



Effectively running IBM Cognos BI for Linux on z Systems in a z/VM environment

Mario Held (mario.held@de.ibm.com)

Linux on z Systems Performance Analyst

IBM Corporation



#SHAREorg



SHARE is an independent volunteer-run information technology association that provides **education, professional networking and industry influence.**

Copyright (c) 2015 by SHARE Inc. Except where otherwise noted, this work is licensed under <http://creativecommons.org/licenses/by-nc-sa/3.0/>



Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries.

A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml

The following are trademarks or registered trademarks of other companies.

- Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
- SUSE is a registered trademark of Novell, Inc. in the United States and other countries.
- Red Hat, Red Hat Enterprise Linux, the Shadowman logo and JBoss are registered trademarks of Red Hat, Inc. in the U.S. and other countries.
- Oracle and Java are registered trademarks of Oracle and/or its affiliates in the United States, other countries, or both.

Other product and service names might be trademarks of IBM or other companies.

Agenda

- **Cognos Business Intelligence (BI) - System Under Test overview**
- Cognos BI 10.2 tuning and setup
- z/VM setup and the Resource Manager (VMRM)
- Performance study

Acknowledgements

I wish to acknowledge the help provided by Thomas Weber, performance specialist at the end-to-end performance team in the IBM Boeblingen Lab, Germany

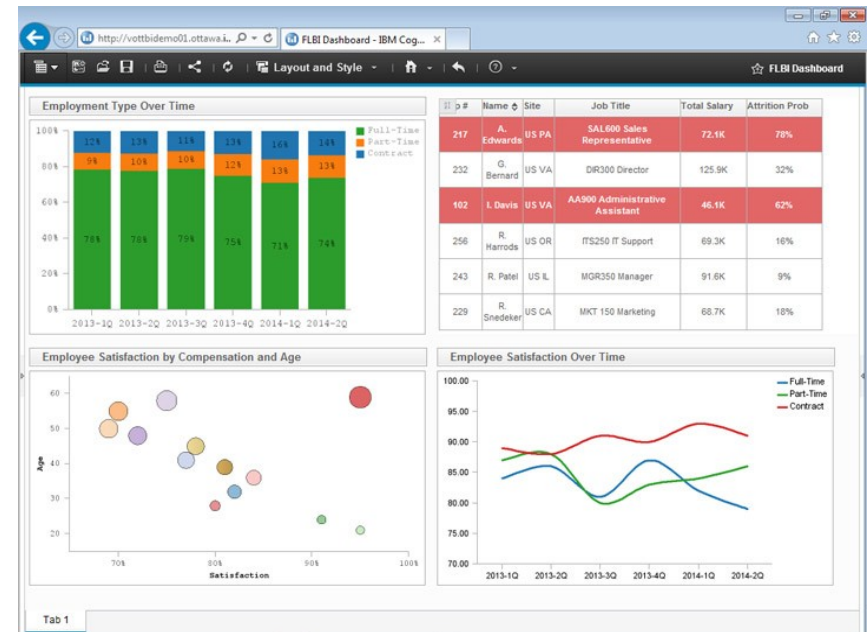
Cognos Business Intelligence for Linux on z Systems



“Know the past,
understand the present,
shape the future.”

IBM Cognos Business Intelligence (BI):

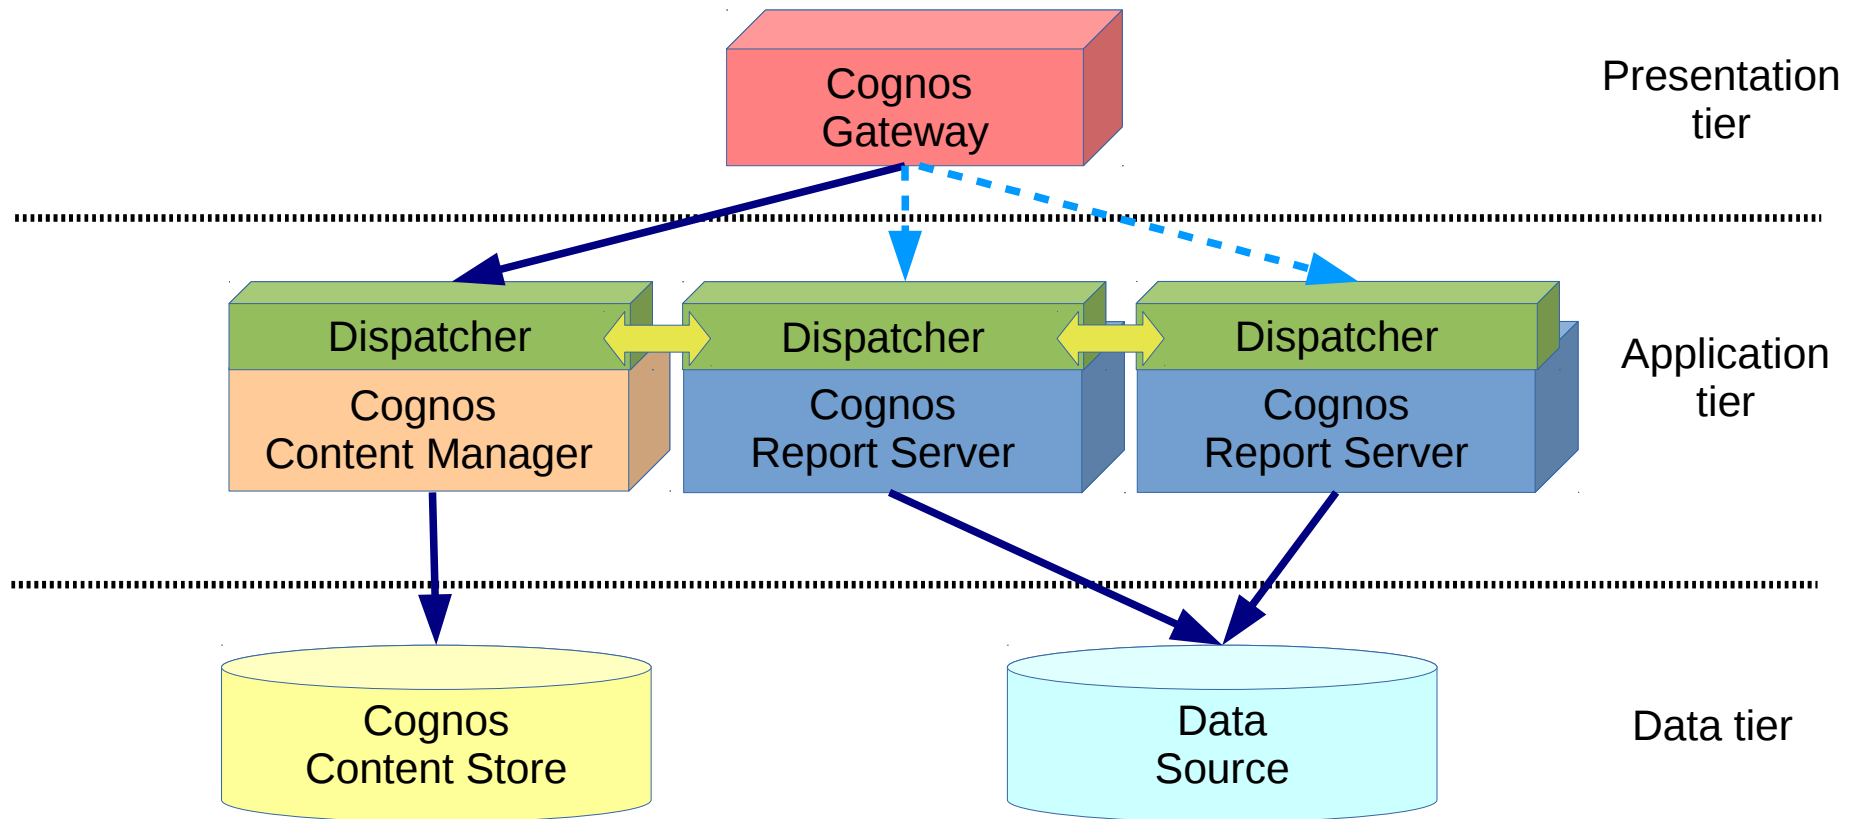
- Web-based application suite
- Software suite with typical components (web portal, content manager, report server)
- Provides techniques and tools to turn raw data into meaningful information
- BI techniques can handle large amounts of unstructured data
- BI allows easy interpretation of large volumes of data
- Helps to identify and develop new strategic business opportunities



Cognos Business Intelligence for Linux on z Systems



Cognos BI Server as a 3-tier model



Linux on z Systems end-to-end project



Involved Teams

- Linux on z Systems end-to-end performance team
- IBM z Systems Cognos BI performance team

Setup

- z/VM virtualized environment
- z/VM Resource Manager setup (VMRM)
- Distributed installation of Cognos BI Server core components
 - Cognos Gateway
 - Cognos Content Manager
 - Cognos Report Servers

Performance study

- Cognos BI workload should get enough CPU resources when constraint
- Concurrent workloads (DayTrader)
- Static and dynamic tuning of z/VM virtual machine CPU resources
- Technical White Paper

Cognos BI Server Components (1)

Cognos BI Gateway

- Web communication is done through a gateway
- Installed on one more webservers
- Webserver extension (CGI script or Apache module)
e.g. for IBM HTTP Server (IHS)
- Receives client requests and passes them to the first registered Cognos dispatcher
- Static connection to a dispatcher
- Does no work balancing or routing of requests

**z/VM Linux
guest**

4 CPUs
4 GiB memory

**Cognos BI
Gateway
10.2.1**

IBM HTTP
Server 8.5.5

SLES 11 SP3



Cognos BI Server Components (2)

Cognos BI Content Manager

- Corner stone in every Cognos BI installation
- Ensures integrity of the Content Store database
- Requires an application server (e.g. WAS)
- Only one active at a time
- Multiple can be configured (failover)

Cognos BI Content Store

- Relational database
- Content Manager maintains the Content Store
- Represents the Cognos BI server meta data repository (e.g. report definitions, security roles, data models,...)

**z/VM Linux
guest**

4 CPUs
8 GiB memory

**Cognos BI
Content
Manager
10.2.1**

IBM Websphere
Application
Server v8.5.5

SLES 11 SP3

**z/VM Linux
guest**

2 CPUs
4 GiB memory

**Cognos BI
Content Store**

IBM DB2 10.5

SLES 11 SP3



z/VM®

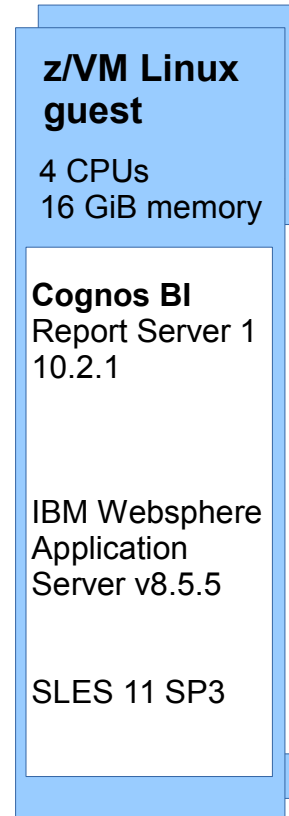
Cognos BI Server Components (3)

Cognos BI Report Server

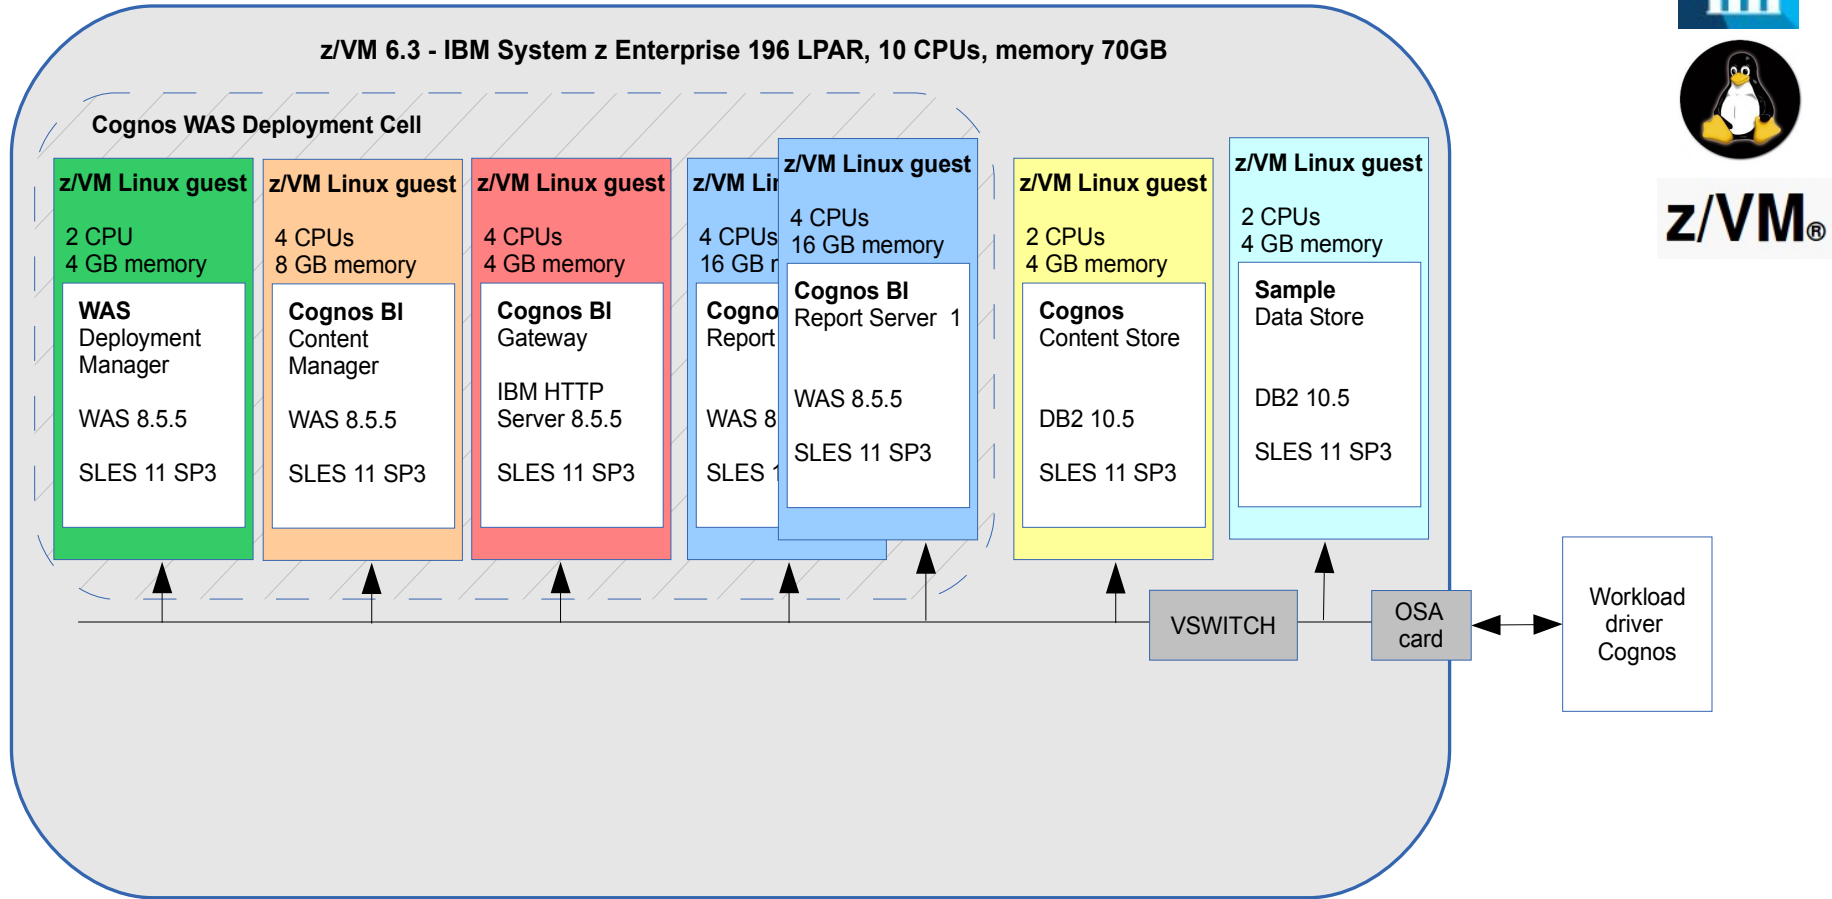
- Runs requests forwarded by the Cognos BI Gateway
 - Reports
 - Analysis
 - DB queries
- Load balancing over the Cognos BI Dispatcher
- Requires an application server (e.g. WAS)
- Multiple instances of Cognos BI Report Servers possible

Cognos BI Data Store

- Relational database
- Holds sample database data for the test

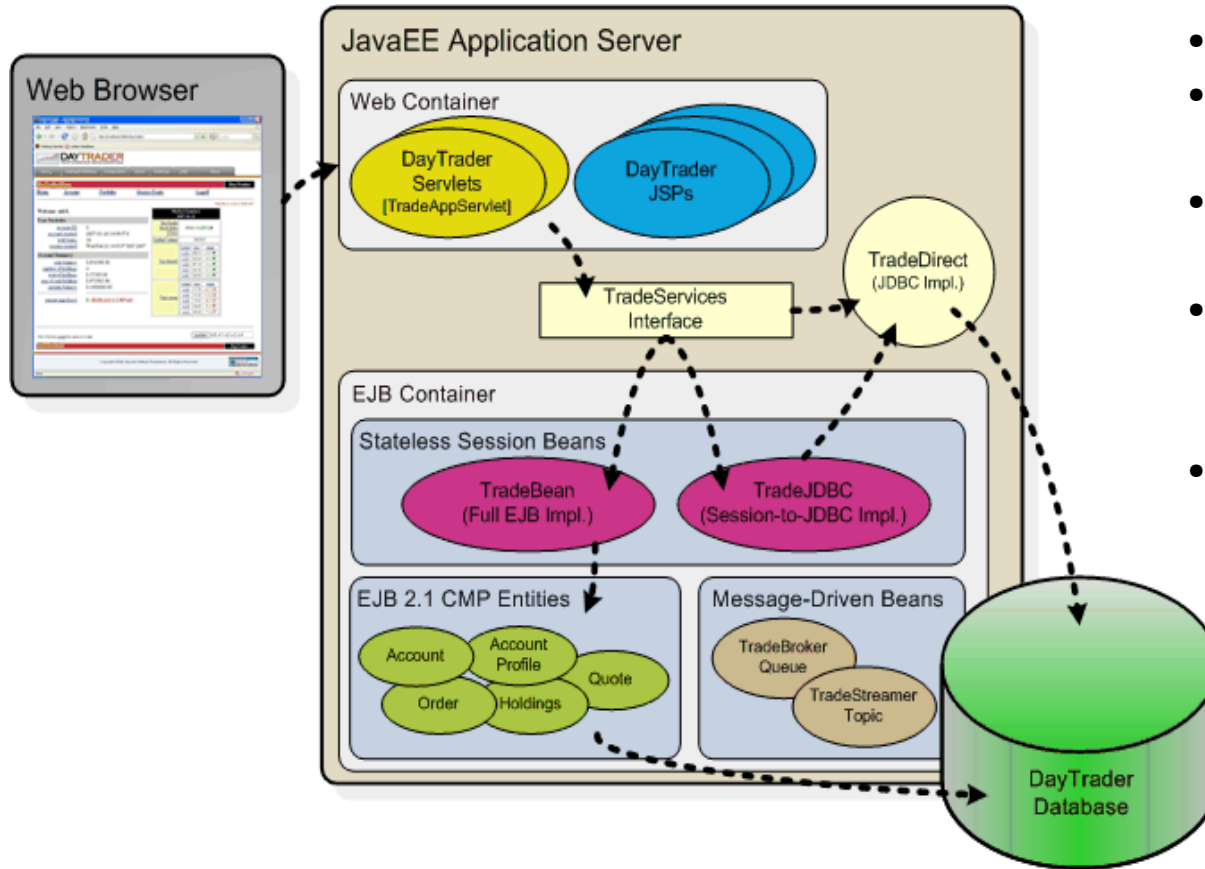


System Under Test (SUT) – Cognos BI



Websphere Application Server Network Deployment:
Cognos BI WAS Deployment Cell: WAS Deployment Manager + Cognos BI WAS Nodes

Concurrent Workloads (DayTrader) - (1)



- Concurrent workload
- Open Source benchmark application
- Emulates an online stock trading system
- End-to-end Java EE (Enterprise Edition) web application
- IBM WAS is a Java EE application server

<http://geronimo.apache.org/GMOxDOC30/daytrader-a-more-complex-application.html>

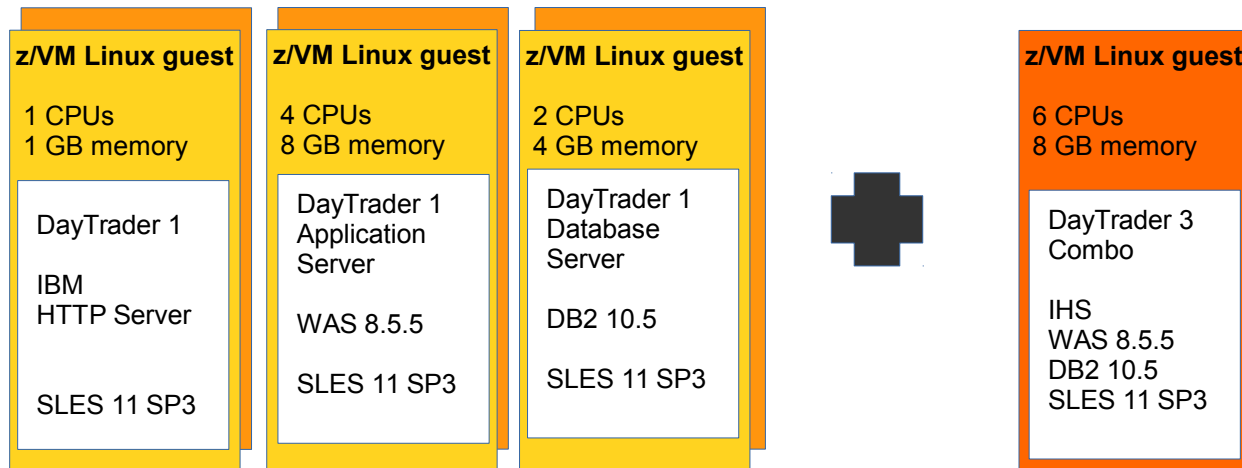
Concurrent Workloads (DayTrader) - (2)



z/VM®

3x DayTrader installations as concurrent workload

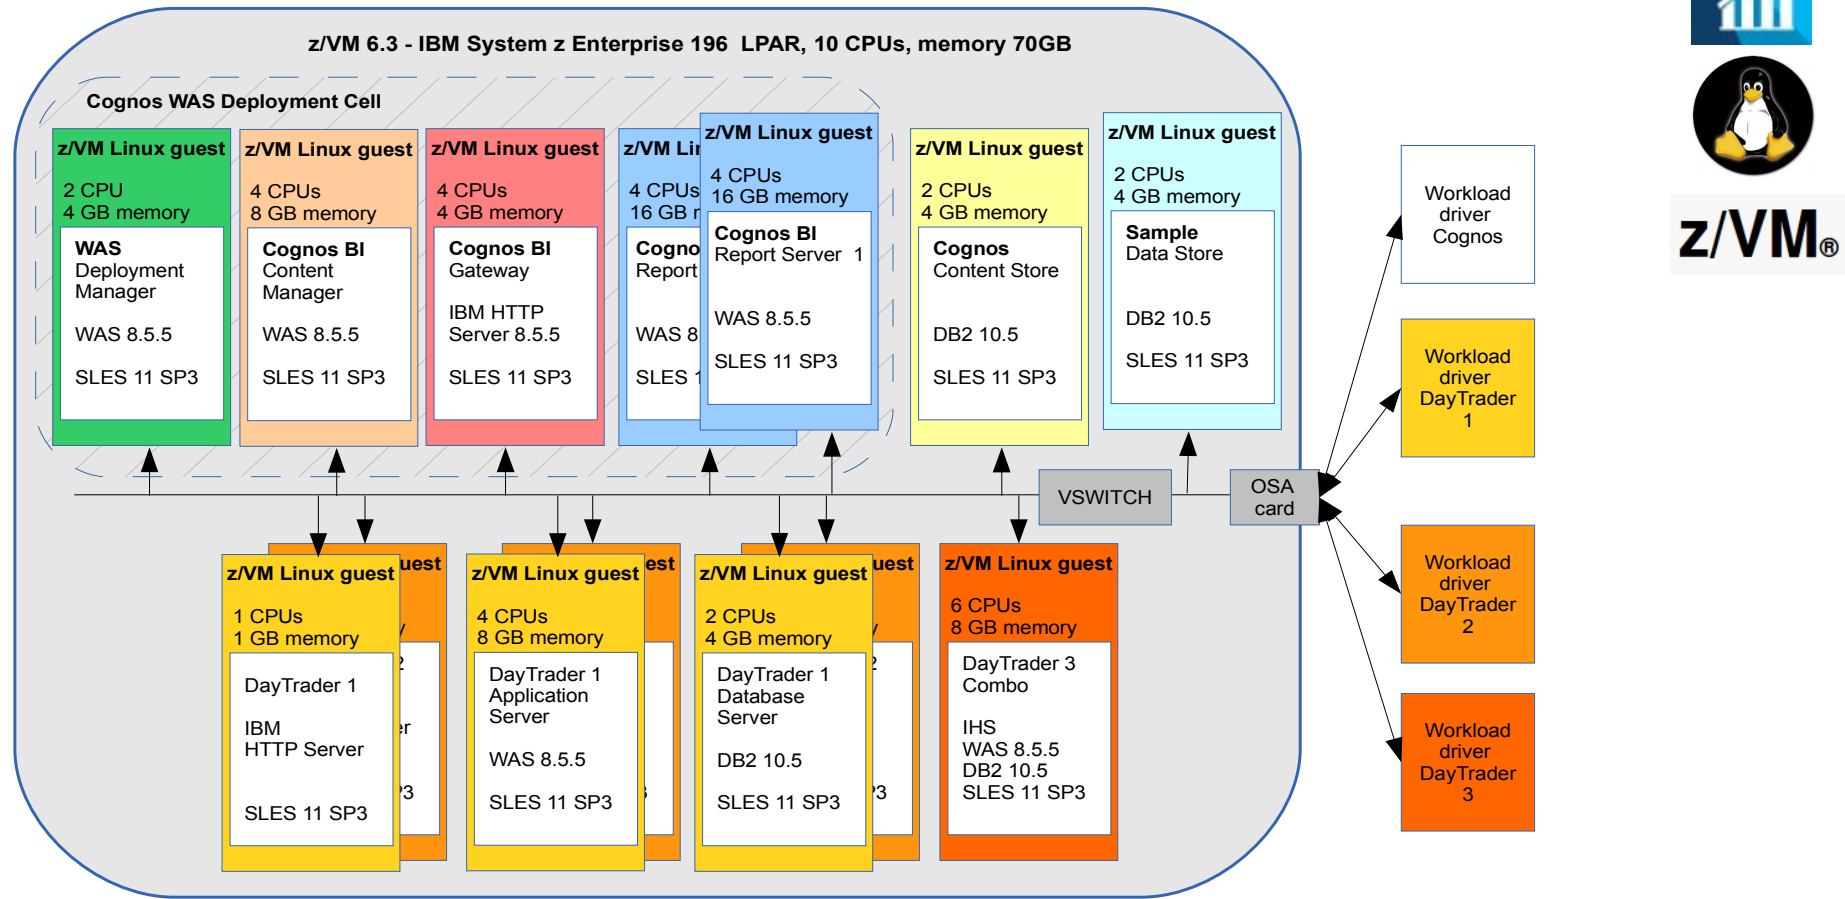
- **2x DayTrader triplets (multiple virtual machines)**
- IBM HTTP server + AppServer + DB Server (each two of them in a single machine)



- **1x DayTrader combo (single virtual machine)**
- IBM HTTP server + AppServer + DB Server in a single machine

<http://geronimo.apache.org/GMOxDOC30/daytrader-a-more-complex-application.html>

Complete System Under Test (SUT)



virtualized SUT running under z/VM:

- 14 Linux guests
- Ratio of CPU overcommitment **4.2 : 1 (42:10)**
- Ratio of Memory overcommitment **1.3 : 1 (90GiB:70GiB)**

Agenda

- Cognos Business Intelligence (BI) - System Under Test overview
- **Cognos BI 10.2 tuning and setup**
- z/VM setup and the Resource Manager (VMRM)
- Performance study

Cognos BI Gateway tuning (1)



- Webserver extension; e.g. for IBM HTTP Server (IHS)
- Gateway knows the location of the primary Cognos BI Dispatcher service

Basic Gateway processing when receiving a request

- Encrypts password to ensure security
- Extracts information needed to submit the request to the Dispatcher
- Attaches environment variables for the web server
- Proper handling of Cognos BI namespaces
- Passes requests to the Dispatcher for processing

Usage of Apache modules on IBM HTTP Server (IHS)

- Replaces the default CGI gateway by an Apache module (apache_mod)
- Requires a change in the IHS configuration file (httpd.conf)
- Add Apache module at the end of the IHS load module list
- Provides enhanced performance and throughput

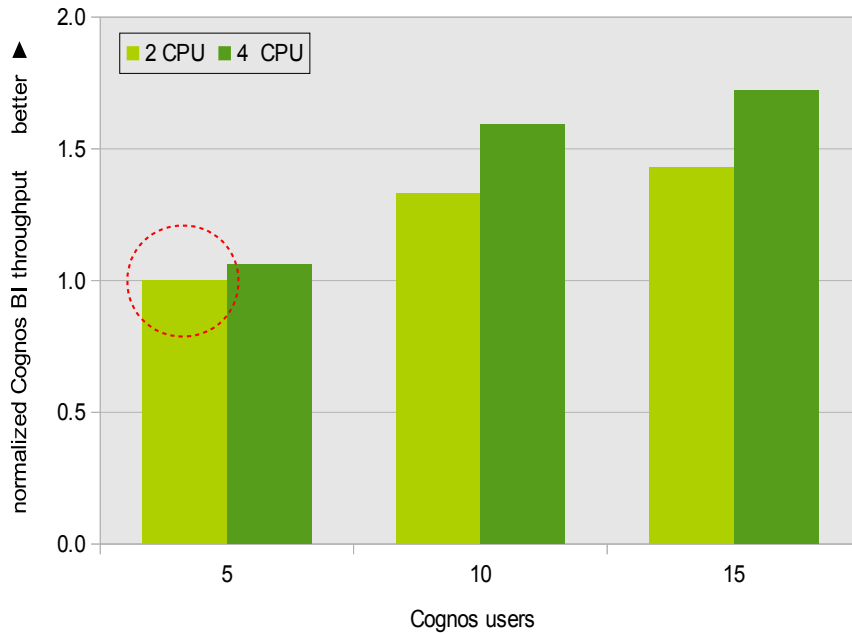
```
LoadModule cognos_module <cognos10_location>/cgi-bin/mod2_2_cognos.so
```

Cognos BI Gateway tuning (2)

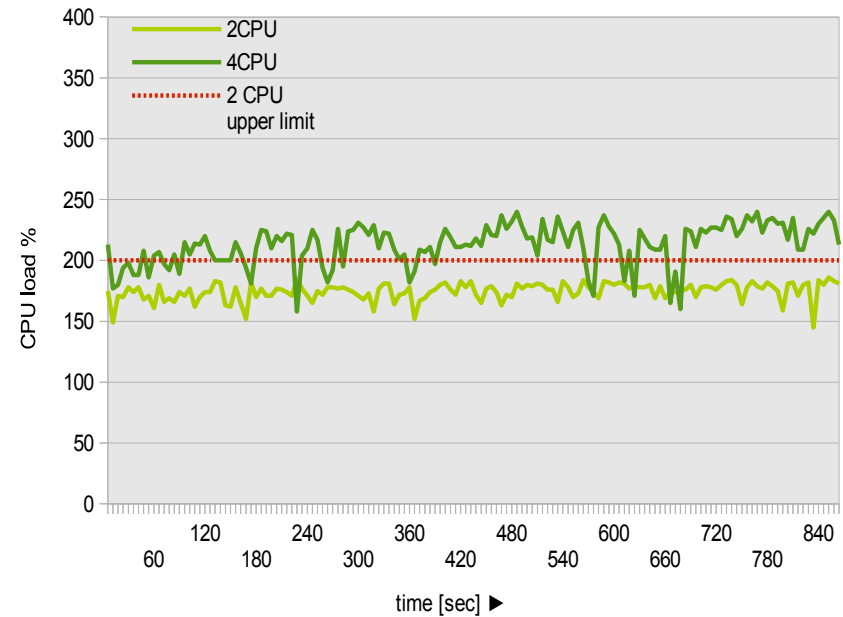


Provide enough (virtual) CPUs!

Cognos BI Gateway - CPUs



Cognos BI Gateway - CPU load for 10 users



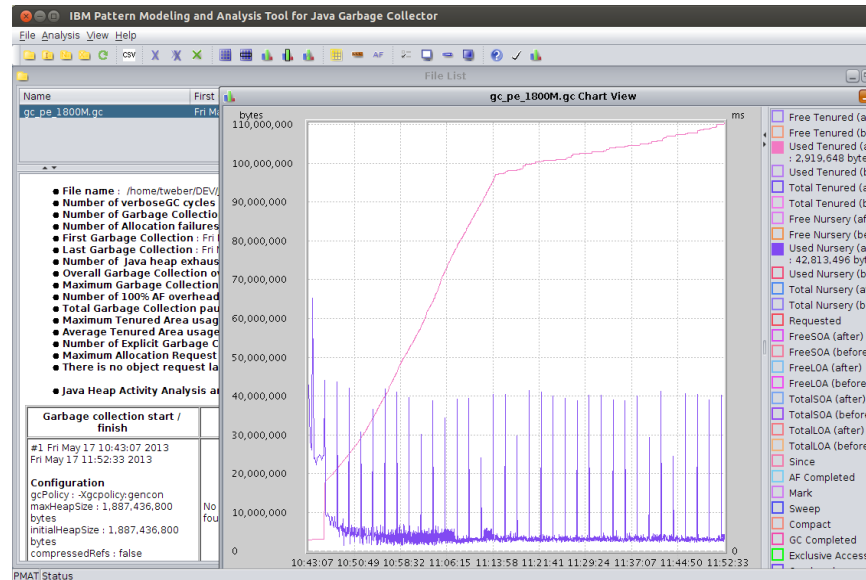
- Same workload level for 2 and 4 CPU measurement series
- Note that the CPU load is not fully bound in the 2 CPU case
- Transaction throughput increases up to 17% by using 4 CPUs in the 10 Cognos users setup

Cognos BI Report Server Java tuning (1)



Websphere Application Server Java Virtual Machine heap size

- Cognos BI Report Server runs as a WAS application
- WAS JVM heap was increased to initial / maximum 1024MB / 2048MB
- A larger JVM heap size often improves the performance
- Monitor the JVM heap over time to find a reasonable heap size
- Consider the memory size of the z/VM virtual machine and other running services



- Screenshot: IBM Pattern Modeling and Analysis Tool for Java Garbage Collector

Cognos BI Report Server Java tuning (2)



WebSphere Application Server web container thread pool

- Web container handle Java server-side code requests for servlets, JavaServer Pages (JSP) etc.
- Initial values for Maximum Size suitable for simple web applications
- Adapted values improve scalability for more complex applications like Cognos BI
- Maximum number of WebContainer threads was increased to 500

WAS administration console path:

Servers -> Application Servers

-> <servername>

-> Thread Pools -> WebContainer

General Properties

* Name

WebContainer

Description

* Minimum Size

50

threads

* Maximum Size

500

threads

* Thread inactivity timeout

60000

milliseconds

Allow thread allocation beyond maximum thread size

Cognos BI Report Server report service tuning (1)




Number of Cognos BI Report Server service processes

- Dispatcher assigns a requests to a report service process
- Processes that the report service can start is limited (BIBus processes)
- Tuning depending on workload patterns
- Maximum number of parallel report service processes is configurable (long running reports)
- High and low affinity is configurable (many short running reports)

Rule of thumb when choosing the number of processes

- Configure based on the available CPUs
- Report processing is mainly a CPU-bound process
- Single report service processes can grow over 1 GiB
- Consider also the memory footprint of the Report Server

Allow twice as much the number of processes for the report service for the SUT (compared to the default):

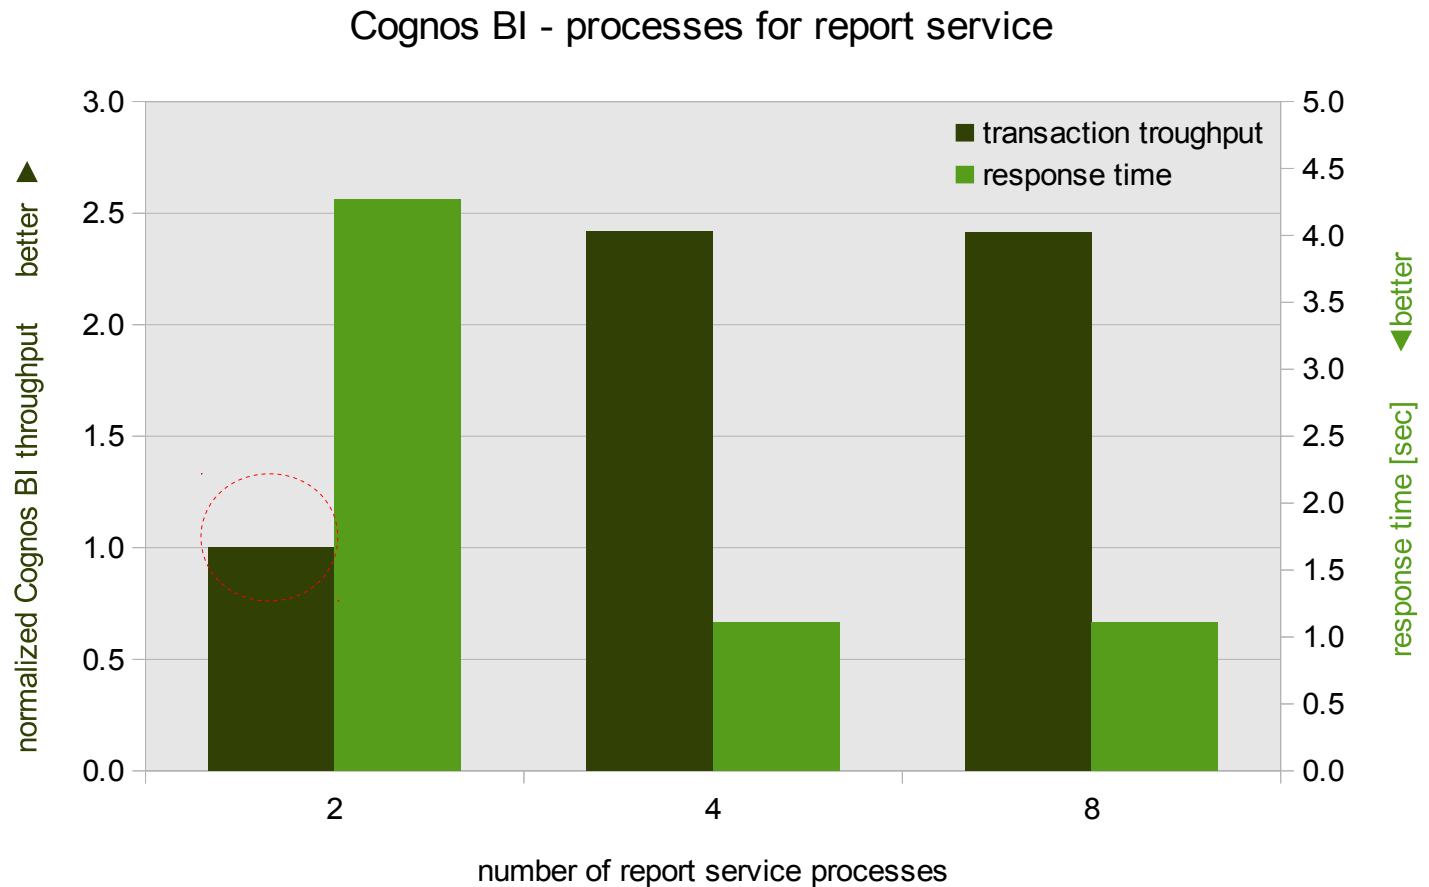
<input type="checkbox"/>	 Tuning	Maximum number of processes for the report service during peak period	<input type="text" value="4"/>	No
--------------------------	--	---	--------------------------------	----

Cognos BI Report Server report service tuning (2)



Scaling the number of Cognos BI report service processes

- Cognos BI Report Server with 4 CPUs showed best throughput / resource ratio




Cognos BI Report Server report service tuning (3)



Memory footprint of 4 active report service processes

Memory consumption of 4 report service processes over time





Cognos BI Report Server Query Service



Query Service – Java heap size

- Initial and maximum heap size was increased to 4096 MiB

<input type="checkbox"/>	 Tuning	Initial JVM heap size for the query service (MB) (Requires QueryService restart)	4096	No
<input type="checkbox"/>	 Tuning	JVM heap size limit for the query service (MB) (Requires QueryService restart)	4096	No

Overall Report Server memory footprint

- Cognos BI Report Server tuning can result in a large memory footprint
- Consider the following memory-consuming components
 - WAS JVM heap size (max. 2 GiB)
 - Query Service JVM heap size (4 GiB)
 - Report service processes can get bigger than 1 GiB per process

==> **Cognos BI Report Server can quickly get a large memory footprint**

==> Report Servers for the SUT were sized with 16GiB each

Agenda

- Cognos Business Intelligence (BI) - System Under Test overview
- Cognos BI 10.2 tuning and setup
- **z/VM setup and the Resource Manager (VMRM)**
- Performance study

Control virtual machine CPU capacity (1)



z/VM introduces shares to control the CPU for a virtual machine

- Share concept applies when the z/VM hypervisor is CPU constraint
- **ABSOLUTE share**
 - Allocates an absolute portion of all z/VM processors for a virtual machine
 - Guarantees a certain percentage of processing time
 - No usage of ABSOLUTE shares for this project
- **RELATIVE share**
 - Allocates a relative portion of the z/VM processors for a virtual machine
(less than any allocations for ABSOLUTE shares)
 - RELATIVE share is an integer number between **1** to **10000**
(larger number means higher share)

Control virtual machine CPU capacity (2)



Definition of the fair share setup

- Used for all non-VMRM tests in this project
- Default RELATIVE share for a virtual machine is **100**
- Fair share setup == assigns a RELATIVE share of **100** per virtual CPU
- Ensures that each virtual CPU has the same weight within z/VM

Example

- Virtual machine VM1 has **2** virtual CPUs: RELATIVE fair share is **200**
- Virtual machine VM2 has **4** virtual CPUs: RELATIVE fair share is **400**

 VM2 can get twice as much CPU time for its 4 CPUs than VM1

z/VM Resource Manager VMRM (1)



VMRM can dynamically tune a z/VM virtual machine CPU shares

- Virtual machines can be members of workload groups
- Workload groups will be managed according to defined goals
- VMRM automatically adjusts performance parameters to achieve the goals
- Only effective in constrained environments

- CP monitor data is used to obtain a virtual machine's resource consumption
- Requires a service virtual machine VMRMSVM

- Other tunable resources are memory and DASD I/O

Note

- VMRM Cooperative Memory Management (CMM) is not used in this project
- VMRM velocity targets for DASDs are not used in this project

z/VM Resource Manager VMRM (2)



Service virtual machine VMRMSVM

- Starts VMRM by calling the IRMSERV EXEC program
- Reads user supplied configuration file (default: VMRM CONFIG A)
 - Definition of workloads, goals, priorities
- VMRM adjusts virtual machine tuning parameters
 - Using CP Command 'SET SHARE' for CPU shares
 - Done for “eligible” virtual machines in a defined workload group
- VMRM interaction with the CP monitor
 - Starts sample monitoring if not already active
 - Default: 60 sec sample monitoring interval
- VMRM screen in the z/VM Performance Toolkit (FCX241)
 - See backup

z/VM Resource Manager VMRM (3)



Sample VMRM configuration file with CPU velocity goals

```
ADMIN MSGUSER VMRMADMN
GOAL COGCPU VELOCITY CPU 70
GOAL DAYCPU VELOCITY CPU 25
WORKLOAD DAYTRADER USER TRDUSR*
MANAGE DAYTRADER GOAL DAYCPU IMPORTANCE 1
WORKLOAD COGNOS USER COGUSR*
MANAGE COGNOS GOAL COGCPU IMPORTANCE 10
```

ADMIN

- identifies a user to receive VMRM messages

GOAL

- used CPU velocity goals (1-100)

WORKLOAD

- describes a workload by userid, account id, or acigroup

MANAGE

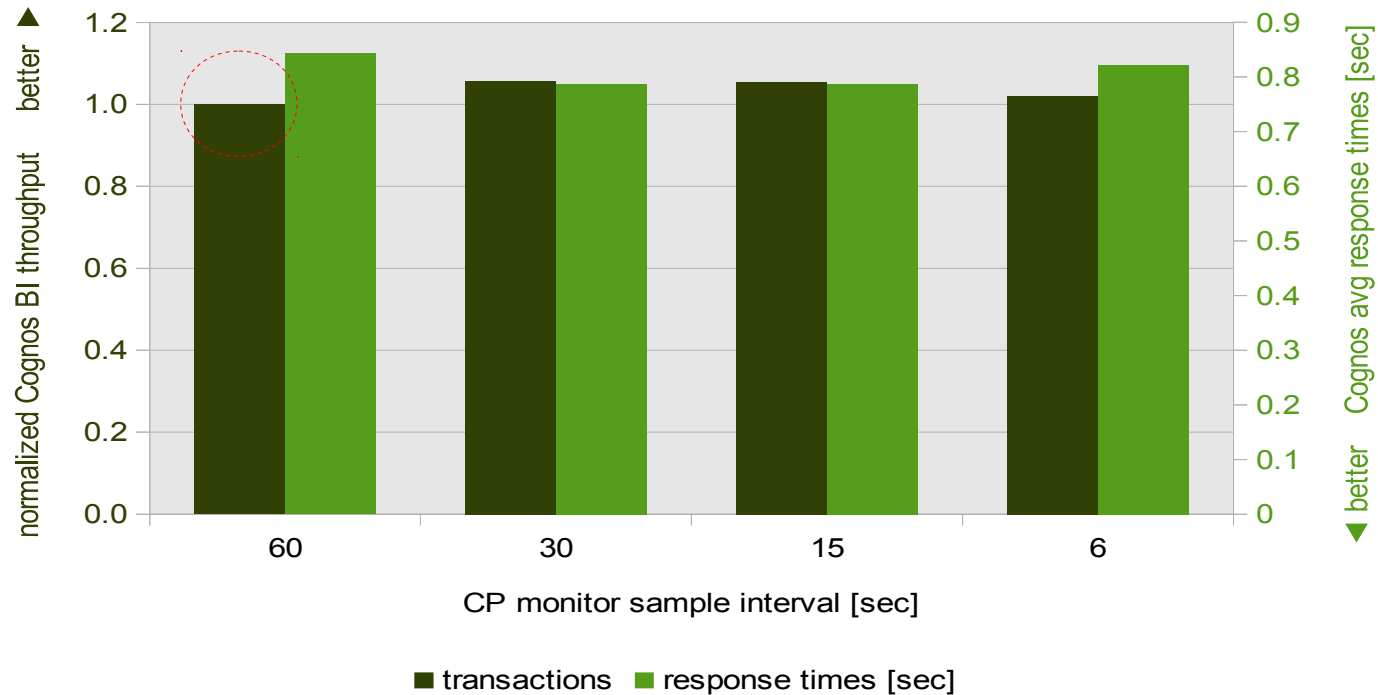
- associates a workload with a goal and assigns an importance value 1(lowest) – 10(highest)

z/VM Resource Manager VMRM (4)



Influence of the CP monitor sample interval

VMRM parameter - CP monitor sample interval



- Measurement series with different sample intervals (default 60 sec)
- Slightly better results with smaller intervals 30 and 15 sec
- Interval of 30 sec has been chosen for the complete VMRM measurement series

Agenda

- Cognos Business Intelligence (BI) - System Under Test overview
- Cognos BI 10.2 tuning and setup
- z/VM setup and the Resource Manager (VMRM)
- **Performance study**

Workload for Performance study (1) – Cognos BI



- Custom benchmark application running on an external x86 machine
- Workload is forming a typical Cognos BI transaction mix
- Multiple users can run the transaction mix

Cognos BI scenario	Workload share
interactive reporting and ad hoc analysis use of Dynamic Query Mode (DQM)	30%
chart and report creation using the new charting engine	20%
dashboard activity users open multiple pages in a webbrowser 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%

Workload for Performance study (2) – Cognos BI



- Different Cognos BI CPU load levels used within the study
- Load varies with number of active users
- Multiple users can run the transaction mix in parallel

Cognos BI CPU load level

Cognos70

5 users
approx 70% IFLs used
default load level for Cognos BI

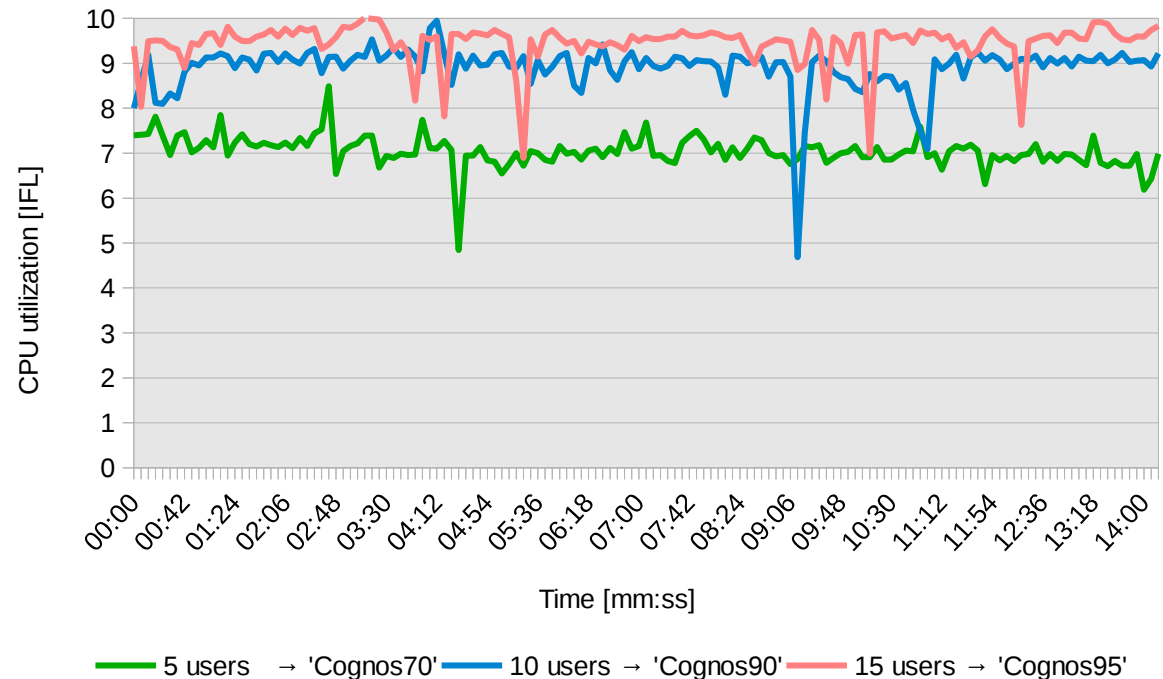
Cognos90

10 users
approx 90% IFLs used
used for some selected scenarios

Cognos95

15 users
approx 95% IFLs used
used for some selected scenarios

Cognos BI Server - LPAR CPU Load (10 IFLs)



Workload for Performance study (3) – DayTrader

- Different DayTrader CPU load levels used within the study
- Load varies with number of acting users
- Load levels are for 3 parallel running DayTrader Apps (2 triplets + 1 combo)

DayTrader CPU load level	Description
DayTrader25	5 users each approx 25% of 10 IFLs used low load level
DayTrader50	125 users each approx 50% of 10 IFLs used medium load level
DayTrader75	190 users each approx 75% of 10 IFLs used high load level

Performance study – Fair share setup

- All virtual CPUs have the same relative share (100 per CPU)
- Mixed workload (Cognos BI + DayTrader workload in parallel)
- Single DayTrader combo has the highest relative share 600
- DayTrader triplets relative shares 100/400/200
- Cognos BI gateway, report servers relative shares 400/400/400

Workload CPU load level

Cognos70 + DayTrader25

high workload with IFL load close to 100%; some peaks above

Cognos70 + DayTrader50


higher workload with IFL load above 100% (full constrained)

Cognos70 + DayTrader75

highest workload with IFL load above 100% (full constrained)

Cognos BI / DayTrader transaction throughput

scaling DayTrader CPU load level at Cognos BI CPU load level 70
100% = single workload for Cognos 70 and for DayTrader 25



■ Cognos 70 ■ DayTrader1 triplet ■ DayTrader2 triplet ■ DayTrader3 combo

Performance study – VMRM setup (1)

- Relative CPU shares are modified by VMRM according to the workload goals
- Mixed workload (Cognos BI + DayTrader) in two different workload groups
- Cognos BI workload has a **high** CPU velocity goal
- DayTrader workload has a **low** CPU velocity goal

Workload CPU load level

Cognos70 + DayTrader25

high workload with IFL load close to 100%; some peaks above

Cognos70 + DayTrader50

higher workload with IFL load above 100% (full constrained)

Cognos70 + DayTrader75

highest workload with IFL load above 100% (full constrained)

VMRM settings used for this study

Fixed (determined by tests)

- CP monitor sample interval: 30 seconds
- Goal importance: 10 (high) for Cognos BI
1 (low) for DayTrader


Varied

- CPU velocity goal

CPU velocity goal	for Cognos BI
Cogn:60 DayTr:25	moderate high
Cogn:70 DayTr:25	high
Cogn:80 DayTr:25	high
Cogn:100 DayTr:1	maximum possible

Performance study – VMRM setup (2)

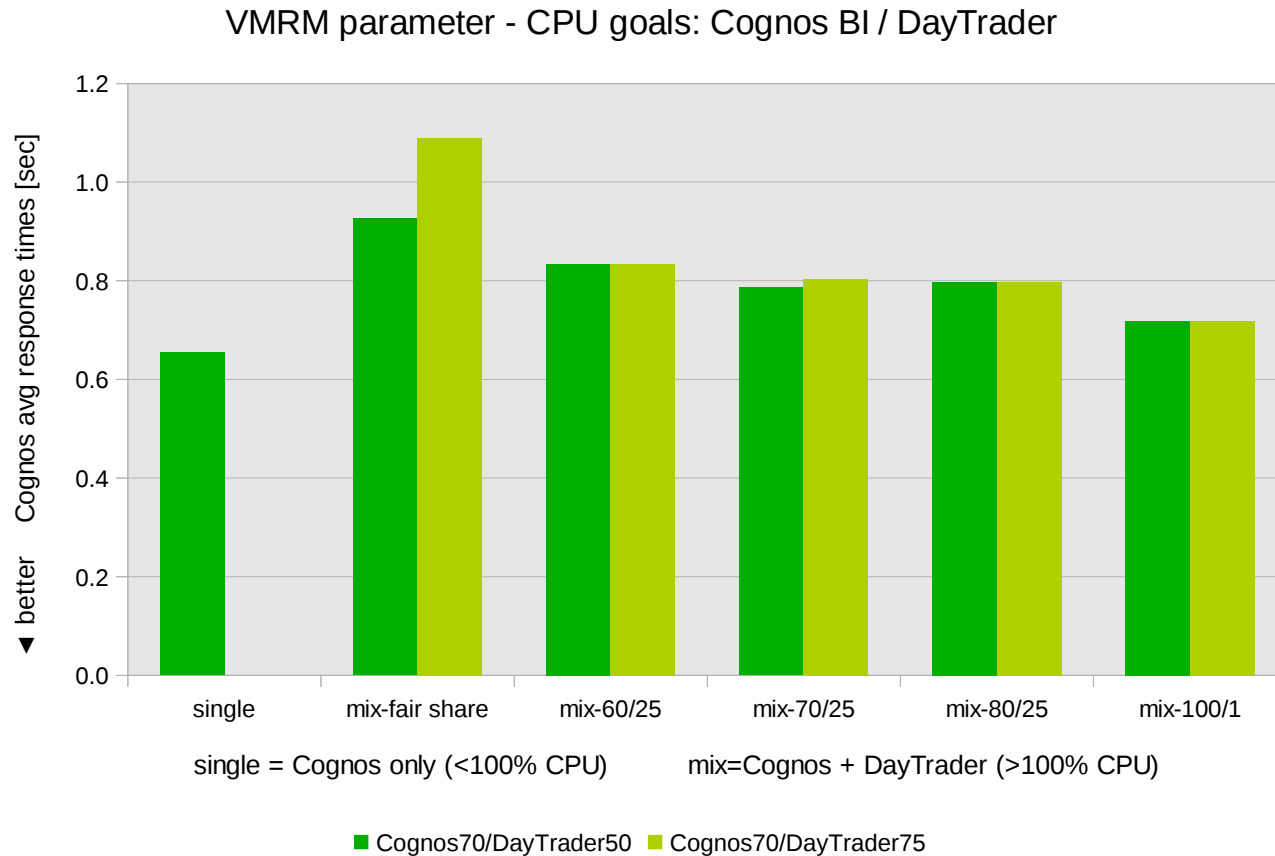
Cognos BI throughput for different VMRM CPU velocity goals



- VMRM CPU velocity goals guarantee a certain throughput level

Performance study – VMRM setup (3)

Cognos BI response times for different VMRM CPU velocity goals



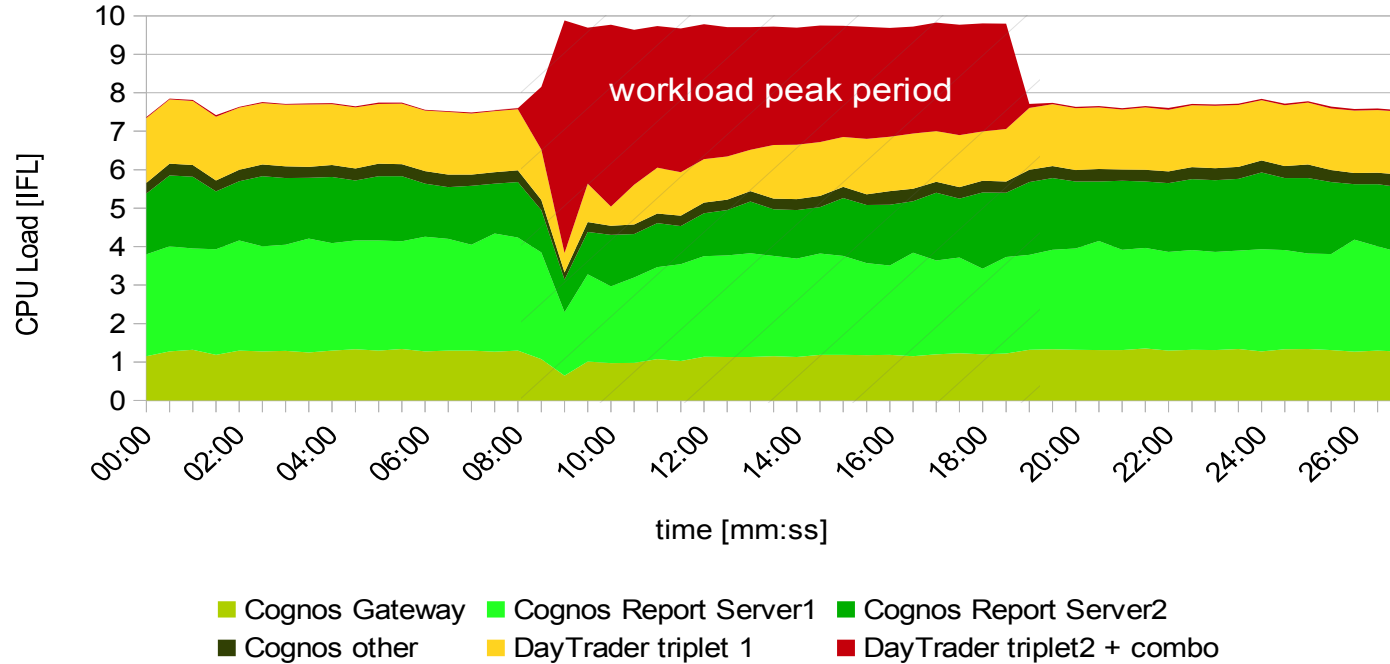
- VMRM CPU velocity goals guarantee a certain response time

Performance study – VMRM managing workload peaks (1)

Scenario 1: “Give others a chance”

Cognos BI + DayTrader workload peak - CPU load

workload peak with Cognos 70 load
 VMRM CPU goals: Cognos 70 / DayTrader 25




- Cognos BI CPU goal 70 and DayTrader CPU goal 25

Performance study – VMRM managing workload peaks (2)

Scenario 1: “Give others a chance”

Cognos BI + DayTrader workload peak - relative shares (selected guests)

workload peak with Cognos 70 load
VMRM CPU goals: Cognos 70 / DayTrader 25




- VMRM relative share adaption according to the CPU velocity goals

Performance study – VMRM managing workload peaks (3)

Scenario 2: “Cognos BI gets all”

Cognos BI + DayTrader workload peak - CPU load

workload peak with Cognos 70 load
VMRM CPU goals: Cognos 100 / DayTrader 1




- Cognos BI CPU goal 100 and DayTrader CPU goal 1

Performance study – VMRM managing workload peaks (4)

Scenario 2: “Cognos BI gets all”

Cognos BI + DayTrader peak workload - relative shares (selected guests)

workload peak with Cognos 70 load
 VMRM CPU goals: Cognos 100 / DayTrader 1

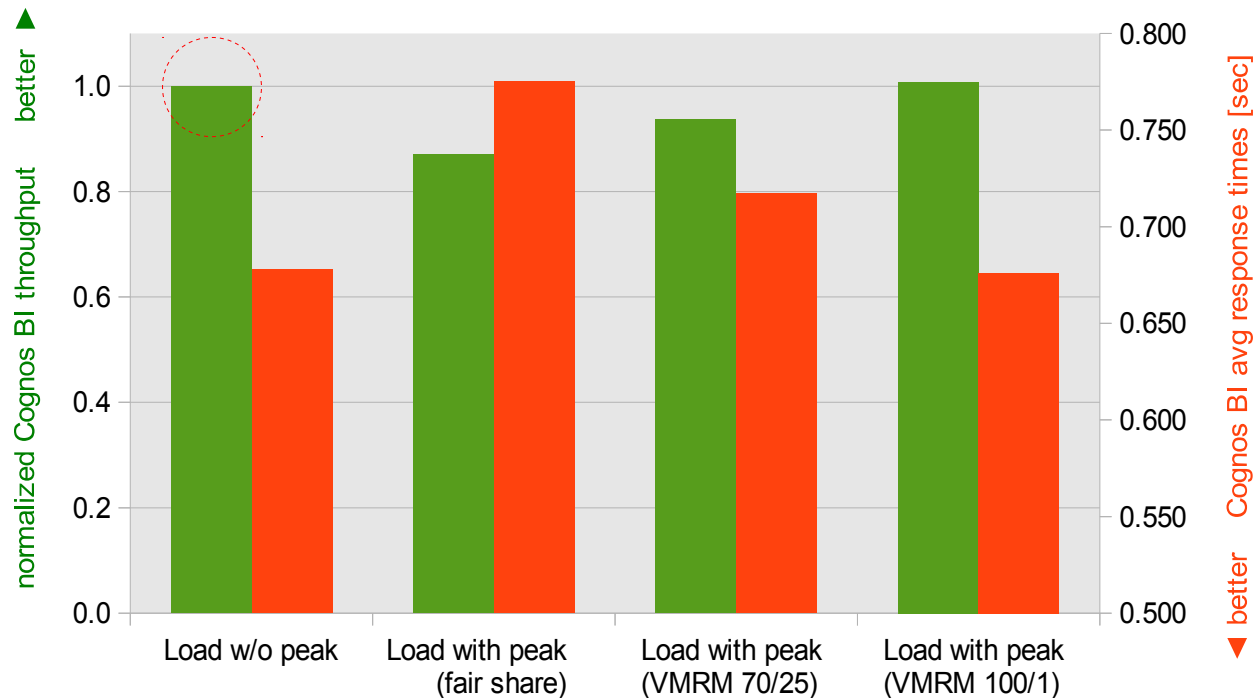


- VMRM relative share adaption according to the CPU velocity goals

Performance study – VMRM managing workload peaks (5)

Cognos BI throughput - compare all scenarios

Cognos BI + DayTrader peak workload impact on throughput and response time



- Maximum velocity goal almost maintains Cognos BI throughput level

Questions ?

- Further information

- For more detailed information see the available White Paper
“IBM Cognos Business Intelligence 10.2.1 for Linux on System z – Performance and z/VM Resource Manangement” (ZSW03268-USEN-00)
http://www.ibm.com/developerworks/linux/linux390/perf/tuning_vm.html#cog

Linux on z Systems – Tuning hints and tips

<http://www.ibm.com/developerworks/linux/linux390/perf/index.html>

- Live Virtual Classes for z/VM and Linux
<http://www.vm.ibm.com/education/lvc/>



Mario Held

*Linux on z Systems
Performance Analyst*

*IBM Deutschland Research
& Development
Schoenaicher Strasse 220
71032 Boeblingen, Germany*

*E-mail:
mario.held@de.ibm.com*

Backup

Complete your session evaluations online at www.SHARE.org/Orlando-Eval

FCX241, VM Resource Manager Screen – VMRM



In the VM Resource Manager Screen (FCX241), the names of workloads which have been active during the last measuring interval are highlighted on the screen.

Layout of VM Resource Manager Screen (FCX241)

FCX241	CPU nnnn	SER nnnnn	Interval	hh:mm:ss	-	hh:mm:ss	Perf. Monitor				
VM Resource Manager			Impor	<--	DASD	-->	<--	CPU	---	>	Active
Server	Workload		tance	D-Goal	D-Act		C-Goal	C-Act			Samples
VMRMSVM	WORK1		10	100	...		100	...			1
VMRMSVM	WORK2		5	50	...		50	...			1
VMRMSVM	WORK3		1	1	...		1	...			1
VMRMSVM	WORK4		10	100	100		100	87			1
VMRMSVM	WORK5		5	50	100		50	43			1
VMRMSVM	WORK6		1	1	100		1	7			1
VMRMSVM	WORK7		10	100	100		100	83			1
VMRMSVM	WORK8		5	50	100		50	41			1
VMRMSVM	WORK9		1	1	...		1	...			1
Command	====>	_									
F1=Help	F4=Top	F5=Bot	F7=Bkwd	F8=Fwd	F12=Return						

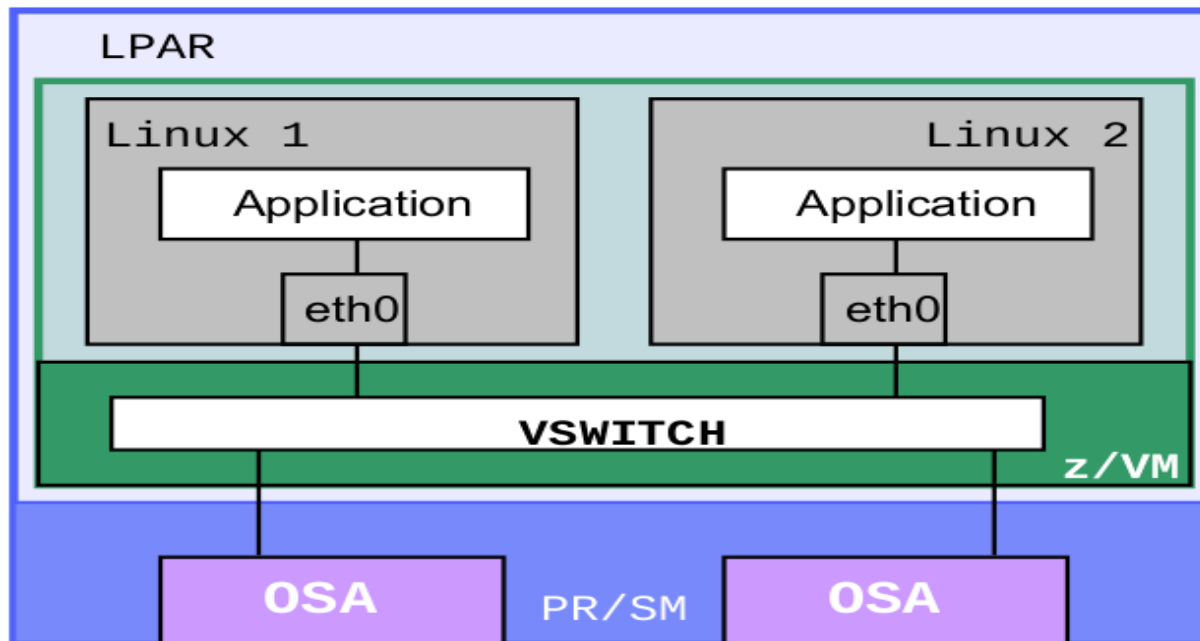
The information shown is based on CP monitor application data domain SAMPLE data.

Networking



z/VM Virtual Switch (VSWITCH) for virtual machine communication

- Good performance
- Recommended method for internal and external z/VM network connectivity
- Special purpose GuestLAN
- Linux reports (lsqeth) it as card_type "GuestLAN QDIO"



More information: <http://www.vm.ibm.com/virtualnetwork>