

Cloud Performance Testing

Cloud Performance Metering as Pre-Requisite for Performance Testing in the Cloud

Dipl.-Wirtsch.-Inf. Jan Thielscher , LL.M.

Ehningen, Juni 2012

Agenda



EACG in brief

Why (Performance) testing in the Cloud?

User Space Monitoring

Insights from test results

ascamso Demonstration

EACG – Your architecture consultancy for IT-based business models

EACG Value Proposition

We partner with our customers from business model design to execution supporting with reflection, expertise, skills, & resources where necessary

EACG Service Offering

	Financial Services	Retail / eCommerce	...
IT-Strategy/ IT-Business Alignment			
Enterprise Architecture/ BPM/BRM/CEP-Initiatives			
Cloud Transformation & Cloud Architecture			

EACG Project Samples



Design Enterprise Architecture & Support Implementation



Design corporate wide E-Commerce Target Architecture



Pan-European Rollout coordination of Import Express Online



Design, Implementation & Operation of Middleware for Shop Platform

Customer Statements

„Mich haben neben der pragmatischen Herangehensweise das gute Verständnis unserer Fachprozesse und die klaren Empfehlungen begeistert“

Frank Dupuis, CEO
Dupuis GmbH & Co. Asset-Management KG

„Wir konnten sowohl in fachlichen wie in technischen Aspekten neue, wertvolle Perspektiven gewinnen.“

Harm Behrens,
Leiter E-Commerce Competence Center
Group technology Partner (gtp) Otto Group

Using cloud solutions may have a strong impact on budget

However, testing requires a reliable base for comparison

Motivation (from a real scenario)

Scenario:

Selection of a middleware platform as central backbone for an E-commerce-provider

Business benefit from project:

0,5 Mil. EBIT/month*)

Resource requirements:

20 x 4-Core-Systems, 16GB RAM and 250GB disk

Further assumptions:

- Purchase of a certain amount of systems will be required to continue current operations and operations of test environment anyway
- All candidates are installable in a cloud environment
- It is not necessary to prove the overall capability to handle required performance by stressing the systems to their maximum extent

**) modified from real case. however, case looks positive even with a few EUR only*

Business case

A rough comparison shows the benefits from using a cloud approach

Real	Comparison class	Cloud
30Tage	Time to order	0.5PD
30Tage	Time to supply	0.5PD
100k	Costs of purchase	10k
1-2PT	Time to recover lost systems	0.5PD
0 EUR	Benefits from saving 3 months	1.5 Mil.

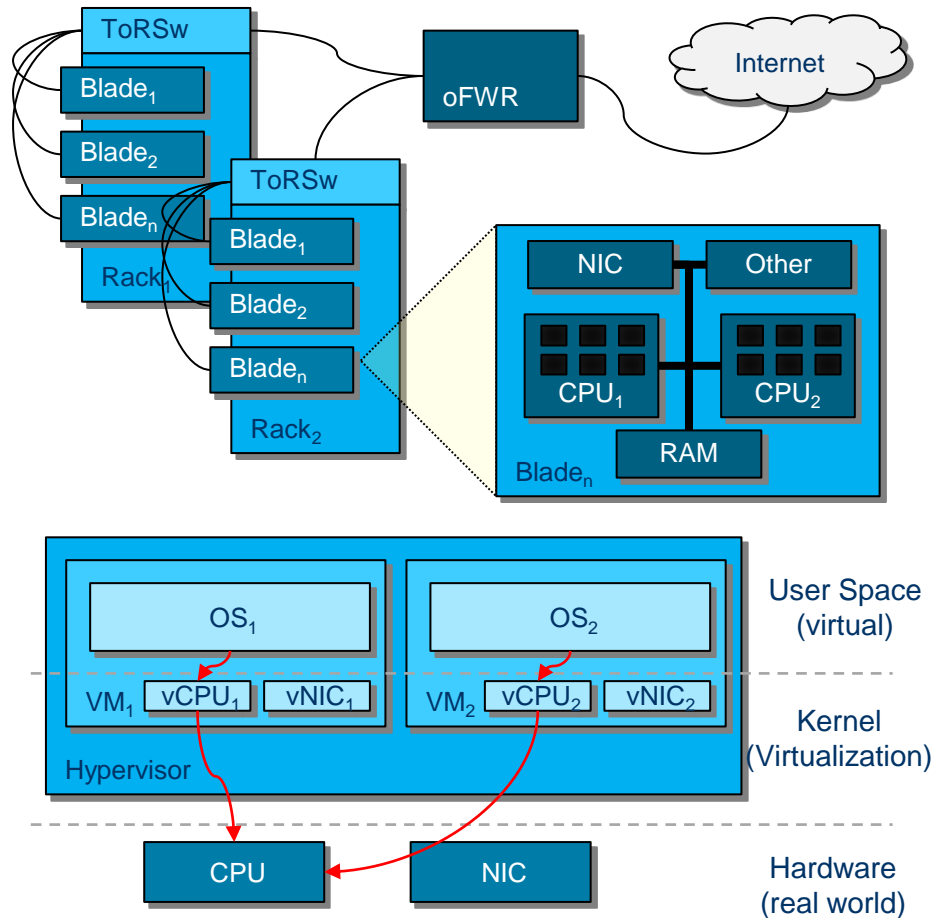
Further benefits from using images instead of real installations:

- Whole environment is archive-able
- It is possible to re-setup the whole environment even after weeks for re-runs with low efforts
- Re-use of tests for later purposes is simple

Performance of cloud infrastructures is driven by cloud design

Poor designed architectures leave room for inconsistencies

Cloud Infrastructure— schematic view



Challenges to master

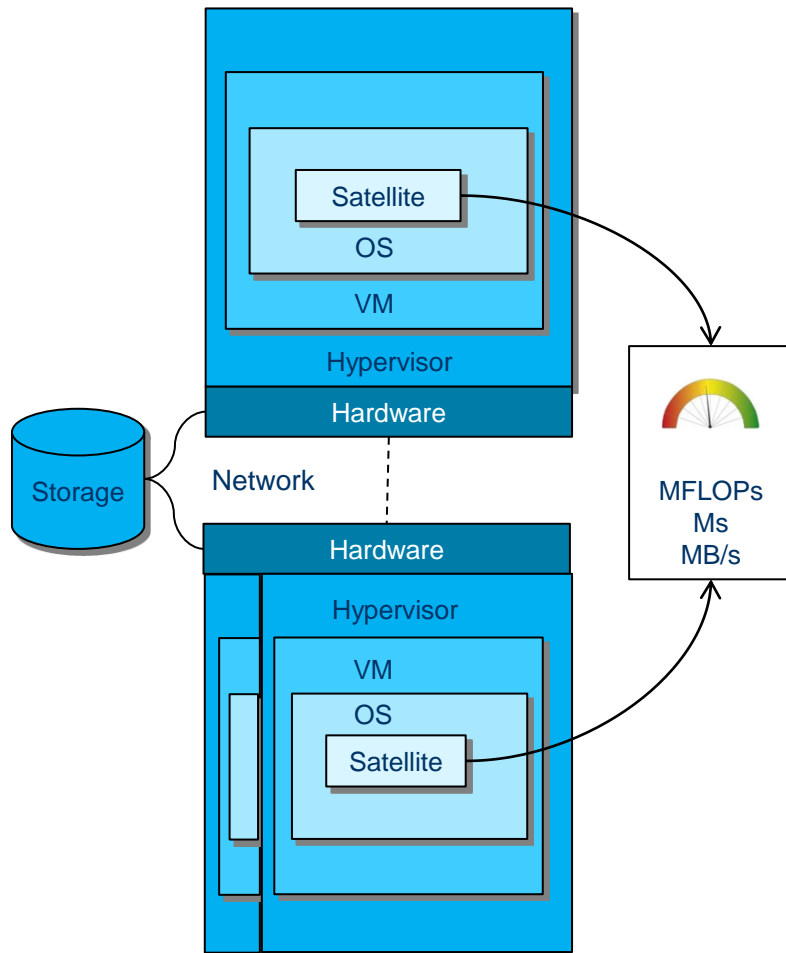
Performance of cloud is driven by many factors

- Hardware-platform: What is the underlying hardware that is to be shared?
- Overprovisioning rate: How many virtual cores reside on a real core?
- VM-isolation: How well provides the hypervisor the existing capacities to its tenants? Will heavy computing workloads on one VM impact neighboring VMs?
- Disc(Hdd): Only few applications do not require at least swap or log-space on hard drives during runtime. How well is the connectivity to either network or virtual disc?
- Network: How many virtual systems share how much bandwidth and how many real interfaces? Will the accessibility of real NIC be given in the right time? How is internet-access, how is inter-system-traffic routed through the VLAN isolating the virtual systems?

To evaluate purchased capacities properly the price / performance ratio has to be measured. This requires normalized metrics

Only User Space Monitoring unveils the real capacities provided

Measuring capacities and achievable throughput at point of delivery encapsulates complexity



Measuring provided compute capacity is complex

- Standard CPU usage reports are useless for virtual environments
- Cores and GHz are only weak indicators as compute capacity is strongly driven by simulated CPU chipset
- Real capacity indication only can be retrieved from calculations inside simulated machine

Network latency can be another bottleneck

- Many factors influence network latency
 - Size of backbone (1, 10, 40GB?)
 - Virtualization rate of NICs (#instances/NIC)
 - VLAN configurations & routing priorities
 - Hypervisor security
- To overcome complexity package transfer between locations gives the correct answer

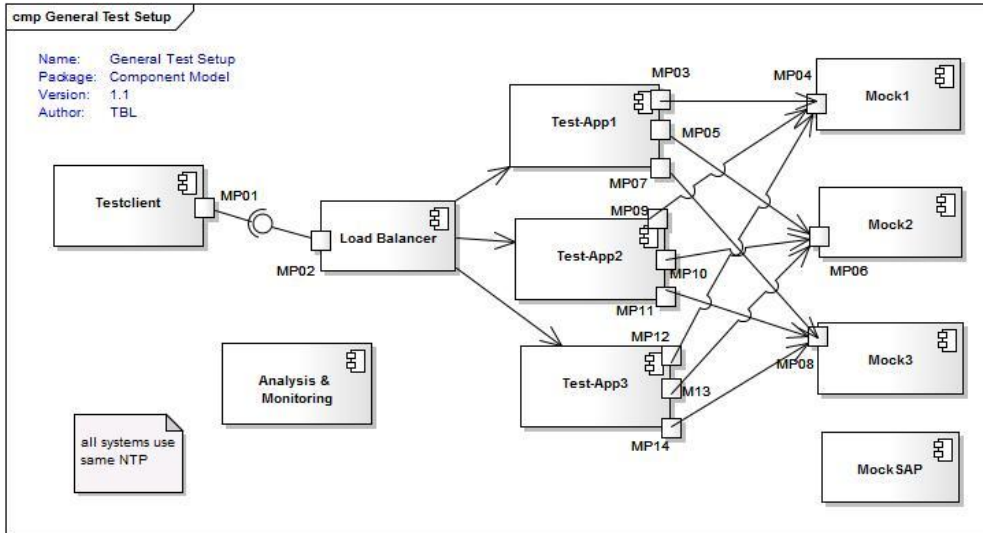
Most applications require acceptable disk speed

- IOPS are not measurable w/o access to hardware
- Cloud infrastructures know many storage types (NAS, SAN, etc.) and technologies (iSCSI, Infiniband, etc.)
- Performance on same hardware might be different depending on usage pattern (Tx-log, DB or File Server)

Just User Space experience gives a comparable measurement

Monitoring of environment performance throughout whole testing requires tool support

Test environment comprises of about 10 instances



Two business scenarios defined to provide many technical test cases


- Transformation and mapping capabilities
- Content based routing
- Usage of different protocol adapters
- Publish/subscribe mechanism
- Scalability 1,2 and 3 node clusters
- Disaster behavior
 - 1 cluster member lost
 - 2 cluster member lost
 - One member recovery

Test duration per vendor: 2 days

Same could quality and repeatable test settings throughout the whole test need to be guaranteed to allow acceptable test results

CPU capacity has been very stable throughout testing period

Sample of metering results from CPU test



CPU capacity remained stable:


- MFLOPS are calculated by executing a Fast Fourier Transformation as often as possible in given timeframe (metering interval)
- Here test interval was set to 5 seconds
- Metering interval was set to 100ms

Single fallouts are usual in capacity metering

- Single break downs are common for systems running background tasks
- Critical is a line that drops in performance for several metering attempts in a row

Network testing showed three different qualities

Sample of metering results for network latency



Network latency metered showed three levels:

- Three classes were metered:
 - 3-4ms
 - around 140ms
 - around 270ms

Please note: figures represent roundtrip times

- Due to early metering the results were used to distribute machines accordingly
- Disabling standard anti-virus features on several levels reduced times further
- However, results were communicated to IBM SCE team and they started assessing the issues from vendor side...

User Space Cloud Metering is made simple and easy to use!

Welcome to ascamso!

(ascamso - A Simple Cloud Assessment & Monitoring Solution!)

The
Cloud performance benchmarking
tool

Did you ever encounter doubt whether your VMs really deliver the performance you bought them for?
If yes, this is the place to end your doubts: By performing specific user space monitoring, ascamso unveils the real service quality of your IaaS.

To learn more, just become a member of the ascamso tester cloud. It's free to join, but it will grow your insight!

[Sign Up](#)

Jan Thielscher

*Enterprise Architecture Consulting Group –
EACG GmbH*

Bockenheimer Landstraße 2-4
OpernTurm, 16.OG
D-60306 Frankfurt am Main

T: +49 69 667 748 280

E-mail: info@eacg.de
