation of setting up two report service processes (31-bit) per available report server CPU. We used 64-bit report service processes for the SUT and there was no performance difference between one process per CPU versus two processes per CPU. But with two processes per CPU, the memory requirements for the processes may also double.

The Cognos BI report servers were enlarged to 16 GiB memory each in order to satisfy the memory requirements of already 11 GiB for Cognos services and WAS Java heaps for one instance.

Cognos BI Content Manager and DB2 Content Store

The Cognos BI Content Manager is a service that manages the storage of customer application data, including security, configuration data, models, metrics, report specifications, and report output. It is needed to publish packages, retrieve or store report specifications, and manage scheduling information.

The Content Manager stores its information in a Content Store database, which is typically located in tier three of the three-tier architecture model.

The Content Store for the SUT is an IBM DB2 database. However, other relational databases are supported as well.

The Content Manager uses JDBC (Java DataBase Connectivity) to access the Content Store.

Much of the information in the Content Store, such as report results and XML report specifications, is stored as binary large object (BLOB) fields.

For more information on the Cognos BI Content Providers, refer to the IBM Knowledge Center for Cognos Business Intelligence 10.2.1:

- Cognos Business Intelligence 10.2.1 >Install > Business Intelligence Architecture and Deployment Guide 10.2.1 > The Multi-tiered Architecture > Tier 3. Data: Content Providers

Content Store tuning (DB2)

The following DB2 settings were adjusted according to the guideline to create a Content Store database.

Database configuration parameters

- Application heap size (`applheapsz`)
  - Recommendation: `AUTOMATIC` or at least 1024 KiB
  - Improvement: If the application heap size value is too small, out of memory errors may occur when there are many users.
- Lock timeout (`locktimeout`)
  - Recommendation: 240 seconds
  - Improvement: Do not set this parameter to an infinite value. The default is -1, which means that lock timeout detection is turned off and a lock will be waited for until it is granted or a deadlock occurs.

For more information on the Cognos BI Content Providers, refer the IBM Knowledge Center for Cognos Business Intelligence 10.2.1:
Cognos Business Intelligence 10.2.1 > Install > Business Intelligence Installation and Configuration Guide 10.2.1.1 > Preparing to install > Guidelines for creating the content store > Suggested settings for creating the content store in DB2 on Linux, Windows, and UNIX operating systems

DB2 registry variables

- `DB2_INLIST_TO_NLJN`
  - Recommendation: set to `YES`
  - Improvement: Setting this query compiler variable to `YES` improves performance by controlling optimization decisions about your database, such as forcing optimization decisions and `SQL` query operations.

For more information on DB2 tuning, refer to the IBM Knowledge Center for Cognos Business Intelligence 10.2.1:

- *Cognos Business Intelligence 10.2.1 > Install > Business Intelligence Installation and Configuration Guide 10.2.1 > Performance Maintenance > Tuning a DB2 Content Store*
- *Cognos Business Intelligence 10.2.1 > Install > Business Intelligence Installation and Configuration Guide 10.2.1 > Preparing to install > Guidelines for creating the content store*

Workload descriptions

Performance testing was done using workload generation applications on an IBM Flex System client machine. Two different workloads were required for the tests discussed in this paper.

The workload generator for Cognos BI is an IBM developed application. The competing workloads are an Open Source benchmark application emulating Online Stock Trading systems (DayTrader).

Cognos BI workload

The workload generator for Cognos BI allows to put a couple of scenarios together until a typical Business Intelligence transaction mix is formed. The different scenarios have different shares in the workload mix and are executed in a given sequence by every emulated user over and over again.

The workload shares for the different scenarios are selected to focus on the latest technologies introduced with Cognos 10, for example, *Dynamic Query Mode (DQM)* and the new charting engine.

Two different DB2 databases contain some sample data used for BI processing. The sample databases reside in their own virtual machine separated from the Content Store database.

Table 13 shows the percentages for the different selected workload scenarios.

Table 13. Cognos BI workload scenario percentages

Cognos BI scenario	Workload share
interactive reporting and ad hoc analysis uses Dynamic Query Mode (DQM)	30%
Chart and report creation uses the new charting engine	20%
Dashboard activity users open multiple pages in a browser session	30%
retrieve PDF reports users access saved PDFs reports in the Content Store	10%
retrieve HTML reports users access saved HTML reports in the Content Store	10%

Adjusting the workload level

The number of users serves to adjust the CPU load level for the Cognos BI Server. For the performance tests and the tests with the z/VM Resource Manager, well-defined CPU load levels for Cognos BI were required.

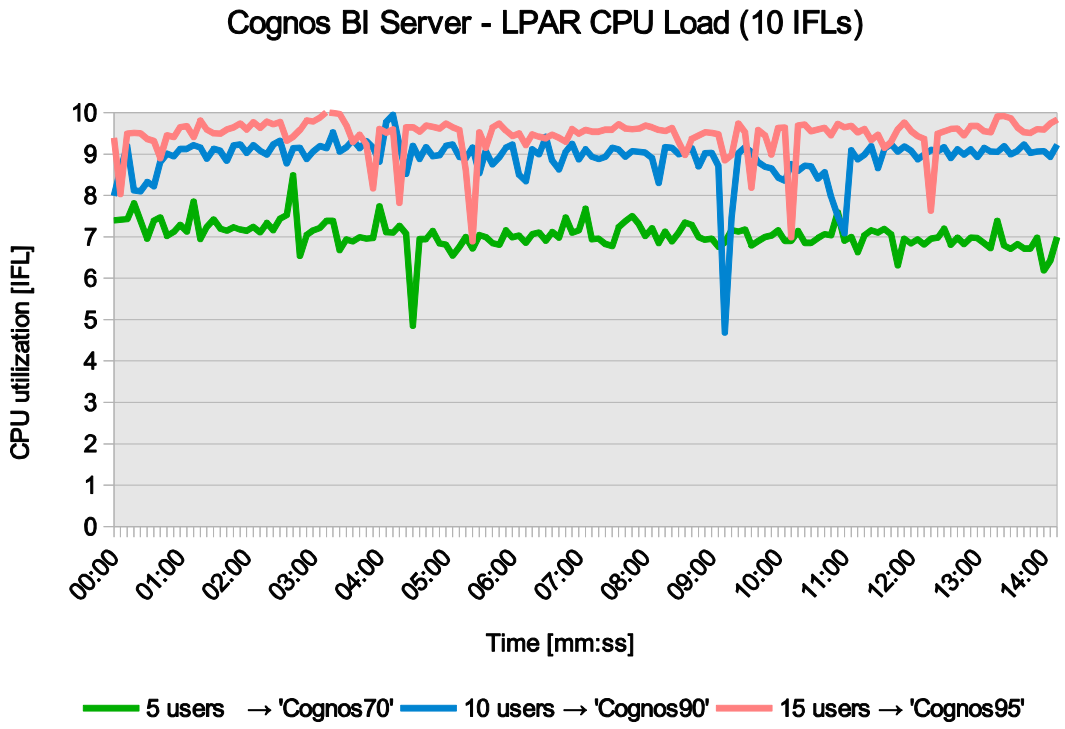

Figure 6: LPAR total CPU load for a different number of Cognos users

Figure 6 shows the overall LPAR CPU load levels with different numbers of users for the Cognos BI Server. The LPAR is defined with 10 IFLs and z/VM virtualizes the environment. For the CPU load level tests, the Cognos BI workload is the only workload running. The z/VM guests with the DayTrader workload are not running at this time. These Cognos BI CPU load levels are used in topic Results. Table 14 introduces three labels representing a certain Cognos BI Server CPU load level.

Table 14. Cognos BI Server CPU load levels

Cognos BI Server CPU load levels	Description
Cognos70	▪ Workload runs with 5 concurrent users ▪ Approx. 70% of the available IFLs are used ▪ Default CPU load level for measurements
Cognos90	▪ Workload runs with 10 concurrent users ▪ Approx. 90% of the available IFLs are used ▪ Used only for selected test scenarios
Cognos95	▪ Workload runs with 15 concurrent users ▪ Approx. 95% of the available IFLs are used ▪ Almost all IFLs are used ▪ Used only for some test purposes

DayTrader benchmark application

DayTrader is an Open Source benchmark application emulating an online stock trading system. Users can log in and view their accounts or portfolios. In the Quotes/Trade section they can buy or sell stock shares out of their portfolios. The benchmark application is currently hosted by the Apache Geronimo project and was originally developed by IBM as the WebSphere Trade performance benchmark. This benchmark was donated to the Apache Geronimo project in 2005.

DayTrader is an end-to-end Java Enterprise Edition (J2EE) web application composed of several Java classes, Java Servlets, Java Server Pages, Web Services and Enterprise Java Beans (EJB). Thus makes it an ideal benchmark application for measuring the scalability and performance of a J2EE application server like IBM WebSphere Application Server (WAS). The presentation layer is based on Java Servlets and Java Server Pages, whereas the back-end business logic uses Java database connectivity (JDBC), Java Message Service (JMS), EJB and Message Bean techniques. The TradeDatabase is hosted on the DB2 database server and is accessed via JDBC from the application server.

Users can be emulated with a workload driver and generate load against the DayTrader application.

For further details, see the DayTrader documentation at:
http://geronimo.apache.org/GMOxDOC30/daytrader-a-more-complex-application.html

The DayTrader benchmark application has been chosen as a competing workload for the Cognos BI Server workload. Therefore the SUT has three independent DayTrader installations capable of running three DayTrader workloads in parallel.

Two of them are setups where IHS, WAS, and the DB2 back-end database each run in separated z/VM virtual machines (DayTrader triplets). The third one combines all components in a single z/VM virtual machine (DayTrader combo).

Adjusting the workload level

The number of users set in the workload generator is used to adjust the CPU load level for all three DayTrader benchmarks running in parallel. Likewise as for the Cognos BI Server, a few well-defined CPU load levels for DayTrader are required for performance testing.

Table 15. DayTrader CPU load levels

DayTrader CPU load levels	Description
DayTrader25	▪ Workload runs with 5 users ▪ Approx. 25% of the available IFLs are used ▪ Low DayTrader CPU load level
DayTrader50	▪ Workload runs with 125 users ▪ Approx. 50% of the available IFLs are used ▪ Medium DayTrader CPU load level
DayTrader75	▪ Workload runs with 190 users ▪ Approx. 75% of the available IFLs are used ▪ High DayTrader CPU load level

Results

The Cognos BI and the WAS DayTrader workloads are executed in parallel in unconstrained and constrained CPU scenarios. A major question for this performance study is the impact on the Cognos BI transaction throughput and response times caused by competing workloads and how good throughput and response times can be guaranteed in such scenarios.

For this purpose, two z/VM test scenarios were considered. The first scenario, z/VM fair share setup, ensures that all virtual CPUs for the z/VM guests have the same relative share. That means all guests get the same weight according to their number of virtual CPUs (see Control virtual machine CPU capacity).

The second scenario is described in two subtopics.

- VMRM parameters at first introduces the z/VM Resource Manager and its CPU management capabilities.
- VMRM - managing workload peaks then describes the test setup and presents the measurement results. In this scenario, VMRM is used to control the relative shares for the guests according to defined management goals. These CPU management goals should define a higher priority for the Cognos BI workload than for the competing workloads. As a result, the Cognos BI workload should run privileged in a constrained CPU environment (see z/VM Resource Manager (VMRM)).

z/VM fair share setup

For this test scenario, all virtual CPUs have the same relative share.

The Cognos BI CPU load level is set to Cognos70 (see Table 14) and the DayTrader workloads are scaled in parallel from CPU load level DayTrader25 to DayTrader75 (see Table 15).

This results in the following LPAR CPU loads for the SUT:

- Cognos70 + DayTrader25 -> high workload with LPAR CPU close to 100% (unconstrained)
- Cognos70 + DayTrader50 -> higher workload with LPAR CPU load above 100% (constrained)
- Cognos70 + DayTrader75 -> highest workload with LPAR CPU load above 100% (constrained)

In the fair share setup, every virtual CPUs gets the same relative share of 100. So the guest weights for the z/VM hypervisor are dependent on the number of virtual guest CPUs. The larger guest for the DayTrader combo has 6 virtual CPUs assigned and therefore has a relative share of 600. The DayTrader triplets have relative shares of 100/400/200 for their three guests. When comparing guest by guest, the relative share for each individual DayTrader triplet guest is lower than for the DayTrader combo guest, however the sum for the DayTrader triplets is higher.

Transaction throughput

Figure 7 shows the transaction throughput for the Cognos BI and the DayTrader workloads for the z/VM fair share setup.

There are three measurement series with different DayTrader CPU load levels (DayTrader25, DayTrader50 and DayTrader75). The Cognos BI CPU load level is at Cognos70.

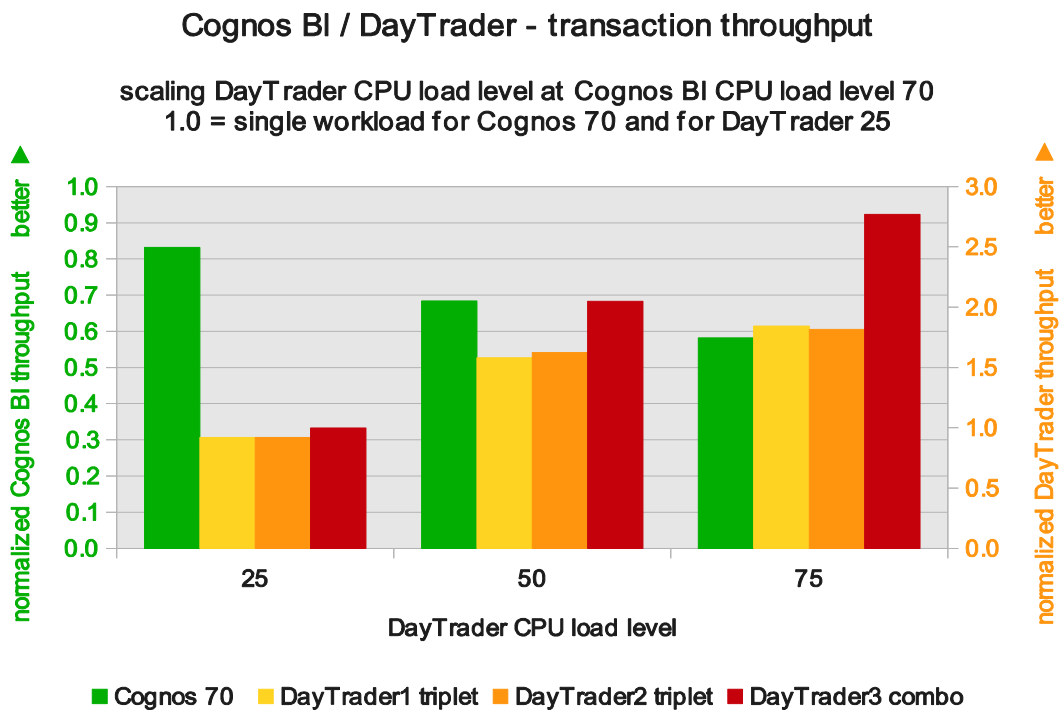

Figure 7: Transaction throughput for Cognos BI and DayTrader workloads

Observations: The Cognos BI transaction throughput is normalized against the throughput number for Cognos70 running as a single workload in the SUT. Hence, the single Cognos70 transaction throughput number is equivalent with 1.0.

The DayTrader transaction throughput is normalized against the throughput number for DayTrader25 running as a single workload in the SUT. The transaction throughput numbers for the DayTrader triplets and the combo are set to 1.0 for each workload.

Example: Cognos70 + DayTrader25 are running in parallel. Cognos70 achieves a normalized transaction throughput of 0.8 and DayTrader achieves a throughput of around 2x 0.9 for the triplets, and 1.0 for the combo compared to their single results. With that metric, the maximum possible sum of all relative throughputs is 4.0 for all workloads together.

With an increasing CPU load level for the DayTrader workloads, the Cognos BI transaction throughput goes down from 0.8 to less than 0.6. The DayTrader transaction throughputs increase according to their CPU load levels (higher number of users). For the workload mix Cognos70 + DayTrader75, the DayTrader combo scores a much higher throughput than the DayTrader triplets.

Conclusion: In regard to throughput, the Cognos BI workload shows a significant impact in reaction to the increasing DayTrader workload running in parallel. Both triplets show a moderate increase in throughput. The winner of this competition in regard to throughput is the DayTrader combo. It scales exactly with the increasing load level. That indicates that guests with a larger number of virtual CPUs in a fair share setup have an advantage because of their higher relative shares.

Transaction response times

Figure 8 shows the response times for the Cognos BI and the DayTrader workloads in the z/VM fair share setup.

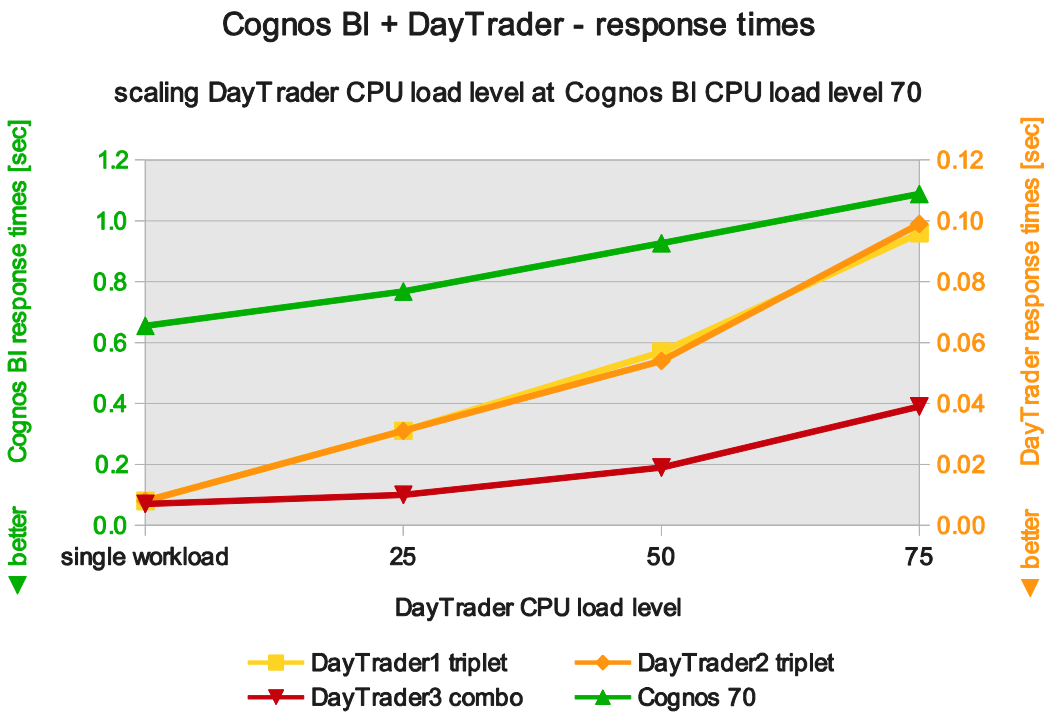

Figure 8: Response times for Cognos BI and DayTrader workloads

Observations: Figure 8 shows the response times for the Cognos BI (green) and the DayTrader workloads (yellow, orange, and red). The single workload response times stand for the non-parallel execution of each workload. When executed in parallel, the response times increase, the higher the CPU load level for both workloads is. For the most constrained scenario Cognos70 + DayTrader75, the DayTrader combo has only a moderate increase in its response time compared to the DayTrader triplets.

The response times for Cognos BI transactions vary from 0.6 to 1.1 seconds.

Conclusion: The Cognos BI and the DayTrader transactions have different characteristics:

- Cognos BI has longer running transactions with response times in the range of a second.
- DayTrader has short running transactions with response times in the range of milliseconds (10x - 50x shorter runtime).

It seems that the different transaction characteristics behave differently in regard to response times, while the short running DayTrader transactions have an advantage in regard to throughput, they are very sensitive in regard to response times when the CPU resource runs short in the fair share setup. Here the impact on the longer running Cognos BI transaction is more moderate.

LPAR CPU load

Figure 9 shows the LPAR CPU loads (used IFLs) for Cognos BI and DayTrader in the fair share setup. The measurement series include the different DayTrader CPU load levels (DayTrader25, DayTrader50, and DayTrader75). The Cognos BI CPU load level is at Cognos70 for all measurements.

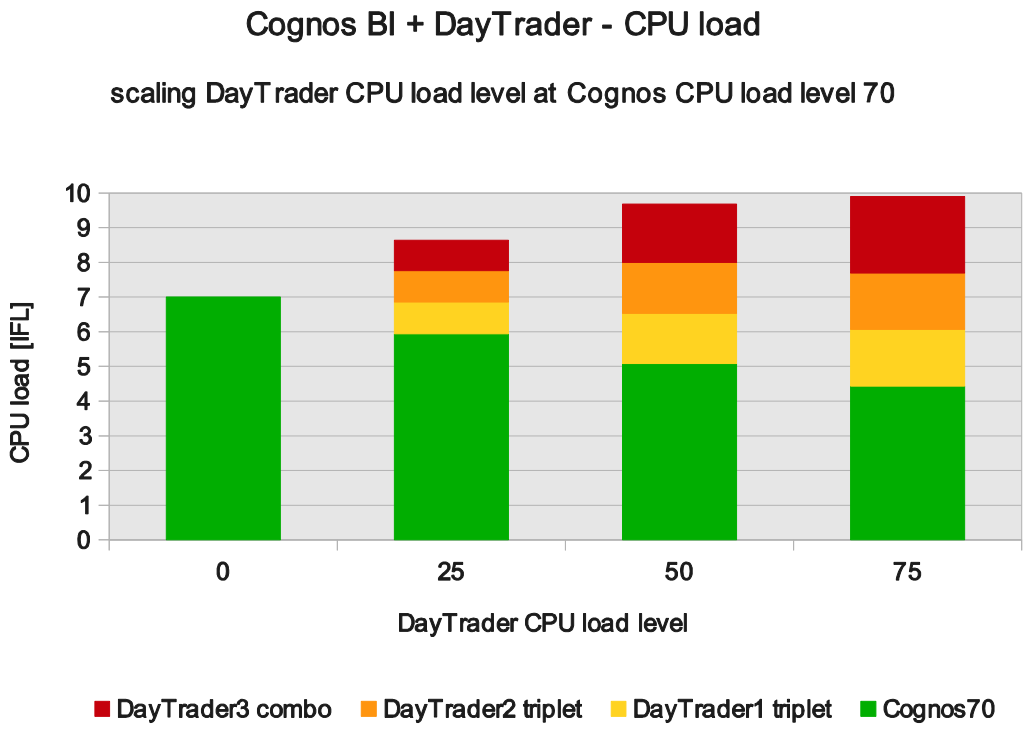

Figure 9: CPU load for Cognos BI and DayTrader workloads

Observations: The Cognos70 single measurement (DayTrader CPU load level 0) consumes exactly 7 out of 10 IFLs.

The other three bars show the LPAR CPU loads for Cognos BI and DayTrader running in parallel.

Cognos70 + DayTrader25 show that the LPAR CPU resources are not yet constrained. All DayTrader workloads consume the same amount of CPU resources.

For Cognos70 and DayTrader50/75, the workload for DayTrader is further increased. LPAR CPU resources are constraint now. The portion of CPU resources required for DayTrader is higher for these combined workload mixes. Cognos BI gets decreasing CPU resources for its transactions. At DayTrader load level 75, there are around 4.5 IFLs remaining for the Cognos BI workload.

The DayTrader combo gets more CPU resources compared to DayTrader triplets in the Cognos70 + DayTrader75 scenario.

Conclusion: In the fair share setup, the DayTrader workloads can get CPU resources for their guests, while the Cognos BI workload cannot hold its CPU load level. This behavior is expected and intended in a fair share setup. The surprising part was that the DayTrader combo, which has the highest relative share of 600 compared to the other guests in the SUT, gets much more CPU resources than the DayTrader triplets. The DayTrader triplets and Cognos BI get less CPU resources according to the relative fair shares for their guests. Even the larger Cognos BI report servers and DayTrader WAS servers, which had relative shares of 400, receive significantly less CPU resources than the DayTrader combo workload.

Figure 10 shows the CPU load per Cognos BI guest at the Cognos70 + Daytrader50 CPU load level.

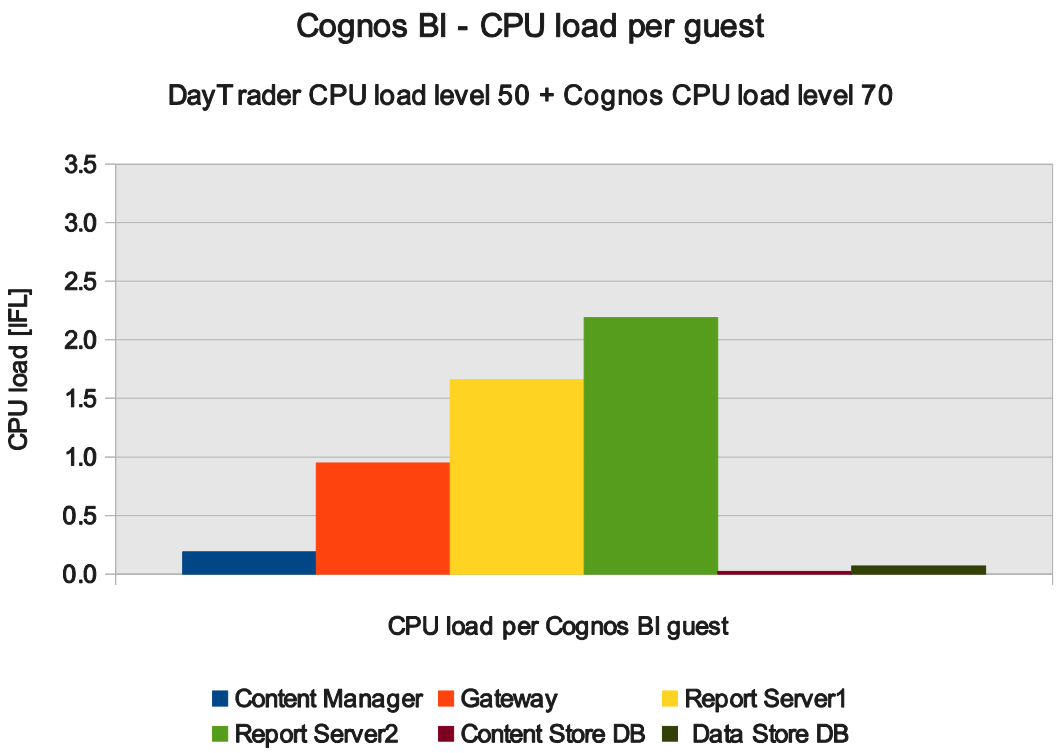

Figure 10: Cognos BI - CPU loads per guest

The Cognos BI workload mix has a quite high portion of transactions using Dynamic Query Mode (DQM) and the new charting engine. This type of workload is CPU intensive and mainly generates load on the Cognos BI Gateway and the two Report Servers. The dispatcher on the report servers pass requests to both report servers for load-balancing needs.

In general it seems that in a fair share setup a larger guest (couple of virtual CPUs) has an advantage over setups with distributed installations in multiple guests that have higher shares in sum. The expectation is that the impact of competing workloads can be mitigated when changing to a prioritized share setup where Cognos BI guests get higher relative shares than the DayTrader guests.

Each manual setup is a fixed setup. This means that when the CPU pressure of the competing guests is just high enough, it will take away a significant amount of resources from the preferred workload. However, when running on the most extreme configuration, it will give the Cognos BI guests the highest share possible, and the DayTrader guests the lowest relative share possible.

But that bears the risk that the DayTrader workload will run in request processing time outs because of lacking CPU capacity. The measurement results for the fair share setup confirm the assumption that competing workloads will have an impact on the Cognos BI workload. This is true either for the transaction throughput or response times and goes back to the reduction of assigned CPU capacity from the z/VM hypervisor according to the guest fair shares.

VMRM parameters

For the second test scenario, the z/VM setup was modified to enable z/VM Resource Manager (VMRM). The Cognos BI and the DayTrader workloads were put into separate VMRM workload groups.

Then a high velocity target for CPU is defined for Cognos BI, and a low one is defined for the DayTrader (see sample definition in z/VM Resource Manager (VMRM).

The first measurement series helps to understand the influence of the VMRM parameters and shows how they impact the transaction throughputs and response times for both workloads. Especially these parameters were of interest:

- CP Monitor facility sample interval
- VMRM CPU goal importance
- VMRM CPU goals for the workloads

Finally, two VMRM configurations are selected to study the management of a sudden workload peak period resulting in a time frame with a CPU constraint for the SUT.

CP monitor sample interval

VMRM requires CP monitor sample data in order to monitor the performance of a virtual machine. VMRM starts sample monitoring, if it is not active. The default CP monitor sample interval is 60 seconds. Shorter sample intervals are tested to reduce the reaction time for changed CPU load situations and their influence on the Cognos BI transaction throughput and response time.

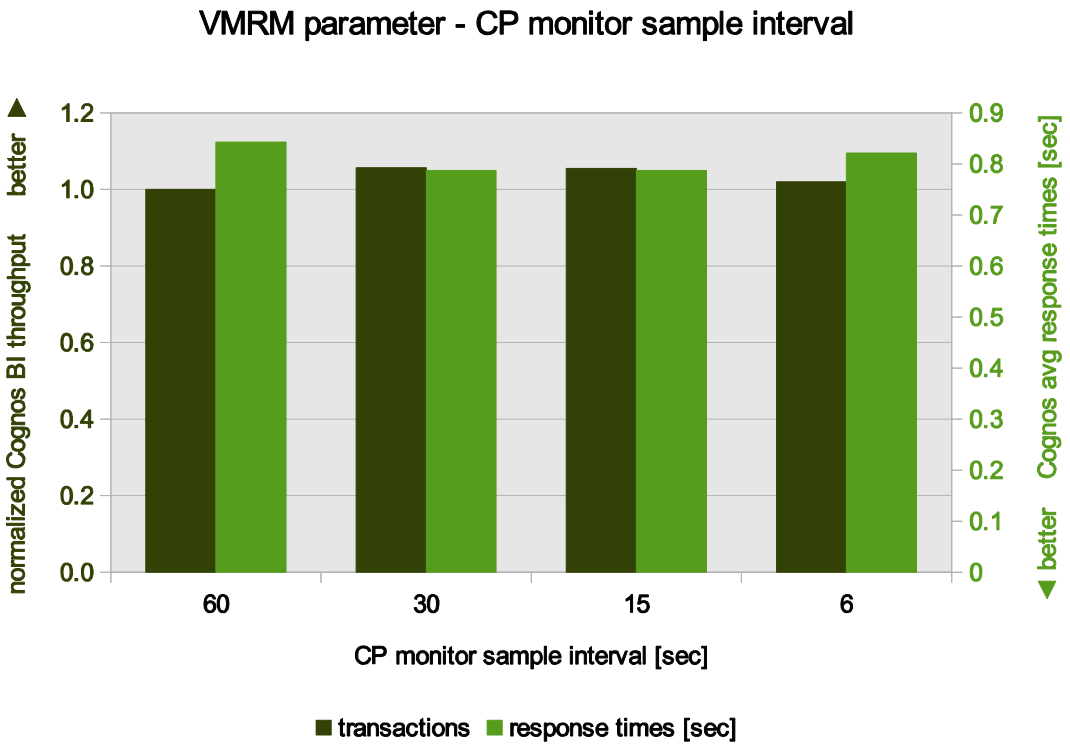

Figure 11: VMRM parameters - influence of the CP monitor sample interval

Observations: The CP monitor sample interval is scaled down from 60 to 6 seconds. The dark bars show the transaction throughput for Cognos BI and the light bars the corresponding response time. The measurements with 15 and 30 seconds for the sample interval show around 6% higher throughput numbers and more than 6% shorter response times.

Conclusion: The default CP monitor sample is 60 seconds, which is a long time for a resource management interval. To shorten the time to react on workload load peaks, a smaller sample interval has been chosen, which improves throughput and response times. A too short interval can introduce unnecessary overhead and degrades the performance. An interval value between 15 and 30 seconds shows the best behavior in our environment. For all other VMRM measurement series, a sample interval of 30 seconds has been chosen.

Goal importance

The VMRM goal importance is an integer number from 1 to 10. The higher the number, the more important is the goal for the workload to be managed.

The following CPU goal importance combinations for the workloads are tested:

- Cognos BI importance 10 / DayTrader importance 1 -> highest importance for Cognos BI
- Cognos BI importance 8 / DayTrader importance 3 -> high importance for Cognos BI
- Cognos BI importance 5 / DayTrader importance 5 -> same importance for both workloads

The workload CPU load levels for this measurement series are Cognos70 + DayTrader50 resulting in a constrained LPAR CPU load. The VMRM CPU goal was set 70 for Cognos BI and 25 for DayTrader for all measurements in this series.

Figure 12 shows the Cognos BI and DayTrader transactions when scaling the goal importance.

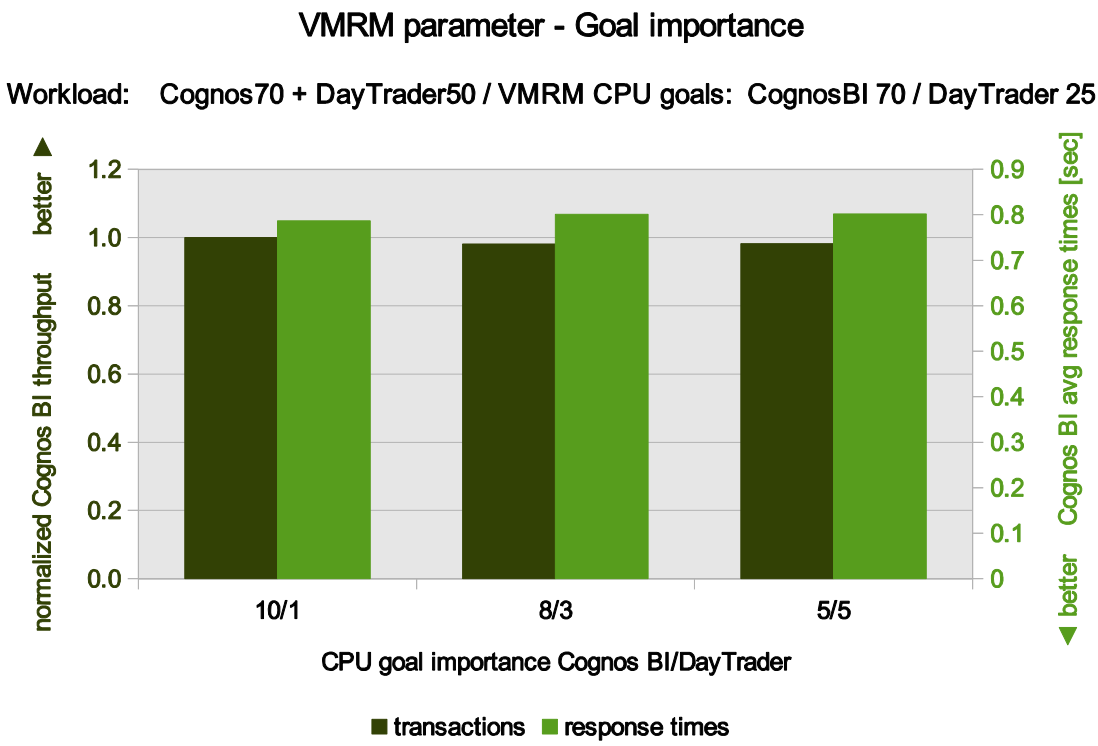

Figure 12: VMRM parameters - influence of the goal importance

Observations: There is a slight advantage for the 10/1 CPU goal importance for the Cognos BI workload. However, the impact of this parameter is not very relevant for the tested scenarios.

Conclusion: The 10/1 CPU goal importance (highest priority for Cognos BI) is chosen for the next VMRM measurements. This goal importance also matches best the objective to privilege the Cognos BI workload over all other workloads.

Probably this parameter becomes more important when there are more than two CPU goals and workload groups defined.

Workload CPU goals

The VMRM CPU goals can be an integer number from 1 to 100. The number specifies the percentage of time that a workload should receive CPU resources.

The following CPU goal combinations for the workloads were tested:

- Cognos BI CPU goal 60 / DayTrader CPU goal 25 -> moderate CPU goal for Cognos BI
- Cognos BI CPU goal 70 / DayTrader CPU goal 25 -> high CPU goal for Cognos BI
- Cognos BI CPU goal 80 / DayTrader CPU goal 25 -> high CPU goal for Cognos BI
- Cognos BI CPU goal 100 / DayTrader CPU goal 1 -> maximum CPU goal for Cognos BI

The workload CPU load levels are set to Cognos70 + DayTrader50 / DayTrader75 resulting in a constrained LPAR CPU. The VMRM CPU goal importance is set to 100 for Cognos BI and to 1 for DayTrader for all measurements in this series.

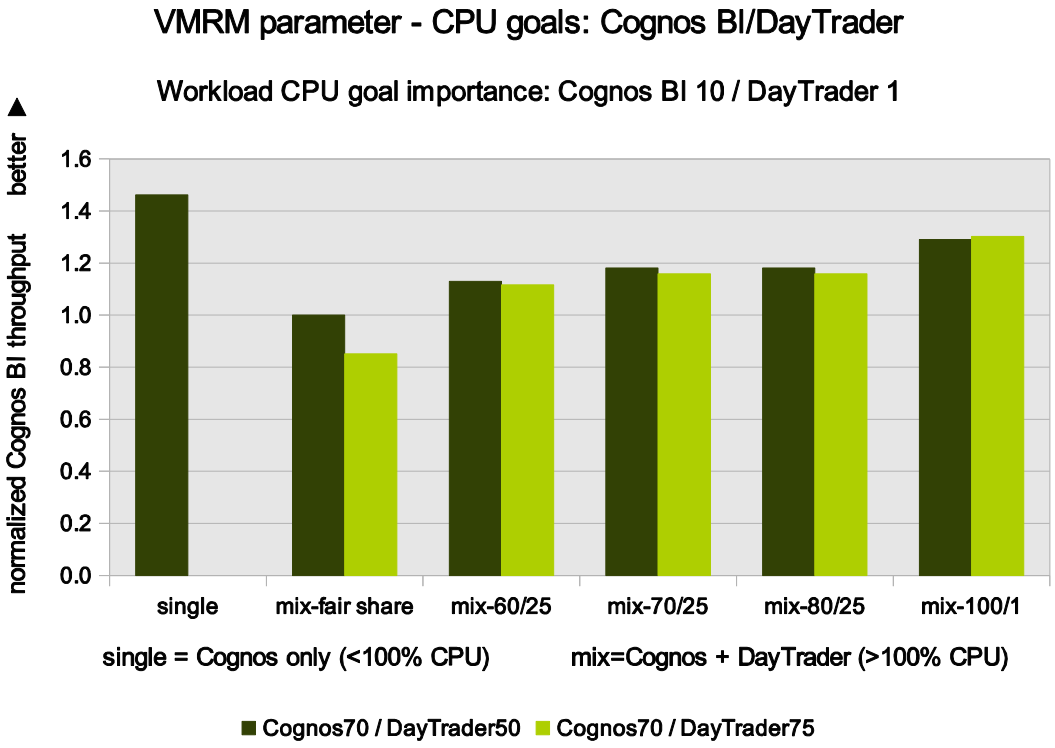

Figure 13: VMRM parameters - CPU goals: Cognos BI transaction throughput

Observations: The Cognos BI transaction throughput is normalized against the fair share mixed workload scenario Cognos70 + DayTrader50. The leftmost single bar shows the transaction throughput for Cognos70 running as a single workload on the system without any CPU constraint. All other mixed scenarios run in a constraint CPU environment.

In the fair share scenario, the transaction throughput decreases due to the parallel DayTrader workloads running. The higher the pressure from the DayTrader workloads (DayTrader75), the greater is the impact on the Cognos BI workload.

The four rightmost bar pairs show the transaction throughputs for the VMRM scenarios with the different CPU goal combinations. The moderate and high CPU goals for Cognos BI score a higher transaction throughput compared to the fair share scenario. The maximum CPU goal shows the best throughput getting closer to the single Cognos BI only throughput.

Even with a higher DayTrader load level (DayTrader75), the Cognos BI transaction throughput stays at the same level as with a lower DayTrader50 load level for a certain CPU goal combination.

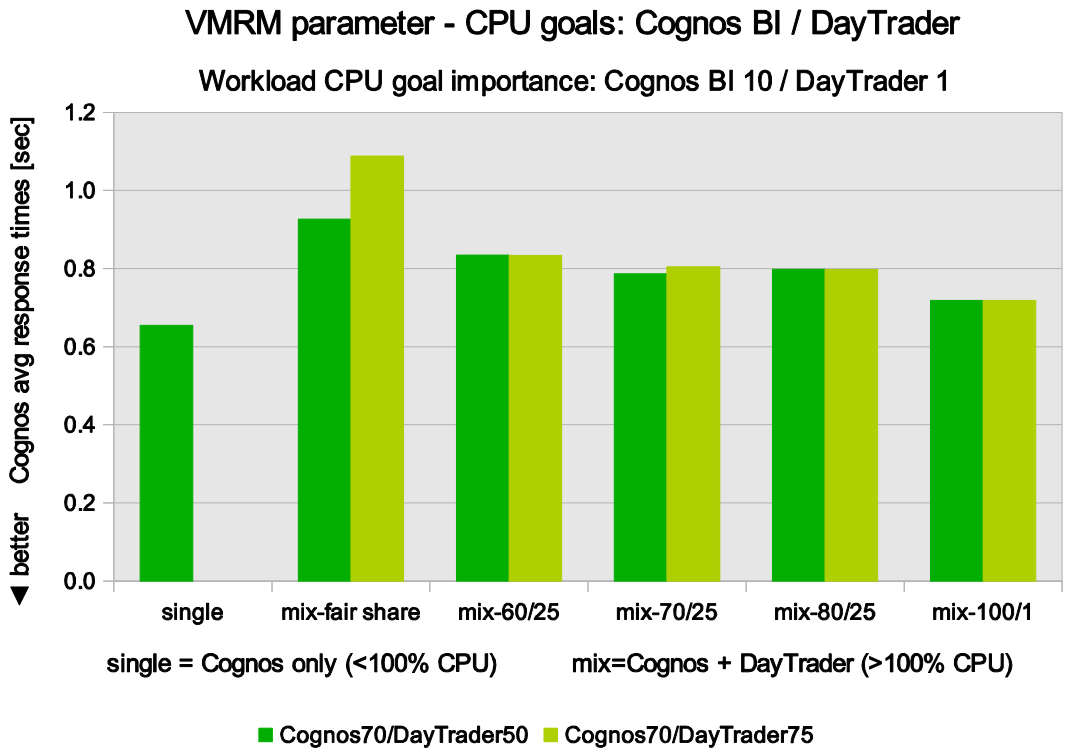

Figure 14: VMRM parameters - CPU goals: Cognos BI response times

Observations: The leftmost single bar shows the best Cognos BI transaction response time for Cognos70, running as a single workload on the system without any CPU constraint. All other mixed scenarios run in a constraint CPU environment.

The fair share scenario shows the highest response times due to the parallel DayTrader workloads running. The higher the pressure from the DayTrader workloads (DayTrader75), the higher is the response time for the Cognos BI transactions.

The four rightmost bar pairs show the response times for the VMRM scenarios with different CPU goal combinations. The moderate and high CPU goals for Cognos BI show better response times compared to the fair share scenario. With the maximum CPU goal, the response time is getting closer to the response time of the single Cognos70 workload.

Even with a higher DayTrader load level (DayTrader75), the Cognos BI transaction response times stay at the same level as with a lower DayTrader50 load level for a certain CPU goal.

Conclusion: In all VMRM configurations (moderate, high and maximum), the Cognos BI workload is preferred over the DayTrader workload and scores higher throughput numbers for Cognos BI compared to the fair share setup. The behavior for response times is similar.

The maximum CPU goal results in the best transaction throughput and response time of all VMRM configurations. This meets the expectation of this goal to give Cognos BI the highest priority. But this setup bears the risk that the DayTrader workloads get nearly no CPU resources and may run into a time-out. The SUT was set up to avoid such a scenario by assigning three CPUs more to z/VM than the Cognos BI workload would use at its maximum load level. But in a more constrained environment with seven CPUs for example, Cognos BI would require all CPUs and the remaining CPU resources might become critical for the lower prioritized workloads.

Another important aspect is that the Cognos BI transaction throughput and response time keep its levels regardless of the DayTrader CPU load level. The relative shares are adjusted, so that the CPU goal is met whatever the pressure from the competing workload is. This demonstrates how useful VMRM really is to guarantee a certain load level for the Cognos BI workload.

The definition of a CPU goal is the most important tuning parameter for the VMRM CPU management capability.

CPU loads for some workload CPU goal combinations

Figure 15 shows the Cognos BI and different DayTrader CPU load total percentages for the system. The mixed workload scenario Cognos70 + DayTrader50 is used for all CPU goal combinations.

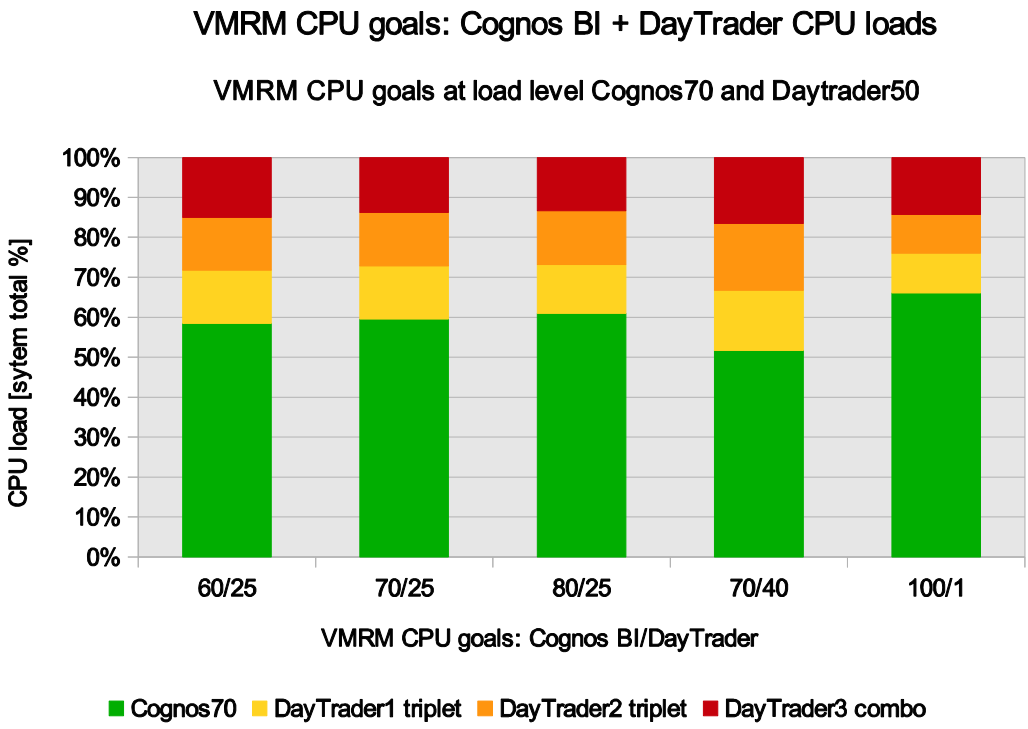

Figure 15: VMRM CPU goals - CPU loads for some VMRM CPU goal combinations

Observations: The five bars show the stacked CPU load percentages for the Cognos BI and DayTrader workloads. The mixed workload scenario Cognos70 + DayTrader50 results in a CPU constrained system. The total system CPU load percentage is at 100% for all VMRM CPU goal combinations.

For the three leftmost CPU goal combinations, the Cognos BI goal was scaled from 60 to 80, and DayTrader was kept at 25. The Cognos BI CPU percentages slightly increases to 60%. With the CPU goal combination 70/40, the DayTrader workloads receive a higher CPU goal. The Cognos BI CPU percentage decreases from around 60% to 50% compared to the CPU goal combination 70/25.

The rightmost bar shows the maximum CPU goal 100/1 in favor to Cognos BI. In this case, Cognos BI gets more than 65% of the CPU resources. Keep in mind that the maximum CPU percentage for the CPU load level Cognos70 would be 70%.

Conclusion: This measurement series shows the influence on the assignment of CPU resources for some different CPU goal combinations. Varying CPU goal combinations allow a fine-tuning for the assignment of CPU resources that a particular workload should receive in a CPU constraint system.

Notice that a CPU goal combination of 70/40 does not mean Cognos BI gets 70% of the available CPU resources. Furthermore this is the percentage of time that the workload should receive CPU resources when it is ready to consume them. This results in 50% of the available CPU resources in this case.

VMRM - managing workload peaks

The measurement series discussed in this topic demonstrate the capabilities of VMRM in managing workload peaks in two scenarios. The first scenario examines the performance when Cognos BI is the preferred workload, but a part of the CPU resources is also assigned to other tasks. The second scenario satisfies all CPU requests of Cognos BI, at the cost of other workloads.

Transaction throughput and response times

Two VMRM configurations are tested how they manage an abrupt workload (CPU load) peak, where all LPAR CPUs are fully utilized and which results in a CPU bound system.

Goal 1 - A chance for others: Whenever the system runs into such a described CPU constrained situation, Cognos BI should be the preferred workload. But other workloads still get a chance to perform some work. The VMRM CPU goals are set as follows:

- Cognos BI CPU goal 70 / DayTrader CPU goal 25

Goal 2 - Cognos BI gets all: Whenever the system runs into such a described CPU constrained situation, Cognos BI should get all CPU resources it requires, even at the cost of other running workloads. The VMRM CPU goals are set as follows:

- Cognos BI CPU goal 100 / DayTrader CPU goal 1

CPU load peak: The overall load starts unconstrained with the Cognos70 workload plus one DayTrader IFL in parallel for 10 minutes. Then the second DayTrader triplet and the DayTrader combo workloads are started. All DayTrader workloads together add up to the DayTrader50 load. From that point on, the LPAR CPUs are constrained for another 10 minutes and VMRM starts to manage the workloads according to their CPU goals. Afterwards, the DayTrader workloads are stopped except for one single DayTrader triplet and the system is unconstrained for the last 10 minutes.

The VMRM CPU goal importance is set to 100 for Cognos BI and is set to 1 for DayTrader for all measurements in this series.

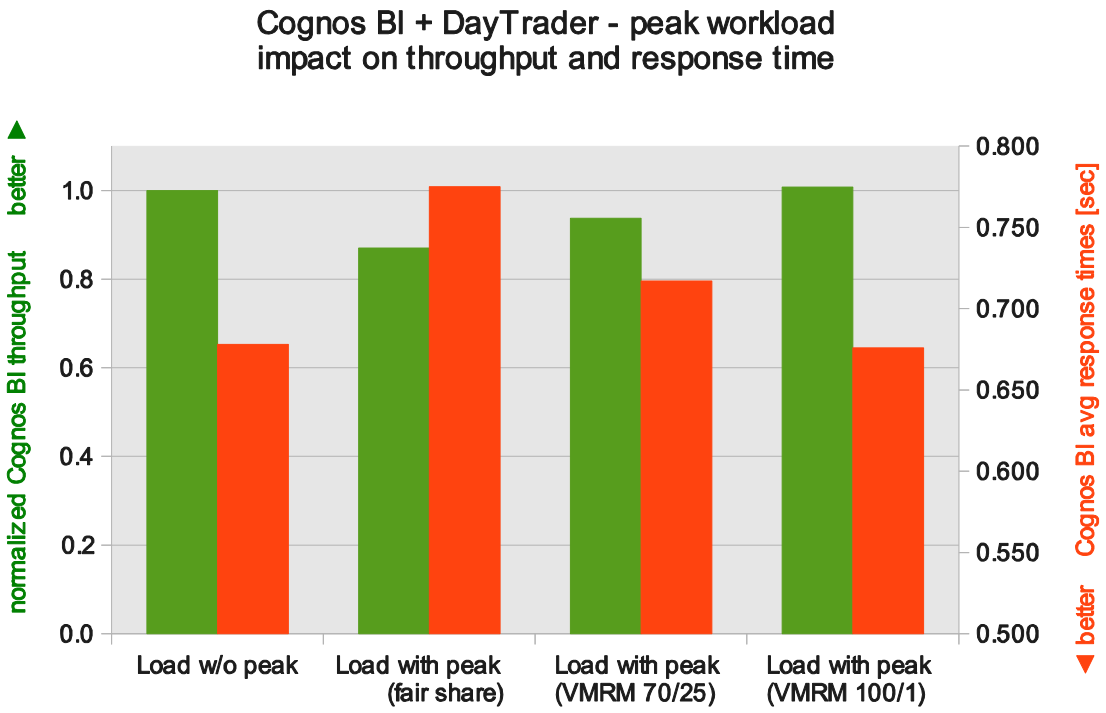

Figure 16: VMRM managing workload peak - impact on Cognos BI throughput and response time

Observations: The Cognos BI transaction throughput is normalized against the unconstrained combined Cognos + DayTrader workload without the peak load. The green bar shows the Cognos BI transaction throughput and the red bar the response time.

The rightmost three bar pairs are the measurements including a 10 minute workload peak for the fair share and the two VMRM configurations (Goal 1 and Goal 2).

The fair share setup shows a clear throughput loss and increased response time.

The VMRM CPU goal configuration (70/25) called **A chance for others** minimizes the throughput loss compared to the fair share setup. There is also a much lower impact on the response time.

The VMRM CPU goal configuration (100/1) called **Cognos BI gets all** shows the same performance as the workload without CPU load peak.

Conclusion: The VMRM configurations match their goals and guarantee Cognos BI a defined transaction throughput level and response time when a sudden high peak load arises and leads to constraint CPU resources for the system.

The **Cognos BI gets all** configuration almost straightens out the CPU peak period and maintains the throughput as with an unconstrained CPU load.

VMRM seems to be very useful in protecting a preferred workload against other workloads running in parallel when the CPU resource becomes a bottleneck. For example, it achieves a much more efficient usage of CPU resources required for important workloads that run only in certain time slots.

VMRM CPU goal 1: a chance for others (70/25)

The VMRM CPU goals are set as follows for the workloads:

- Cognos BI CPU goal 70 / DayTrader CPU goal 25 -> high CPU goal for Cognos BI

Figure 17 shows the CPU load as used IFLs for the Cognos BI and the DayTrader guests during the peak and non-peak workload period, with:

- The non-peak workload period: running from 00:00 to 08:00 and 19:00 to 28:00 minutes
- The peak workload period: running from 08:00 to 19:00 minutes.

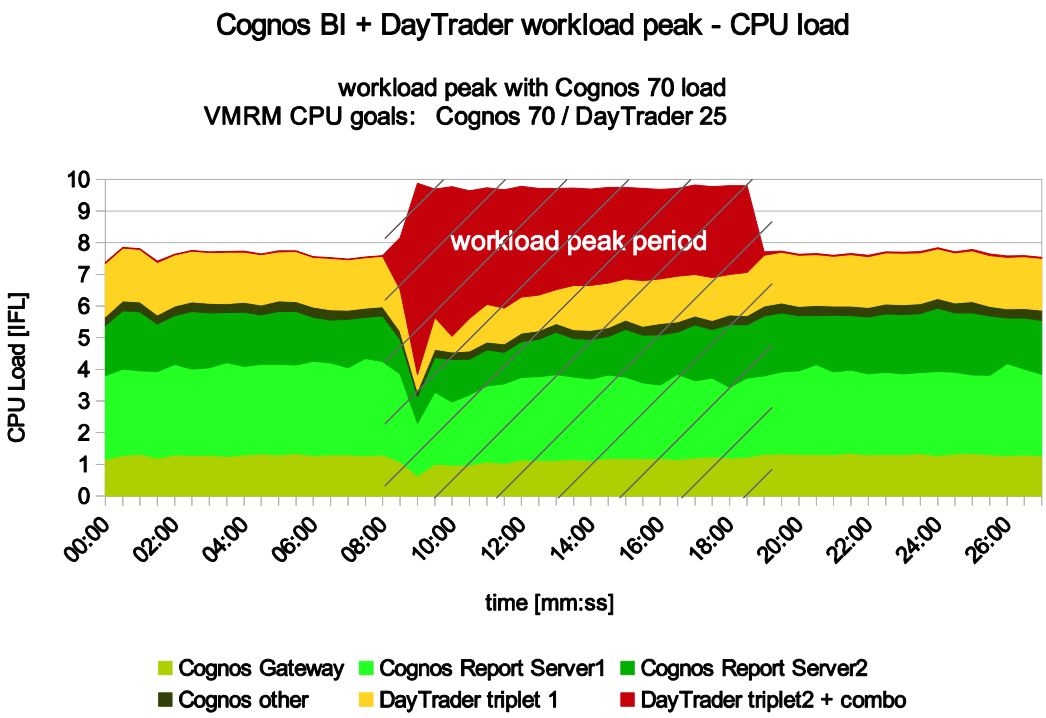

Figure 17: VMRM managing workload peak - Goal 1: CPU loads

Observations: The green areas show the CPU loads for the Cognos BI guests and the yellow area shows the load for the single DayTrader triplet1. These workloads are running all the time.

The red workload peak period in the middle of the figure is composed of DayTrader triplet2 + DayTrader combo.

In the beginning, eight LPAR CPUs are used and the system shows a high CPU utilization, but is still considered as unconstrained. During the workload peak period, all IFLs are used and the system is considered as CPU constrained.

Starting with the peak period, the increased DayTrader workloads are assigned additional CPU resources according to their current relative shares at the expense of the Cognos BI workload. This can be the default shares for the guests, or VMRM already adjusted the shares before. Thirty seconds later, a CP monitor sample interval completes and VMRM starts to adjust the relative shares towards the defined CPU goal. A few sample intervals later, the shares for the Cognos BI guests are upgraded and Cognos BI gets more CPU resources.

Figure 18 shows how VMRM modifies the relative shares for the guests.

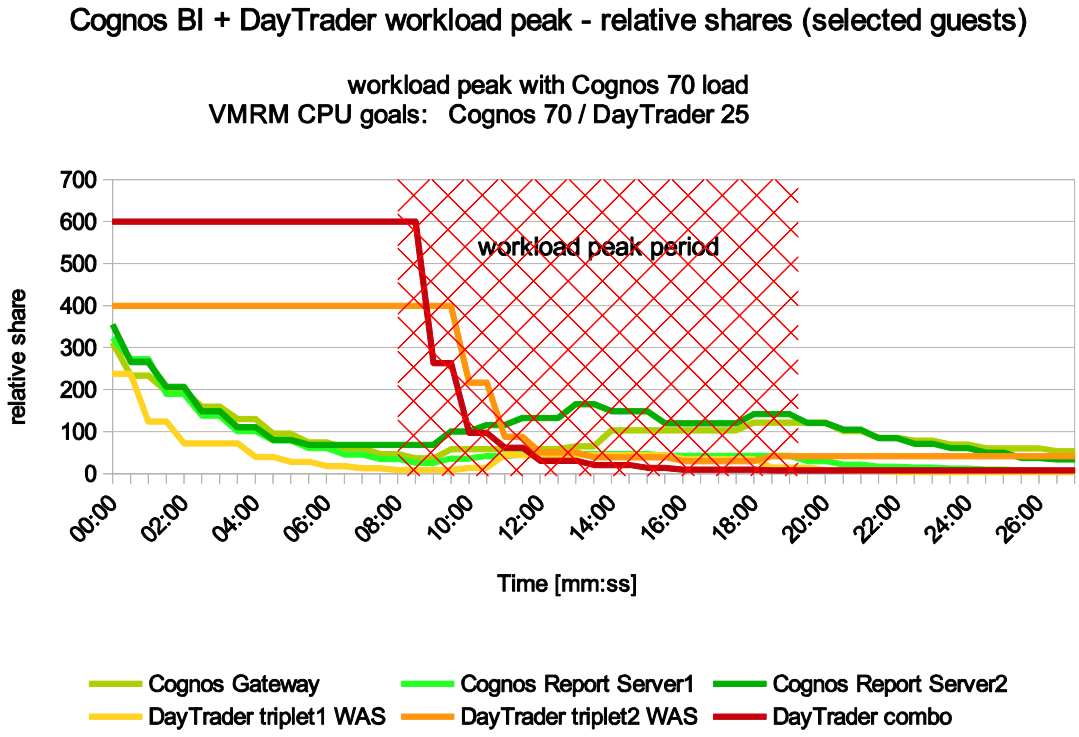

Figure 18: VMRM managing CPU load peak - Goal 1: relative shares

Observations: The green lines show the relative shares for the most important Cognos BI guests (gateway and reports servers). The yellow, orange, and red lines show the relative shares for the DayTrader triplet WAS guests and the combo guest.

VMRM already begins to adjust shares even when the LPAR CPUs are not yet fully constrained. However, this happens apparently only for those guests that already have ongoing workload (Cognos BI and DayTrader triplet1). The relative shares for guests with no load (DayTrader triplet2 and combo) are untouched until the workload peak period starts.

The process of adjusting the shares in accordance to the CPU goals takes some time. The relative shares for all guests are kept at a lower level compared to the fair shares. But in relation, Cognos BI relative shares are higher than the DayTrader shares.

After the workload peak period ends, the Cognos BI guest shares are reduced. There are enough CPU resources available from that point on.

Conclusion: VMRM needs at least one CP monitor sample interval before it starts to adapt relative shares for the guests. With CPU goals below the maximum, the process of adapting relative shares takes longer and the guests with smaller shares still get enough CPU resources. This time period might shorten, when the less prioritized workloads already start with a lower relative share.

VMRM CPU goal 2: Cognos BI gets all (100/1)

The VMRM CPU goals are set as follows for the workloads:

- Cognos BI CPU goal 100 / DayTrader CPU goal 1 -> maximum CPU goal for Cognos BI

Figure 19 shows the CPU load [number of used IFLs] for the Cognos BI and the DayTrader guests during the peak and non-peak workload periods with:

- The non-peak workload period: running from 00:00 to 08:00 and 19:00 to 28:00 minutes
- The peak workload period: running from 08:00 to 19:00 minutes.

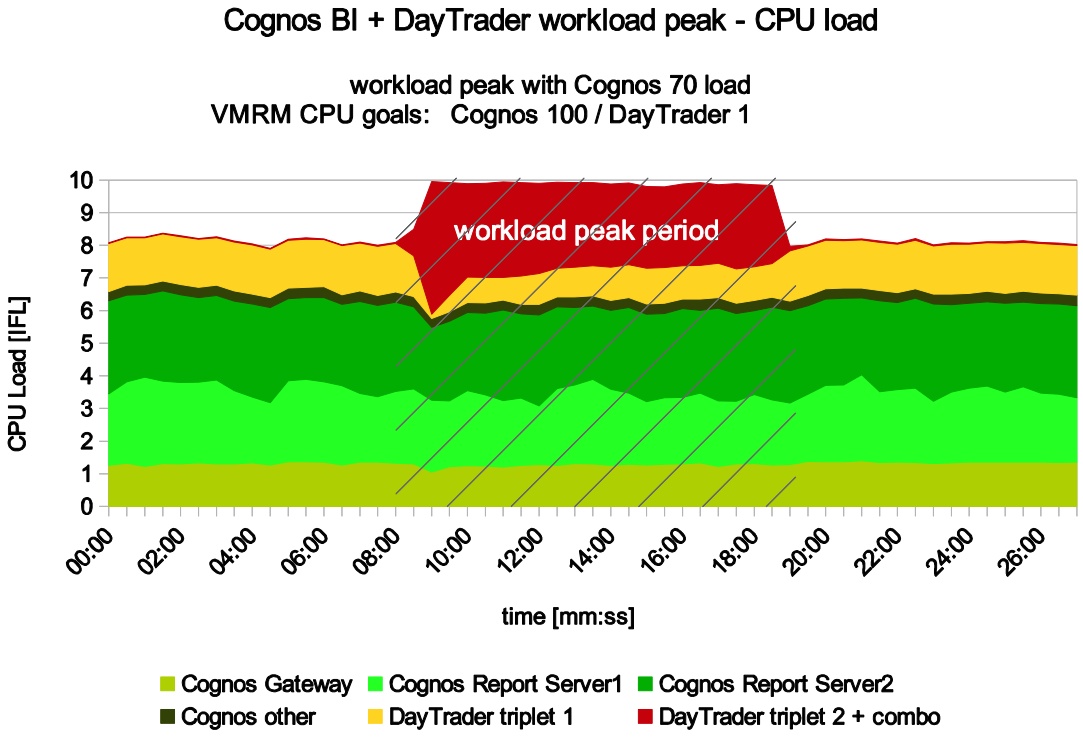

Figure 19: VMRM managing workload peak - Goal 2: CPU loads

Observations: The green areas show the CPU loads for the Cognos BI guests, and the yellow area shows the CPU the load for the single DayTrader triplet1. These workloads are running all the time.

The red workload peak period in the middle of the figure is composed of DayTrader triplet2 plus DayTrader combo.

In the beginning, around eight LPAR CPUs are used and the system shows a high CPU utilization, but is still considered as unconstrained. During the workload peak period, all IFLs are used and the system is considered as CPU constrained.

Starting with the peak period, the increased DayTrader workload gets the CPU resources according to their current relative shares. This can be the default shares for the guests, or VMRM already adjusted the shares before. But this time, the single DayTrader triplet1 has the most severe impact. Momentarily the CPU load for the DayTrader triplet1 goes down to 0.13 IFLs when the load peak started. Thirty seconds later, a CP monitor sample interval completes and VMRM starts to adjust the relative shares towards the defined CPU goal. This time, VMRM adjusts the workloads much faster to their goals. The Cognos BI workload shows a small dip for the CPU load at the beginning of the peak period, and over time, nearly shows no loss of CPU resources.

Figure 20 shows how VMRM modifies the relative shares for the guests.

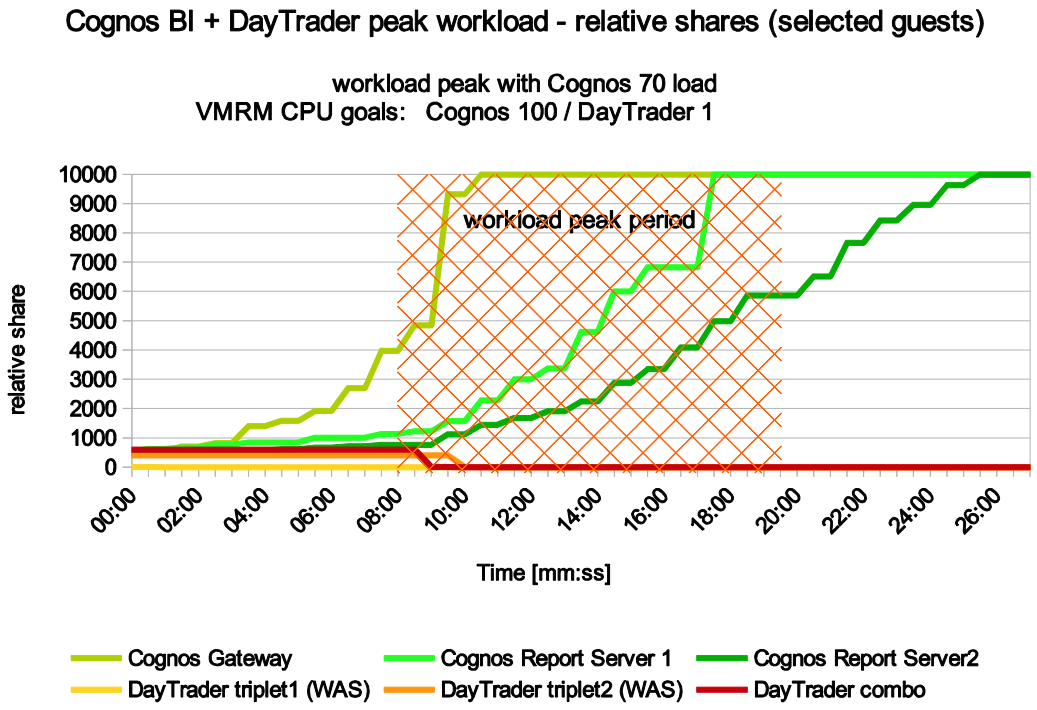

Figure 20: VMRM peak workload - Goal 2: relative shares

Observations: The green lines show the relative shares for the most important Cognos BI guests (gateway and reports servers). The yellow and red lines show the relative shares for the DayTrader triplet WAS guests and the combo guest.

VMRM starts to adjust some shares before the workload peak is issued, even if the LPAR CPUs are not yet fully constrained. With the start of the workload peak period, VMRM begins to tightly increase the shares for the Cognos BI guests and they finally end up at 10000, which is the maximum relative share value.

The CPU shares for the DayTrader guests are kept at a very low level.

After the workload peak period ends, the CPU shares remain more or less unchanged from their latest values. However, there are enough CPU resources available for all workloads from that point on.

Conclusion: VMRM needs at least one CP monitor sample interval before it starts to adapt relative shares for the guests. With this maximum CPU goal for the Cognos BI workload, VMRM controls the relative shares in a completely different manner compared to more moderate CPU goals.

Note: Maximum CPU goals can lead to situations where competing workloads with very low goals get nearly no CPU resources.

References

Cognos Business Intelligence for Linux on System z:
http://www.ibm.com/software/products/en/cognbusiinteforlinuonsystz

How to Improve Performance with PAV, SC33-8414

Linux on System z - Device Drivers, Features, and Commands SC33-8411-23:
http://www.ibm.com/developerworks/linux/linux390/development_documentation.html

Details about Linux for System z network tuning:
http://www.ibm.com/developerworks/linux/linux390/perf/tuning_networking.html

IBM Knowledge Center for Cognos Business Intelligence 10.2.1 documentation:
http://www.ibm.com/support/knowledgecenter/SSEP7J_10.2.1

IBM Knowledge Center for WebSphere Application Server, Network Deployment, V8.5.5 documentation:
http://www.ibm.com/support/knowledgecenter/SSAW57_8.5.5

IBM Knowledge Center for DB2 10.5 for Linux, UNIX and Windows documentation:
http://www.ibm.com/support/knowledgecenter/SSEPGG_10.5.0

Input/Output Configuration Program User's Guide, SB10-7037-11:
http://www.ibm.com/servers/resourcelink

IBM Technical White Paper:
Oracle Database on Linux on System z - Disk I/O Connectivity Study

z/VM V6R3 Performance, SC24-6208-04

z/VM V6 R3.0 Information Center:
http://pic.dhe.ibm.com/infocenter/zvm/v6r3/index.jsp

Customization and Tuning > z/VM: Performance > Tuning z/VM Performance > VMRM Tuning Parameters

Article on IBM developerWorks:
IBM Business Analytics Proven Practices: IBM Cognos BI Dispatcher Routing Explained

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.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing 2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan

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

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

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

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

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Software Interoperability Coordinator, Department 49XA
3605 Highway 52 N
Rochester, MN 55901
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

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.

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

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

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.

Terms and conditions

Permissions for the use of these publications is granted subject to the following terms and conditions.

Personal Use: You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these publications, or any portion thereof, without the express consent of the manufacturer.

Commercial Use: You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of the manufacturer.

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any data, software or other intellectual property contained therein.

The manufacturer reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by the manufacturer, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

THE MANUFACTURER MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THESE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

IBM Systems and Technology Group
Route 100
Somers, New York 10589
U.S.A.
Produced in the United States of America,
10/2014

ZSW03268-USEN-00