


Apache Cloudstack as Cloud-Management-Environment

Claus Kalle

RRZK Uni Köln, 27.03.2015



Disclaimer



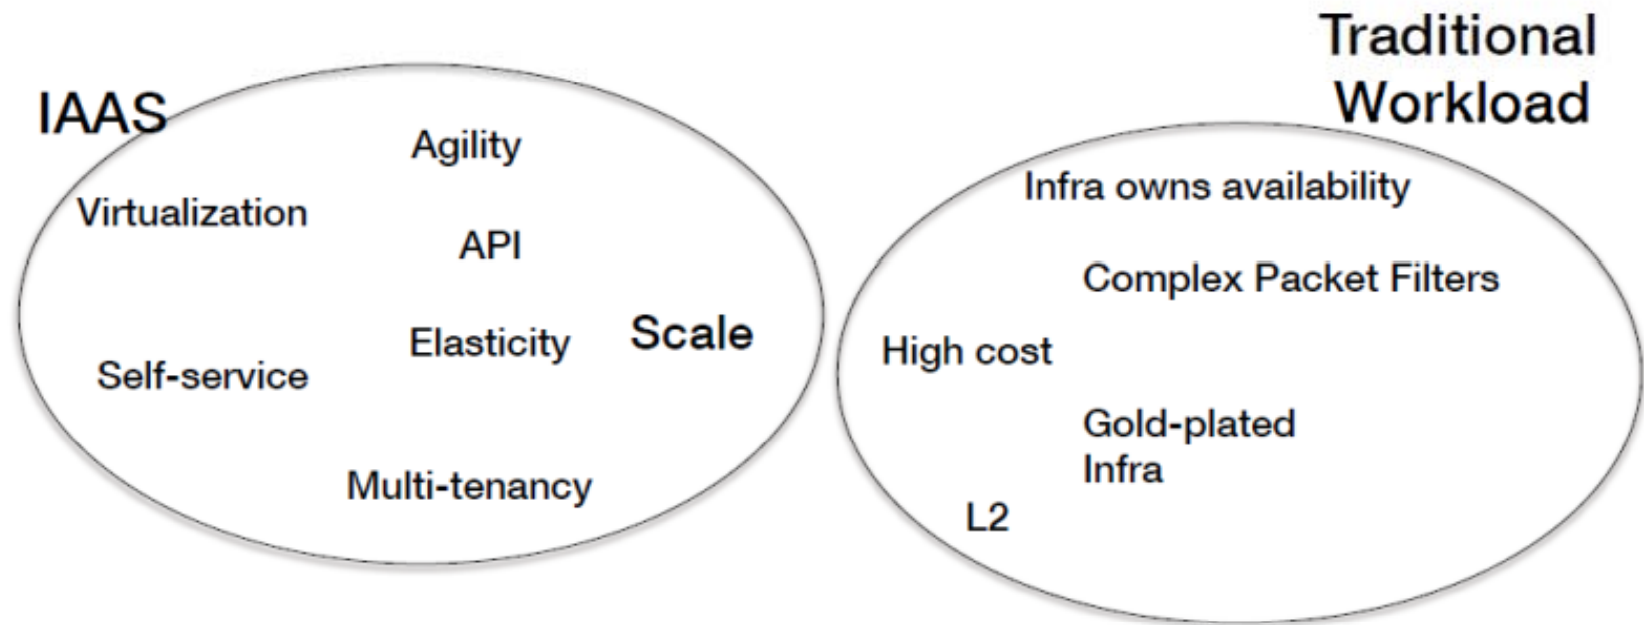
Agenda

- **Cloud**
- **VMware**
- **OpenSource: Apache Cloudstack**
 - Functions
 - Network-Options
- **Functional Comparison**



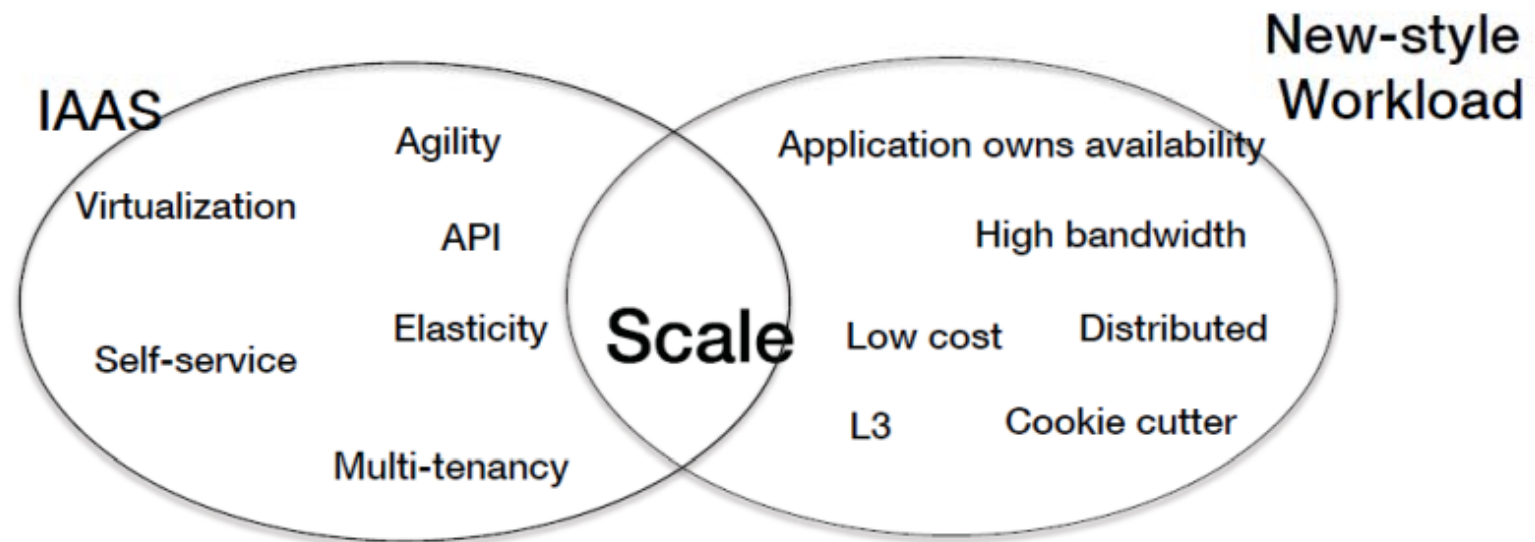
Cloud - Recap 1

IAAS – PAAS – SAAS



It is possible to realize some of the benefits of IAAS for traditional workloads

Cloud - Recap 2



These classes of drivers (IAAS and new-style workloads) are highly complementary and therefore most new-style applications operate on IAAS

Cloud - Recap 3

Scale

“At scale, everything breaks”

- Urs Hölzle, Google

8%

Annual Failure Rate of servers

Kashi Venkatesh Vishwanath and Nachiappan Nagappan, *Characterizing Cloud Computing Hardware Reliability*, SoCC'10

Server failure comes from:

- 70% - hard disk
- 6% - RAID controller
- 5% - memory
- 18% - other factors

Application can still fail for other reasons:

- Network failure
- Software bugs
- Human admin error

Redundancy helps a little

40%

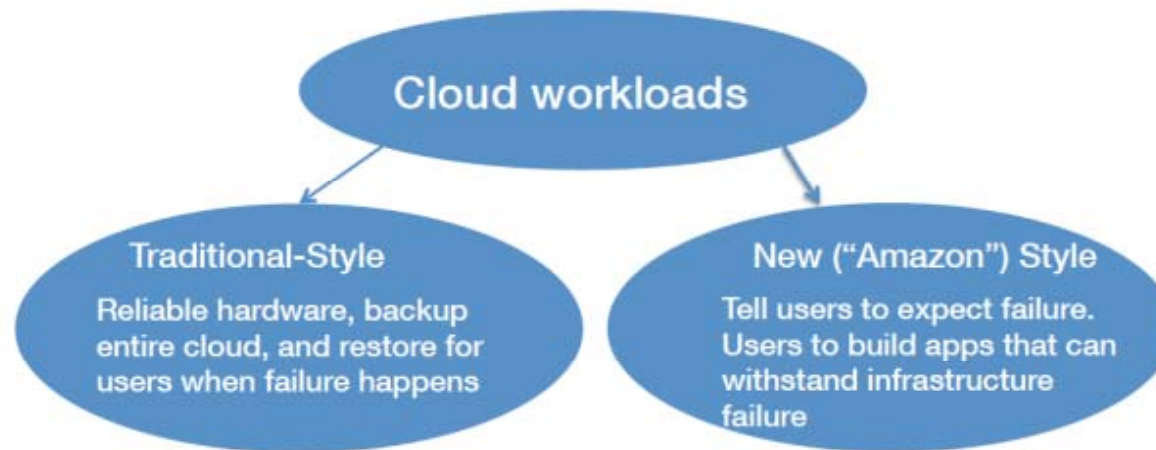
Effectiveness of network redundancy in reducing failures

Phillipa Gill, Navendu Jain & Nachiappan Nagappan, *Understanding Network Failures in Data Centers: Measurement, Analysis and Implications*, SIGCOMM 2011

- Bugs in failover mechanism
- Incorrect configuration
- Protocol issues such as TCP back-off, timeouts, and spanning tree reconfiguration

Cloud - Recap 4

Reliability Strategies



Both styles of workloads must run reliably in the cloud

Reliability Styles

Traditional workload

Link aggregation

Storage multi-pathing

VM HA, fault tolerance

VM live migration

Expect reliability. Back-up entire cloud. Admin controlled failure handling
Think Server Virtualization

Cloud workload

VM backup/snapshots

Ephemeral resources

Chaos monkey

Multi-site redundancy

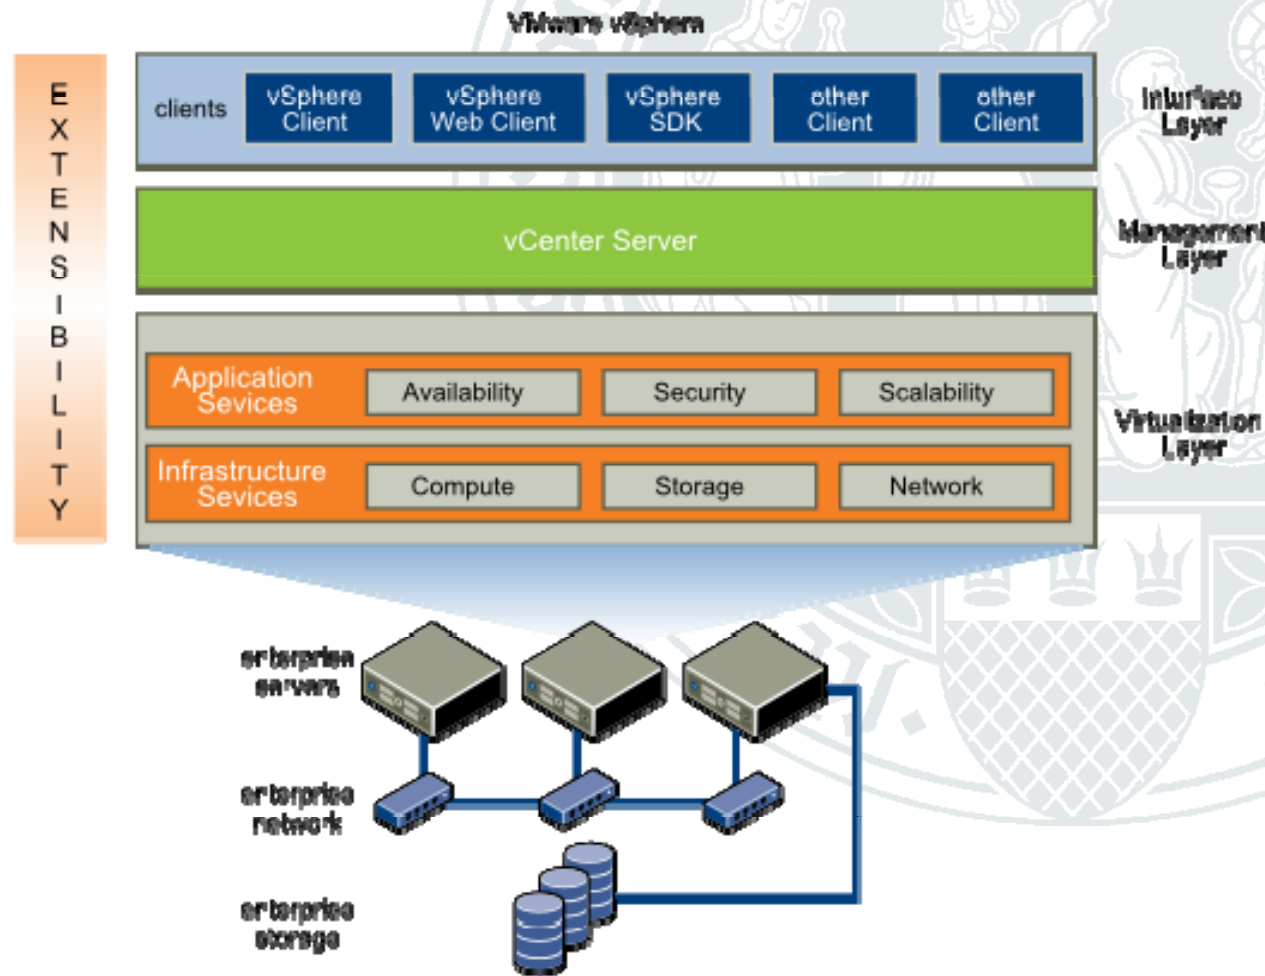
Expect failure. Design app for failure. Self-service failure handling
Think Amazon Web Services

Agenda

- Cloud
- **VMware**
- OpenSource: Apache Cloudstack
 - Functions
 - Network-Options
- Functional Comparison



VMware Components Breakout into Layers



VMware Products

Produkte >

Virtualisierung von Rechenzentren und für Cloud-Infrastrukturen

vCloud Suite

Computing

vSphere

vSphere with Operations Management

Netzwerke und Sicherheit

NSX

Storage und Verfügbarkeit

Virtual SAN

vCenter Site Recovery Manager

vSphere Data Protection Advanced

Virsto

vSphere Storage Appliance

Infrastructure as a Service

VMware vCloud Hybrid Service

Management von Rechenzentren und Clouds

Cloud-Betriebsmanagement

vCenter Operations Management Suite

vCenter Operations Manager

vCenter Configuration Manager

vCenter Hyperic

vCenter Log Insight

vCenter Server

vCloud Director

Cloud-Automatisierung

vCloud Automation Center

vCenter Orchestrator

IT-Business-Management

IT Business Management Suite

IT-Benchmarking

Desktop- und Anwendungsvirtualisierung

Horizon (mit View)

Horizon DaaS

Mirage

Personal Desktop

Fusion

Fusion Professional

Workstation

Player Plus

Kostenlose Produkte

vSphere Hypervisor

vCenter Converter

Compliance Checker für PCI

VMware Compliance Checker für vSphere

Capacity Planner

Infrastructure Planner

Verwandte Themen

Software-Defined Datacenter

Virtualisierung

Cloud Computing

End-User Computing

Virtualisierung von Enterprise-Anwendungen

Branchenlösungen

Lösungen für kleine Unternehmen

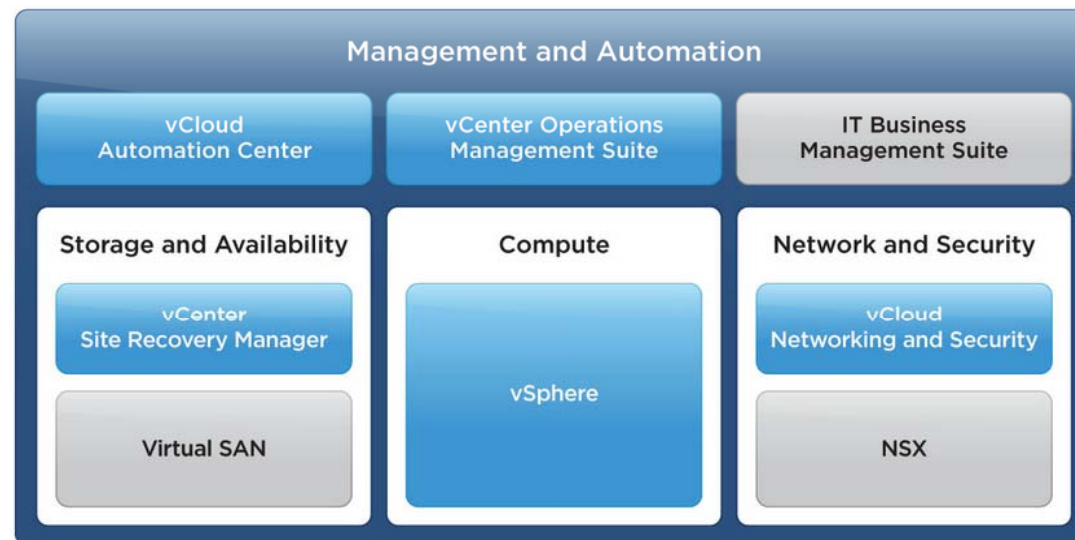


VMware Development

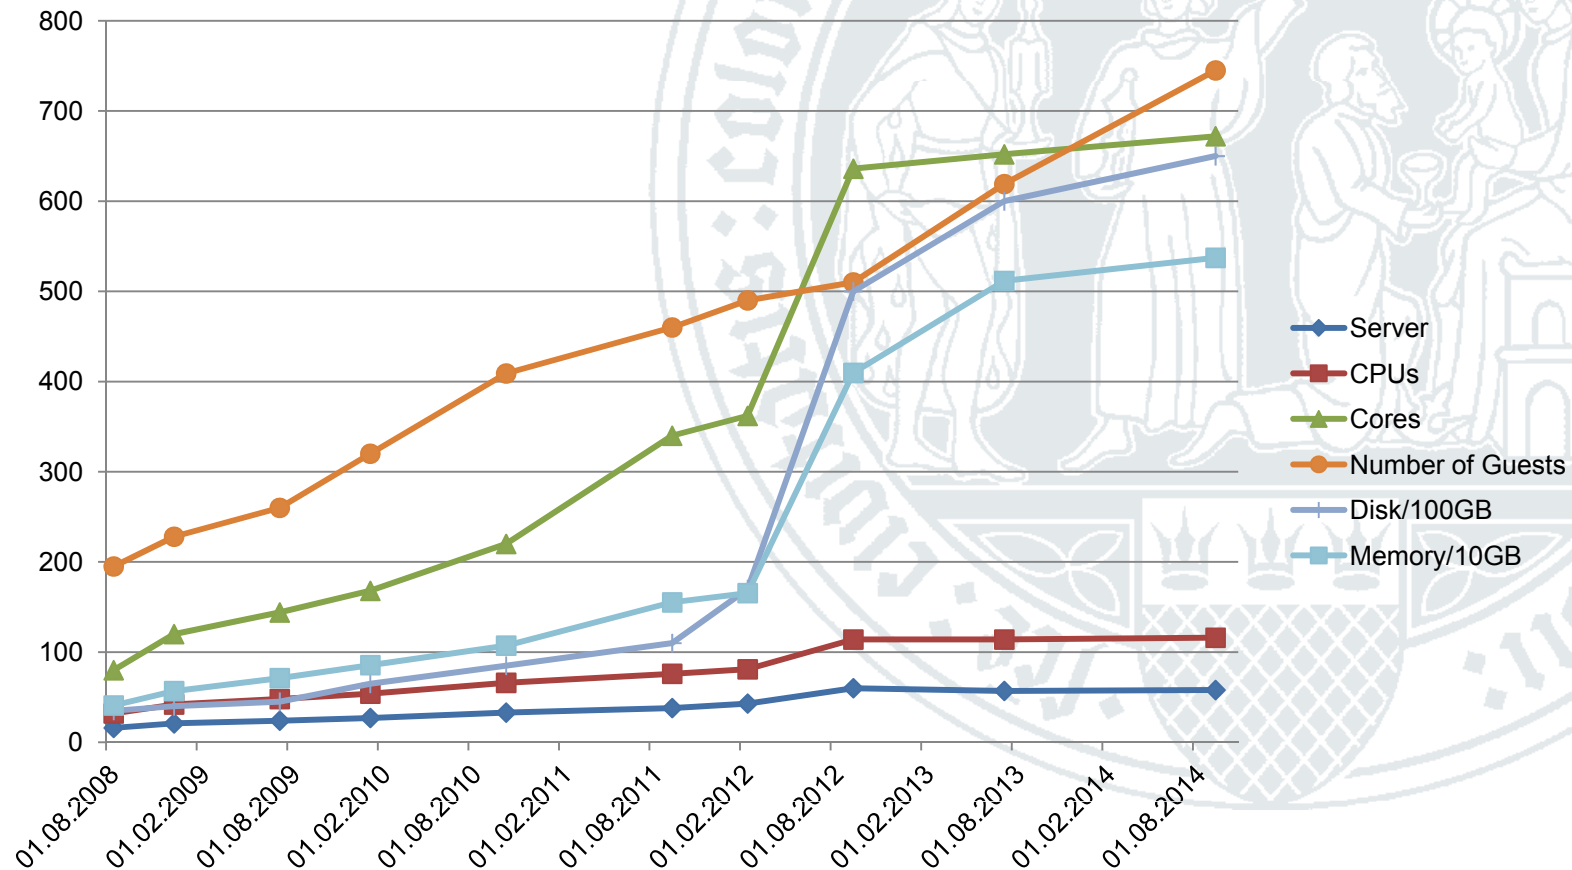
- ESXi 5.5
- vCenter
- Cloud Director

vCloud Suite

The First Step to a Software-Defined Data Center Architecture



VMware at RRZK: Growth



And again new Bundles in 2014



And the end of the vCloud dreams ?

Email as
of
17.02.15
17:04h



The image shows a screenshot of an email notification from VMware. At the top left is the VMware logo. To the right are social media icons for Facebook, Twitter, LinkedIn, and Blogger, followed by a YouTube icon and a 'My VMware' link. Below this is a blue header bar with a geometric pattern. The main content of the email is a notice titled 'Important Notice: vCloud Director and vCloud Networking and Security End of Availability'. The text reads: 'Dear Valued Customer, VMware is announcing the End of Availability (EOA) for the VMware vCloud Director® and vCloud® Networking and Security™ products as of the release of vCloud Suite® 6.0. This letter serves as a final notification for both products, as neither will be available as a standalone offering (announced in

VMware Pricing

vSphere Enterprise Plus

- per socket Invest
- ca. 1800 €
- Per socket S&S / year
- ca. 450 € (150 sockets → ~70000 €)

For 2-socket M620 blade, 3 Jahre S&S

- ca. 6300 €

Hardware ca. 5500 €

4 times overcommitment gains 10000 €!

VMware core functions

- vMotion
- HA
- Heterogeneous Guests (caveat MS-Licencing)
- Import of VLANs
- Distributed vSwitch
- NSX: virtual Networking
- LDAP-Integration
- Templates
- **MISSING:**
easy to use Self-Service Function

Comparison with other Solutions

	VMware vSphere 5.5 Enterprise Plus + vCenter Server 5.5	Microsoft Windows Server 2012R2 + System Center 2012 R2 Datacenter Editions	RHEL6 KVM (lt. Redhat Webpage)
Unterstützung physikalischer CPUs pro Lizenz	1	2	
Managed OSE's pro Lizenz	unlimited	unlimited	
Lizenz für Enterprise Operations Monitoring und Management der Hosts, Guest VMs und Application Workloads, die innerhalb der VMs laufen	Nein, separate Lizenz über vCenter Operations Manager oder Upgrade auf vSphere mit Operations Management et eifoiderlich	Ja, im System Center 2012 R2 enthalten	
Lizenz für Private-Cloud-Management-Fähigkeit	Nein, Private-Cloud-Management-Fähigkeit erfordert eine Lizenz fiber VMWare vCloud Suite.	Ja, im System Center 2012 R2 enthalten	
Maximale Anzahl logischer Prozessoren pro Host	320	320	4096
Maximaler physikalischer RAM pro Host	4 TB	4 TB	64 TB
Maximale Anzahl aktiver VMs pro Host	512	1024	
Maximale Anzahl virtueller CPUs pro VM	64	64	160
Maximaler virtueller RAM pro VM	1 TB	1 TB	2 TB
Maximale Anzahl physikalischer Hosts je Cluster	32	64	
Maximale Anzahl der VMs je Cluster	4000	8000	
Virtual Machine Snapshots	bis zu 32 Snapshots pro VM	bis zu 50 Snapshots pro VM	
Application Load Balancing for Scaling-Out Application Tiers	Nein, erfordert die Versionen vCloud Network and Security or dei vCloud Suite	Ja, über System Center 2012 R2 VMM	
Bare metal deployment of new Storage hosts and clusters	Nein	Ja, über System Center 2012 R2 VMM	
Virtualization of USB Devices	Ja, via USB Passthrough support	Ja, via Remote Desktop connections	
Live Migration of running VMs	Ja, (limited to 4 concurrent vMotions per host when using 1GbE network adapters and 8 concurrent vMotions per Host when using 10GbE)	Ja (unlimited concurrent live VM Migrations)	
Live Migration using compression of VM memory state	Nein	Ja	
Maximale Anzahl der Nodes pro VM Guest Cluster	5	64	
Maximale Anzahl virtueller SCSI HDDs pro VM	60 (PVSCSI) 120 (Virtual SATA)	256 (Virtual SCSI)	
Native 4K Disk Support	Nein	Ja	
Hot-Shrink Virtual SCSI Hard Disks for running VMs	Nein	Ja	
Flash-based Read Cache	Ja, aber nur bis max. 400GB des Caches pro virtueller Disk / 2TB zusammengefassten Caches pro Host für alle virtuellen Disks	Ja, limitiert auf his zu 160 physikalische Disks und 480TB gesamter Kapazität	
Flash-based Writeback Cache	Nein	Ja	
Automated Tiered Storage between SSD and HDD using commodity hard disks	Nein	Ja	
Storage Encryption	Nein	Ja, per BitLocker	



Get out of the Hidden Cost Trap!

webcast

COMPUTERWOCHE

Cloud-Cookbook

ENGINE START

Smarte Cloud Lösung. OpenSource!
Sichere Private Cloud-Infrastruktur im Mittelstand - per Knopfdruck

Agenda

- Cloud
- VMware
- **OpenSource: Apache Cloudstack**
 - Functions
 - Network-Options
- Functional Comparison



4 Opensource-Cloud-Products

AK-ZSYS 2011

OpenNebula.org
The Open Source Toolkit for Cloud Computing

- HPC-oriented

EUCALYPTUS

Gekauft von



- Small scale,
not much network

AK-ZSYS 2014



openstack
CLOUD SOFTWARE



- Toolkit, HPC, roll
your own cloud

cloudstack
open source cloud computing



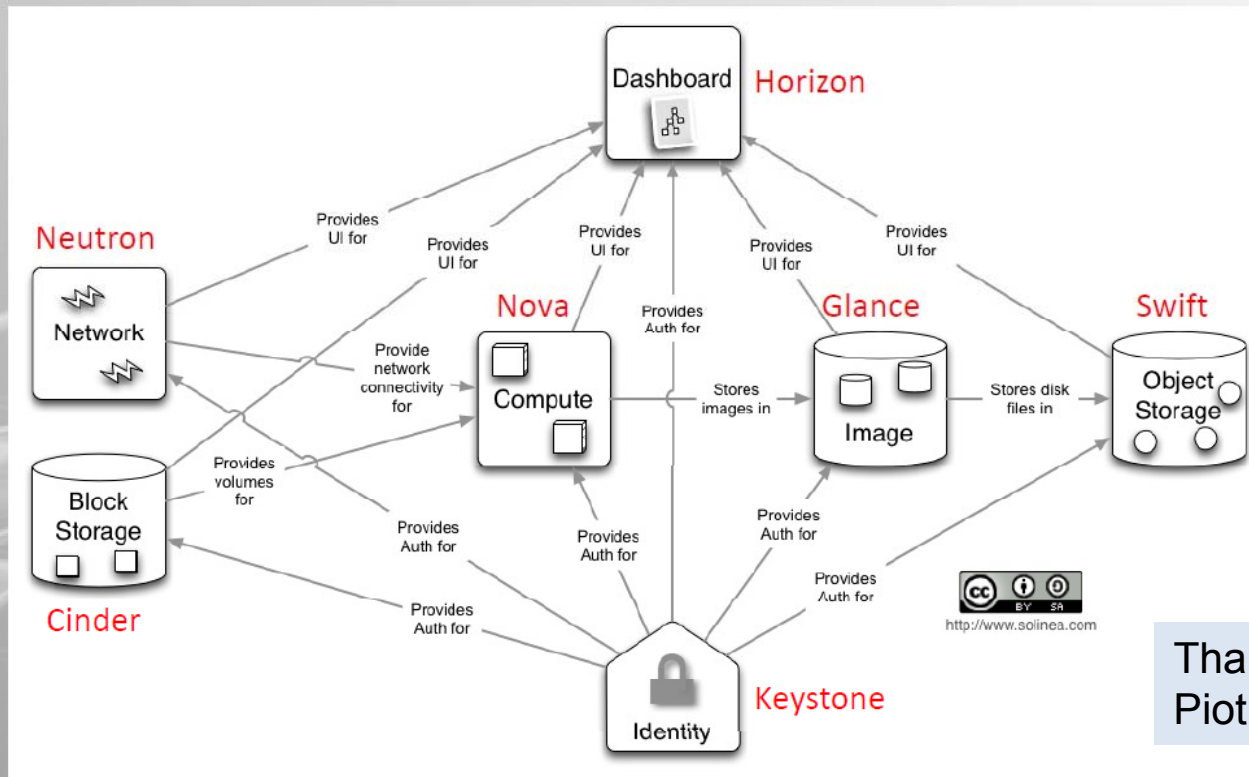
- Mid- to Hi-End,
smaller footprint



Openstack Architecture



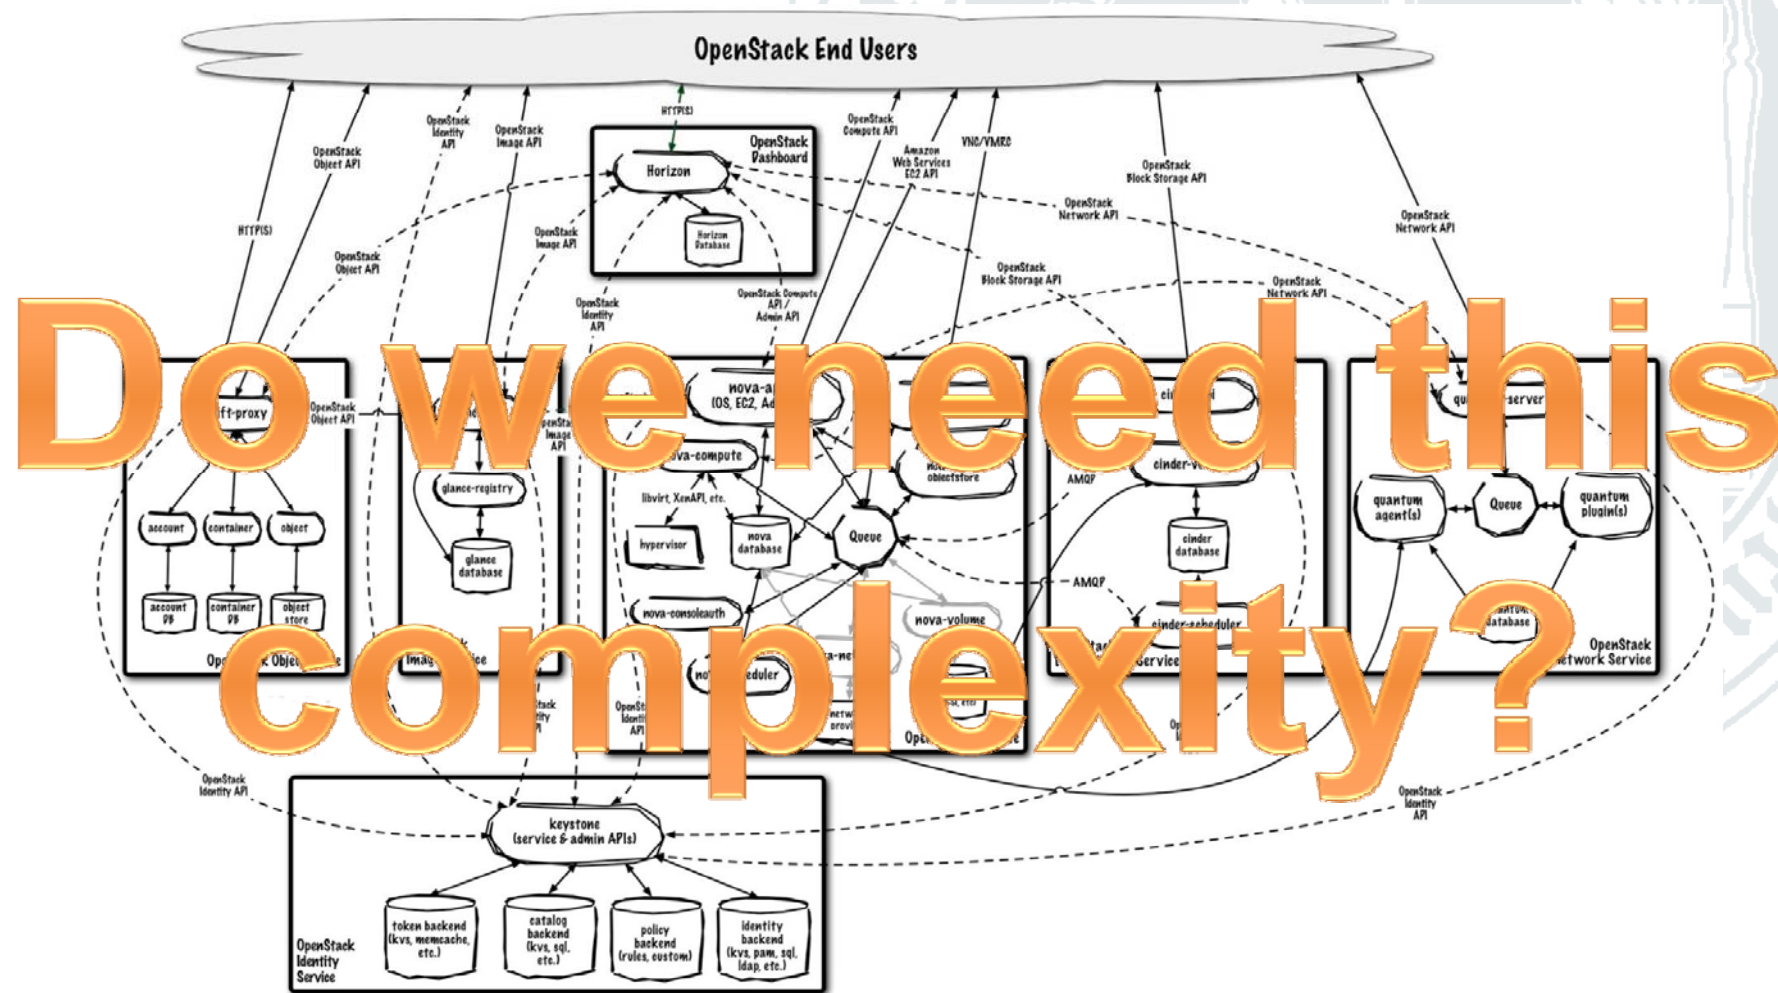
Architektur



Thanks to
Piotr Kasprzak (GWDG)



Openstack Logical Diagram

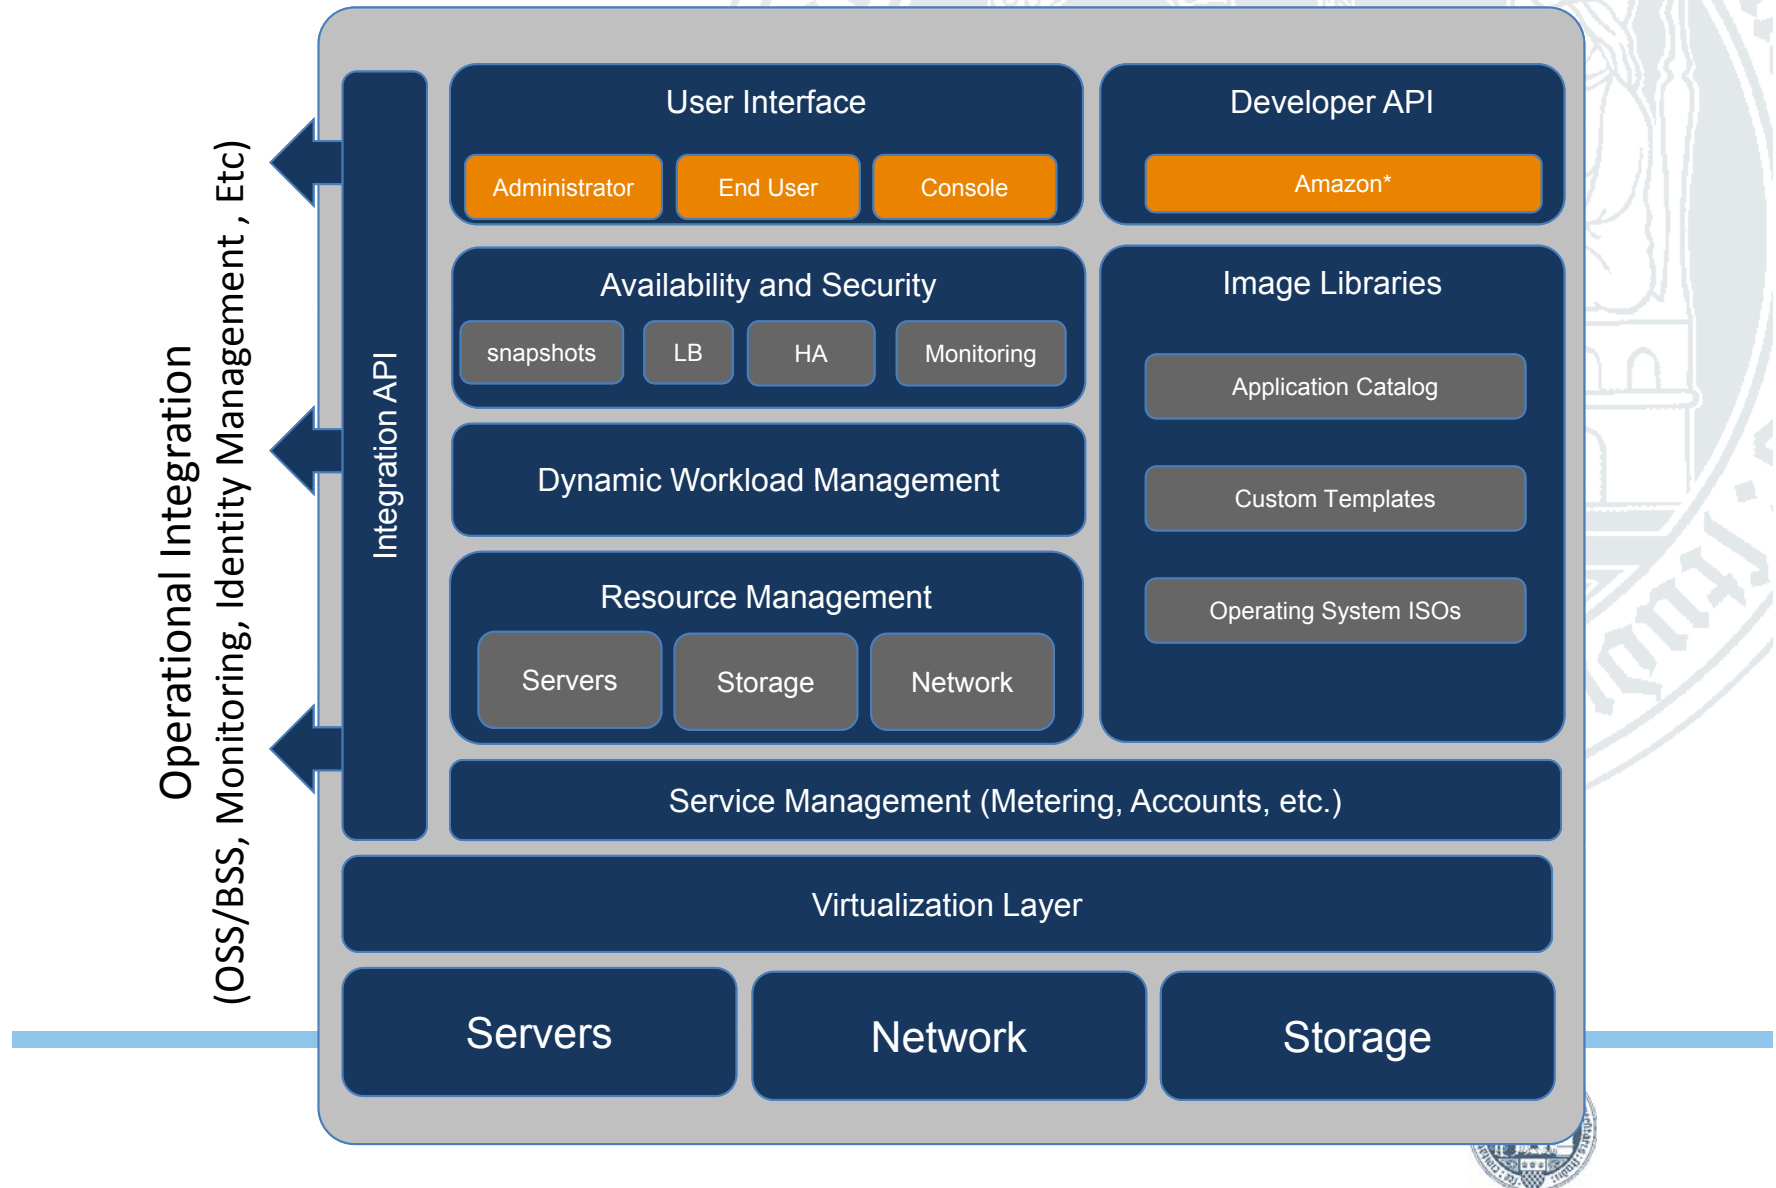


Apache Cloudstack



- Open Source
 - Apache Software License
 - Top Level Project in the Apache Software Foundation since April 2013
 - Open Source since May 2010
- In production since 2009
- Multi-tenant cloud orchestration platform
 - Turnkey platform for delivering IaaS clouds
 - Over 200 commercial deployments: private and public
 - Full featured GUI, end-user API and admin API

CloudStack Architecture



Vokabulary ...

pristine

fresh, clean, and unspoiled

to expunge

To **eliminate** completely; **annihilate**.

curmudgeon

An ill-**tempered** (and frequently **old**) **person** full of **stubborn** ideas or **opinions**.

ephemeral

Ephemeral things are transitory, existing only briefly.

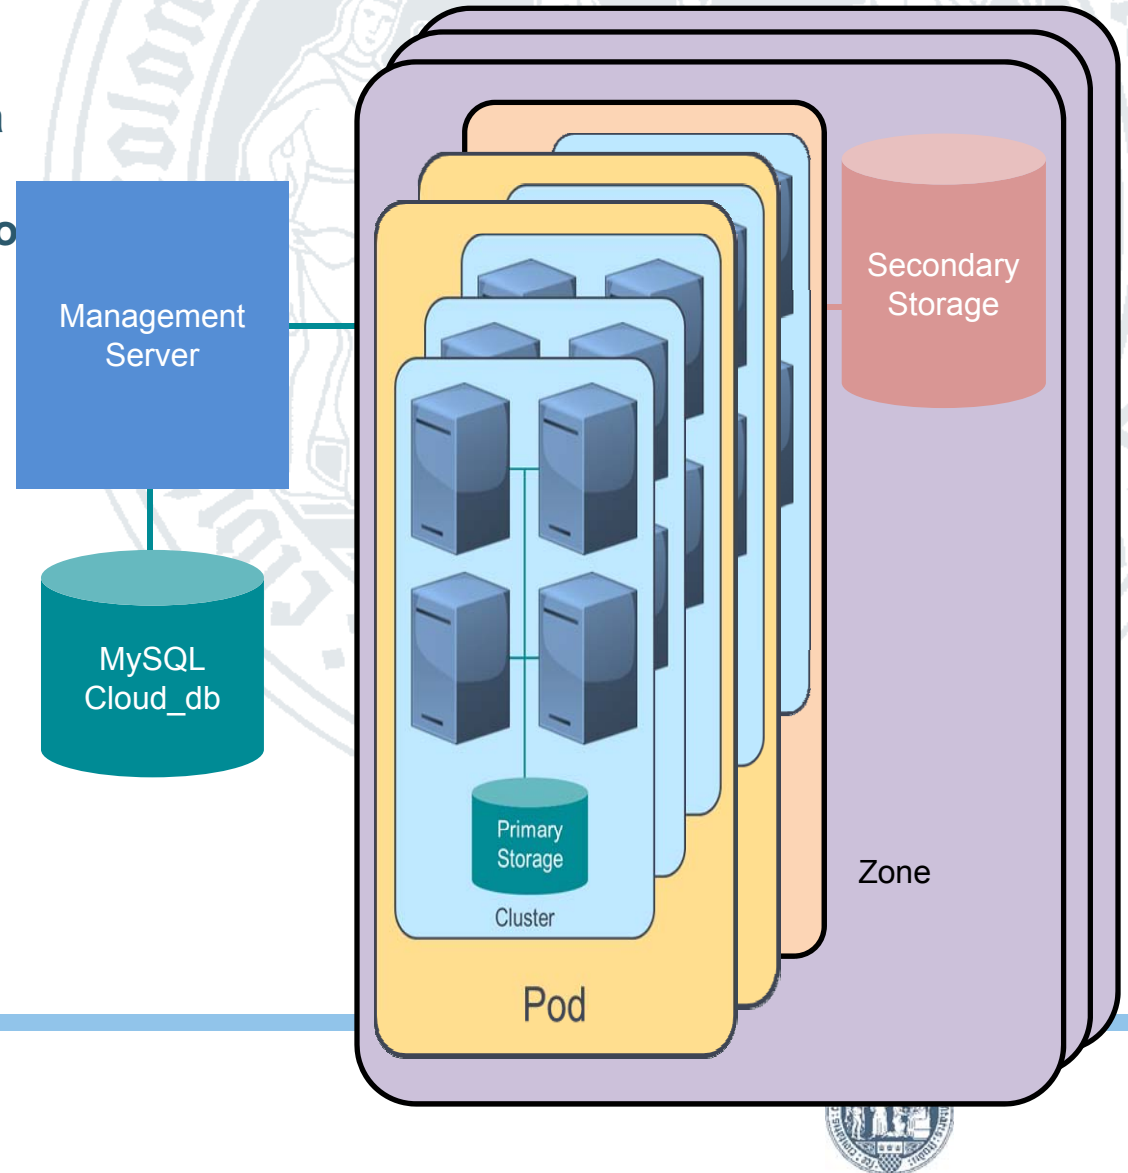
... prior knowledge

VMware-player, Netcfg vSphere chaosmonkey
Centos Ubuntu
ec2/s3 ajax ESX5i mysql iptables
maven java udev kvm
ruby nfs virsh ldap ovs-vsctl ebtables
cloudmonkey soap rest brctl json
xml

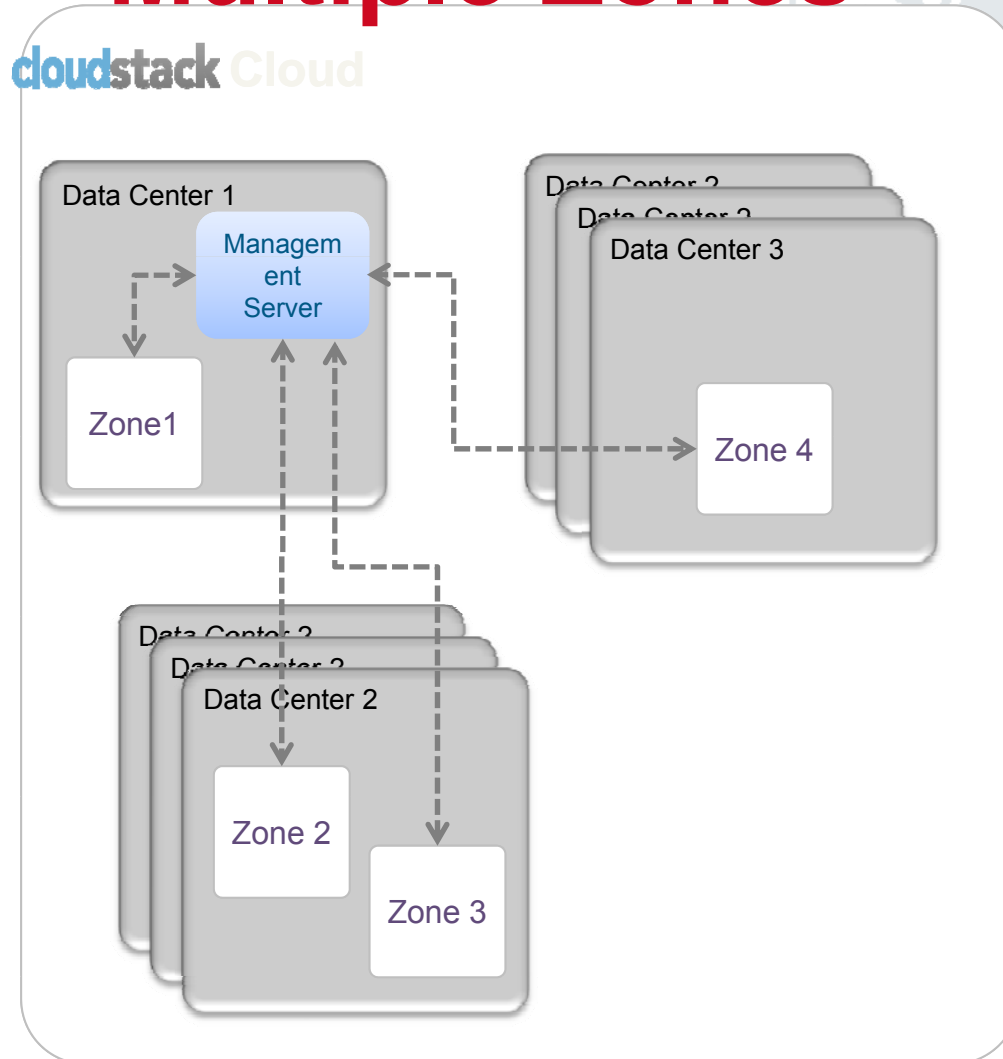


Cloud Infrastructure Overview - Summary

- One or more hosts grouped into a cluster
- One or more clusters grouped into a pod
- One or more pods grouped into a zone
- One or more zones controlled by one management server



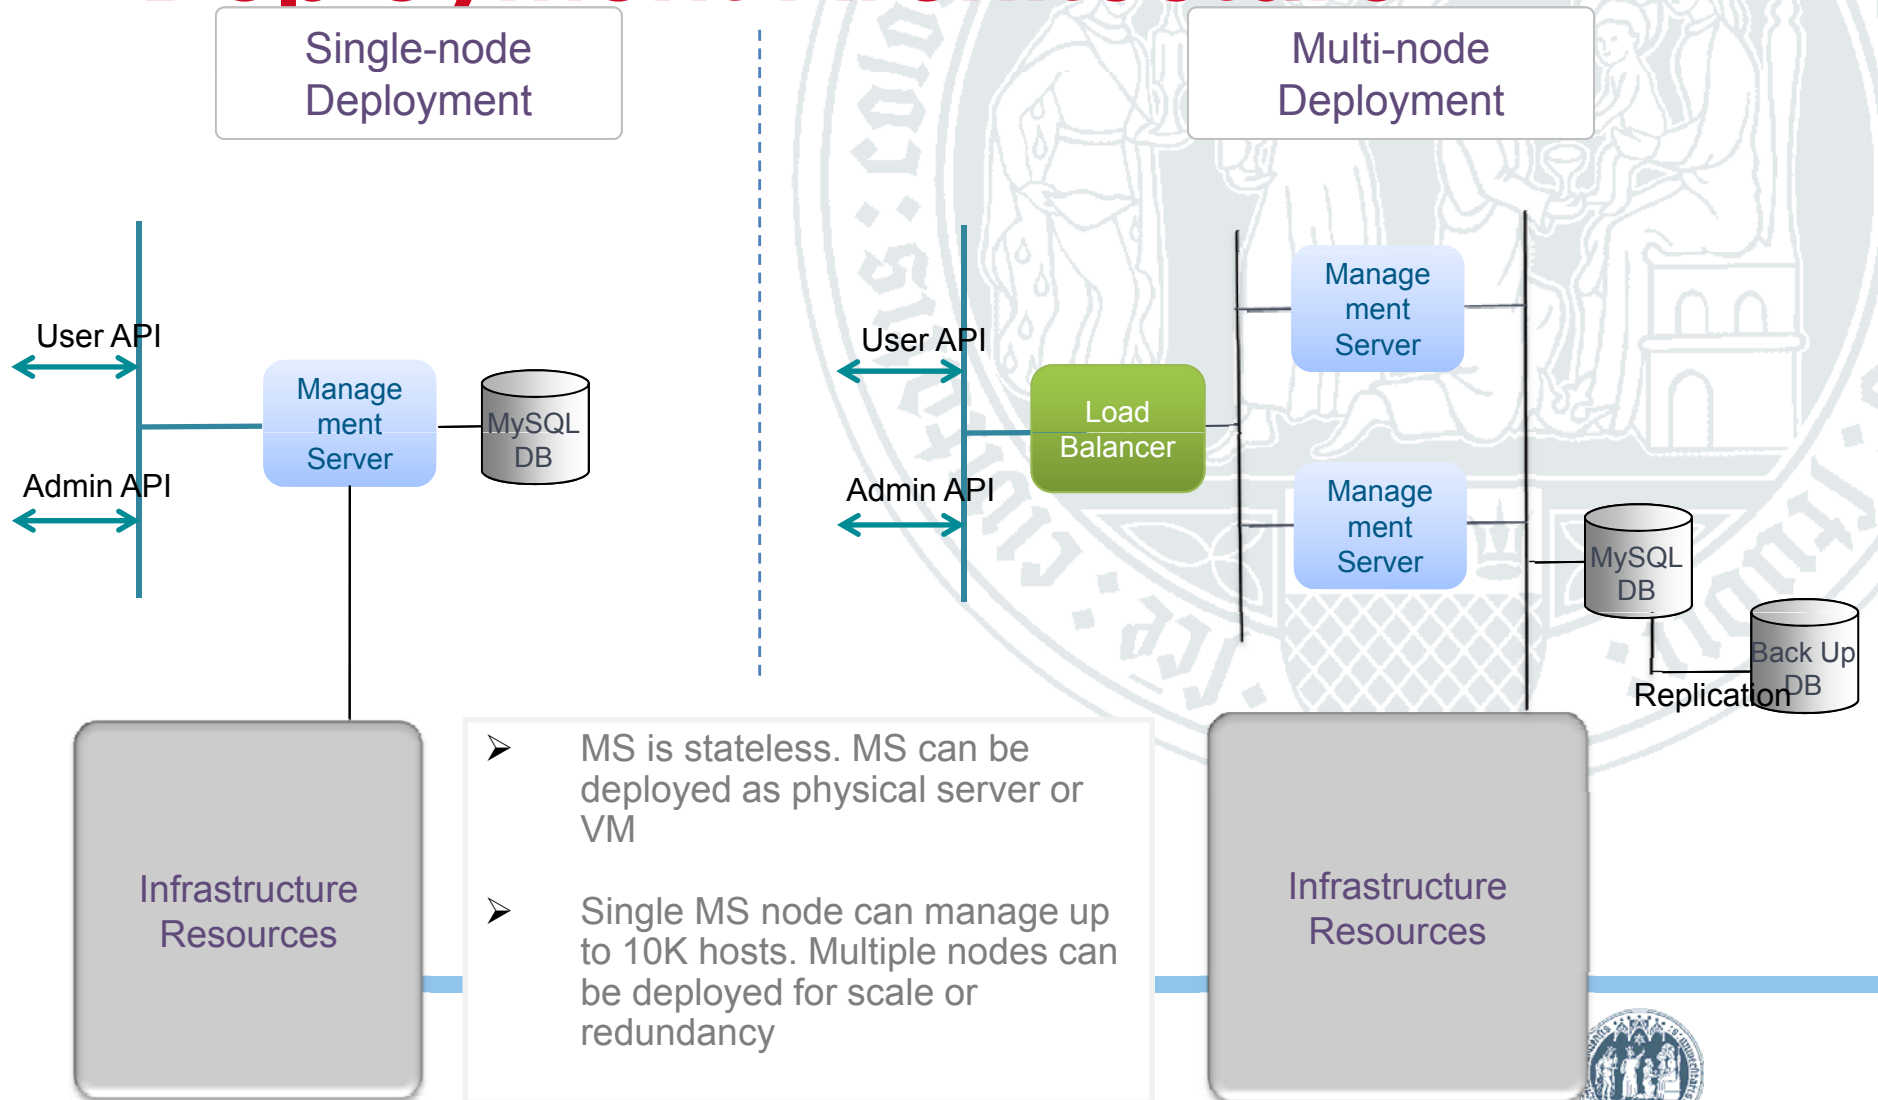
Management Server Managing Multiple Zones



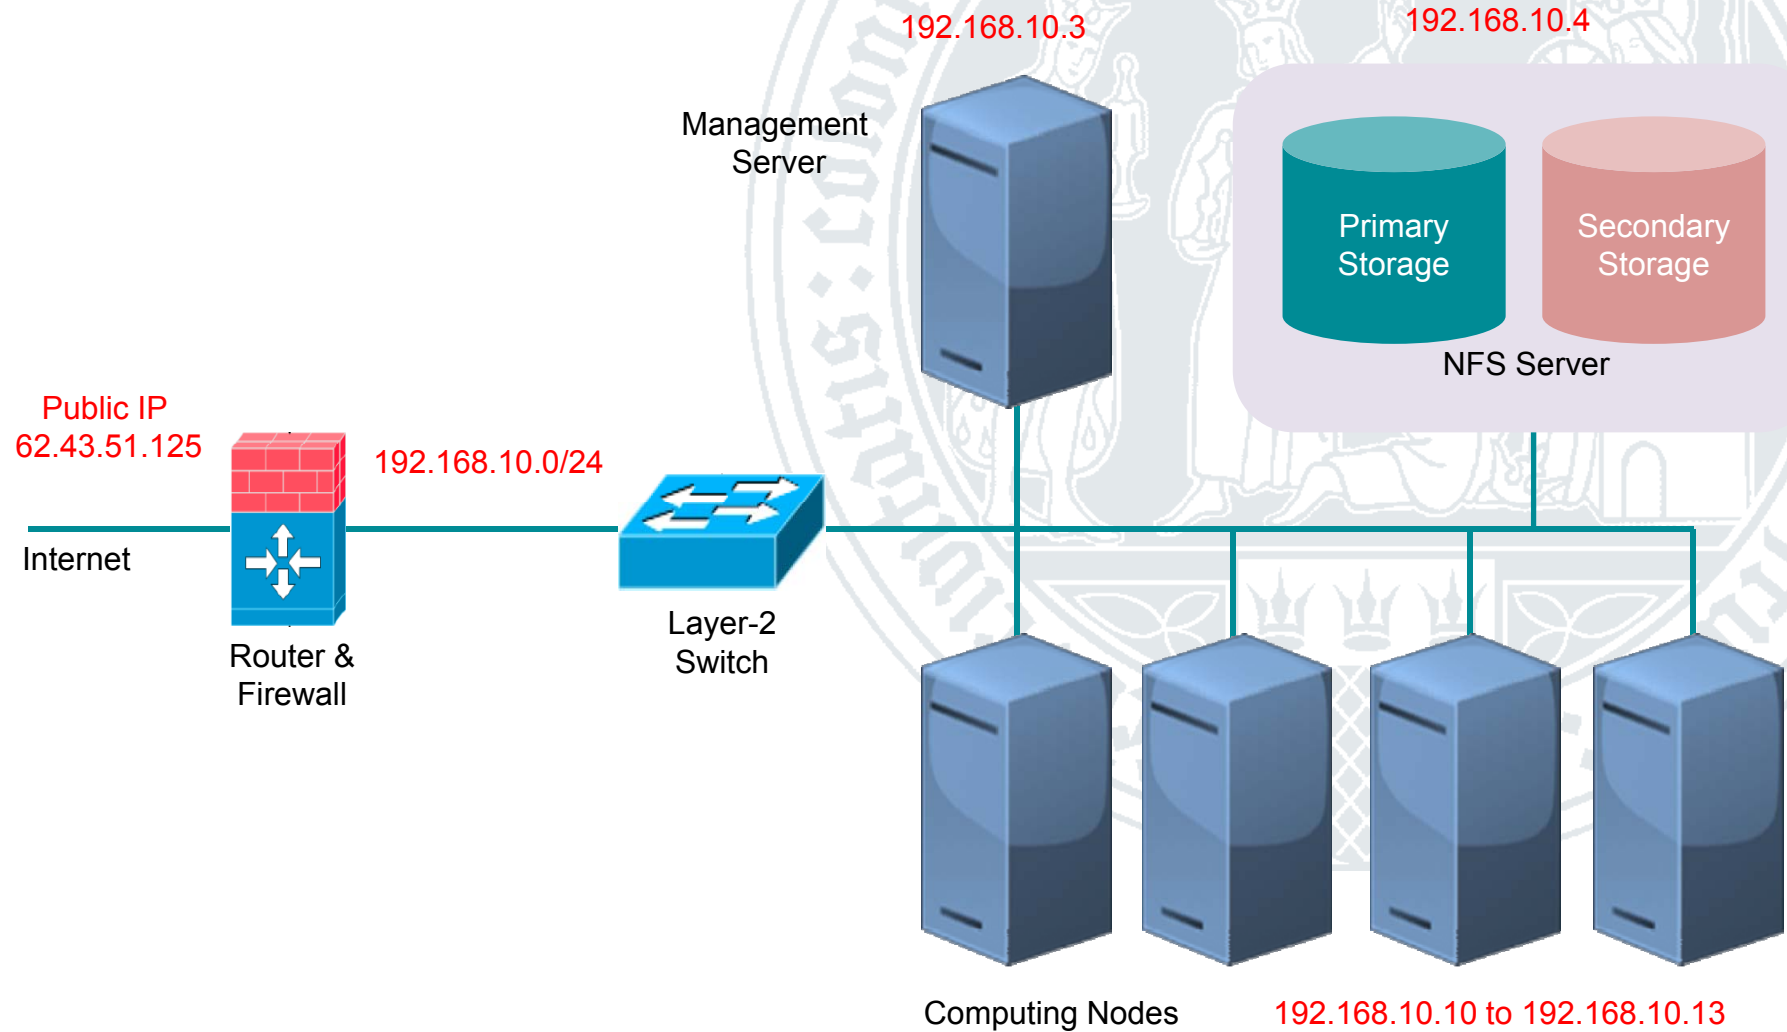
- Single Management Server can manage multiple zones
- Zones can be geographically distributed but low latency links are expected for better performance
- Single MS node can manage up to 10K physical hosts.
- Multiple MS nodes can be deployed as cluster for scale or redundancy



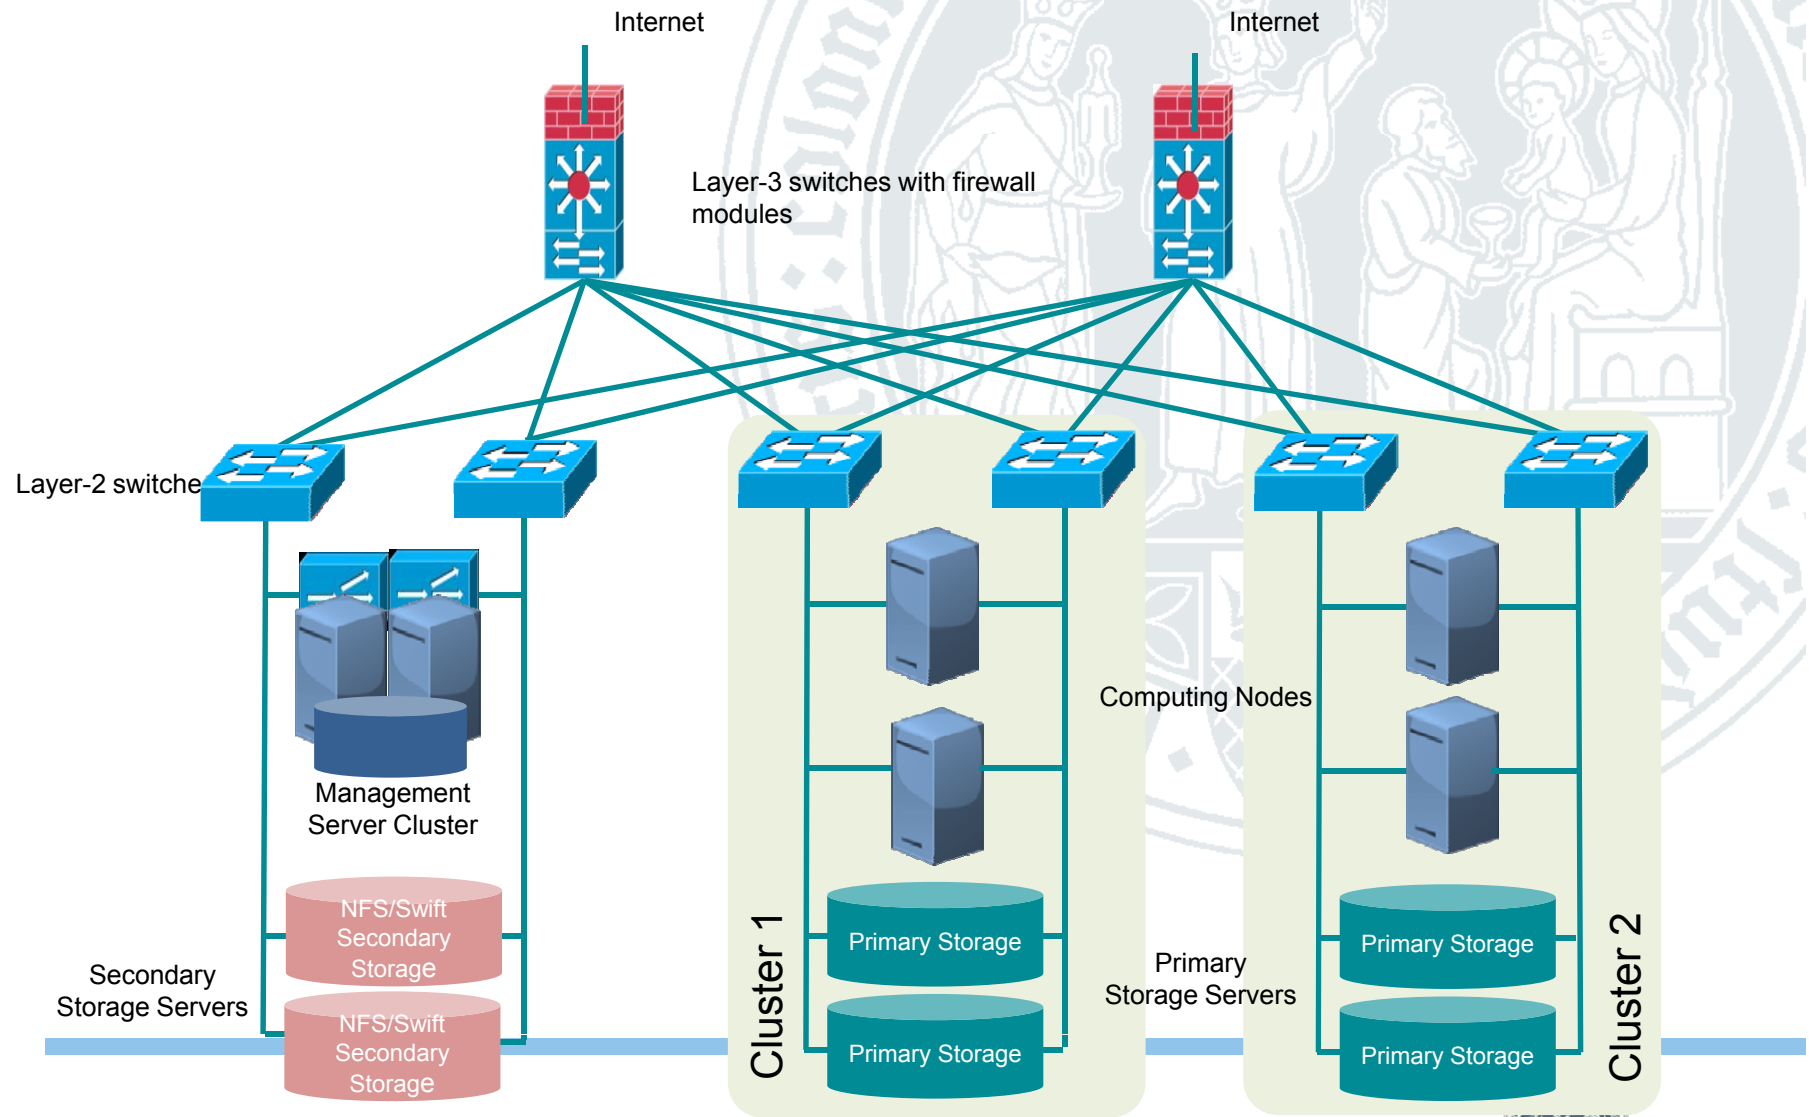
Management Server Deployment Architecture



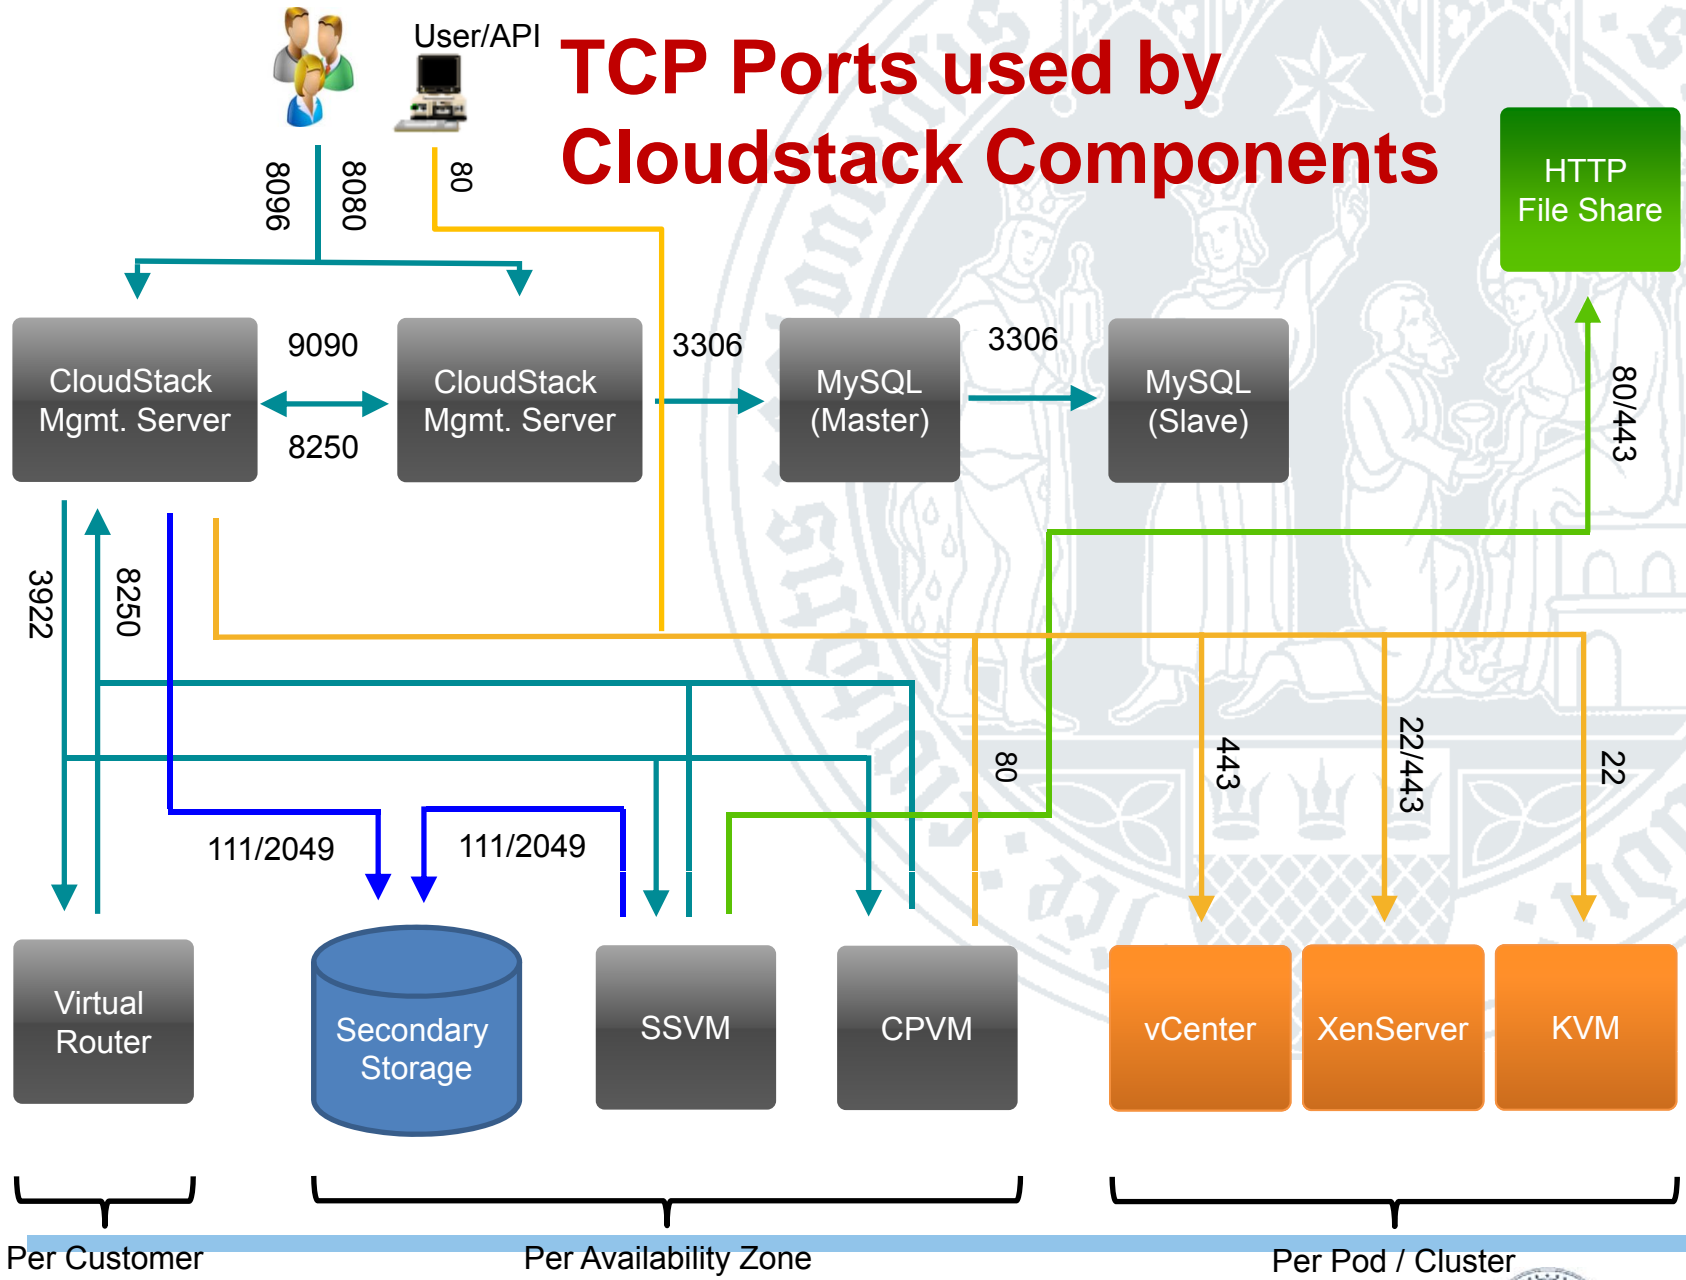
Small-Scale Deployment



Large-Scale Redundant Deployment



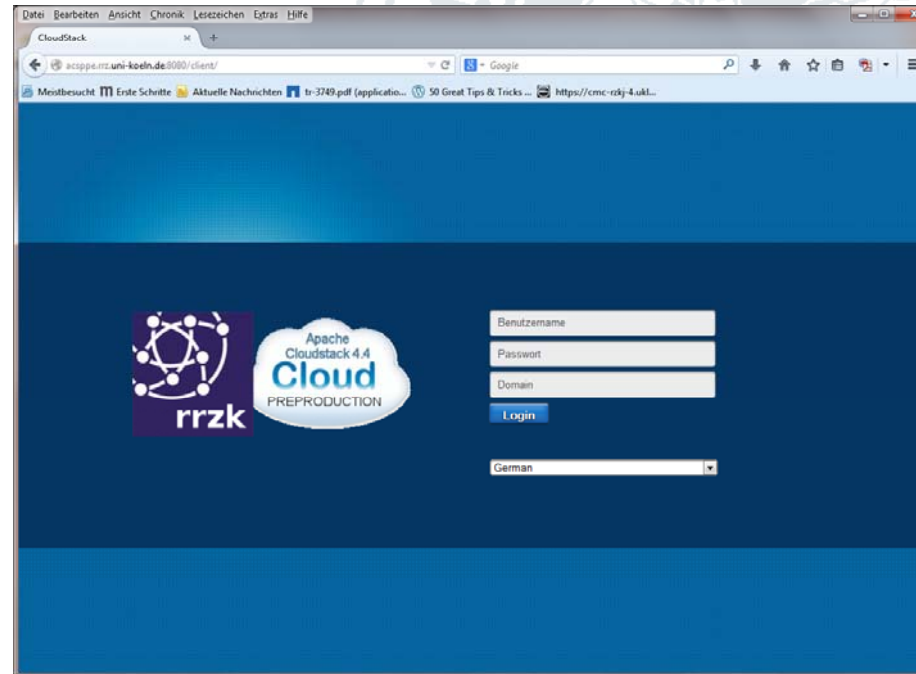
TCP Ports used by Cloudstack Components



Time for a „Reward“



Demo in Preproduction Environment @RRZK



- **deploy W12KR2 and Centos6.5 Guest**
- **To be continued...**

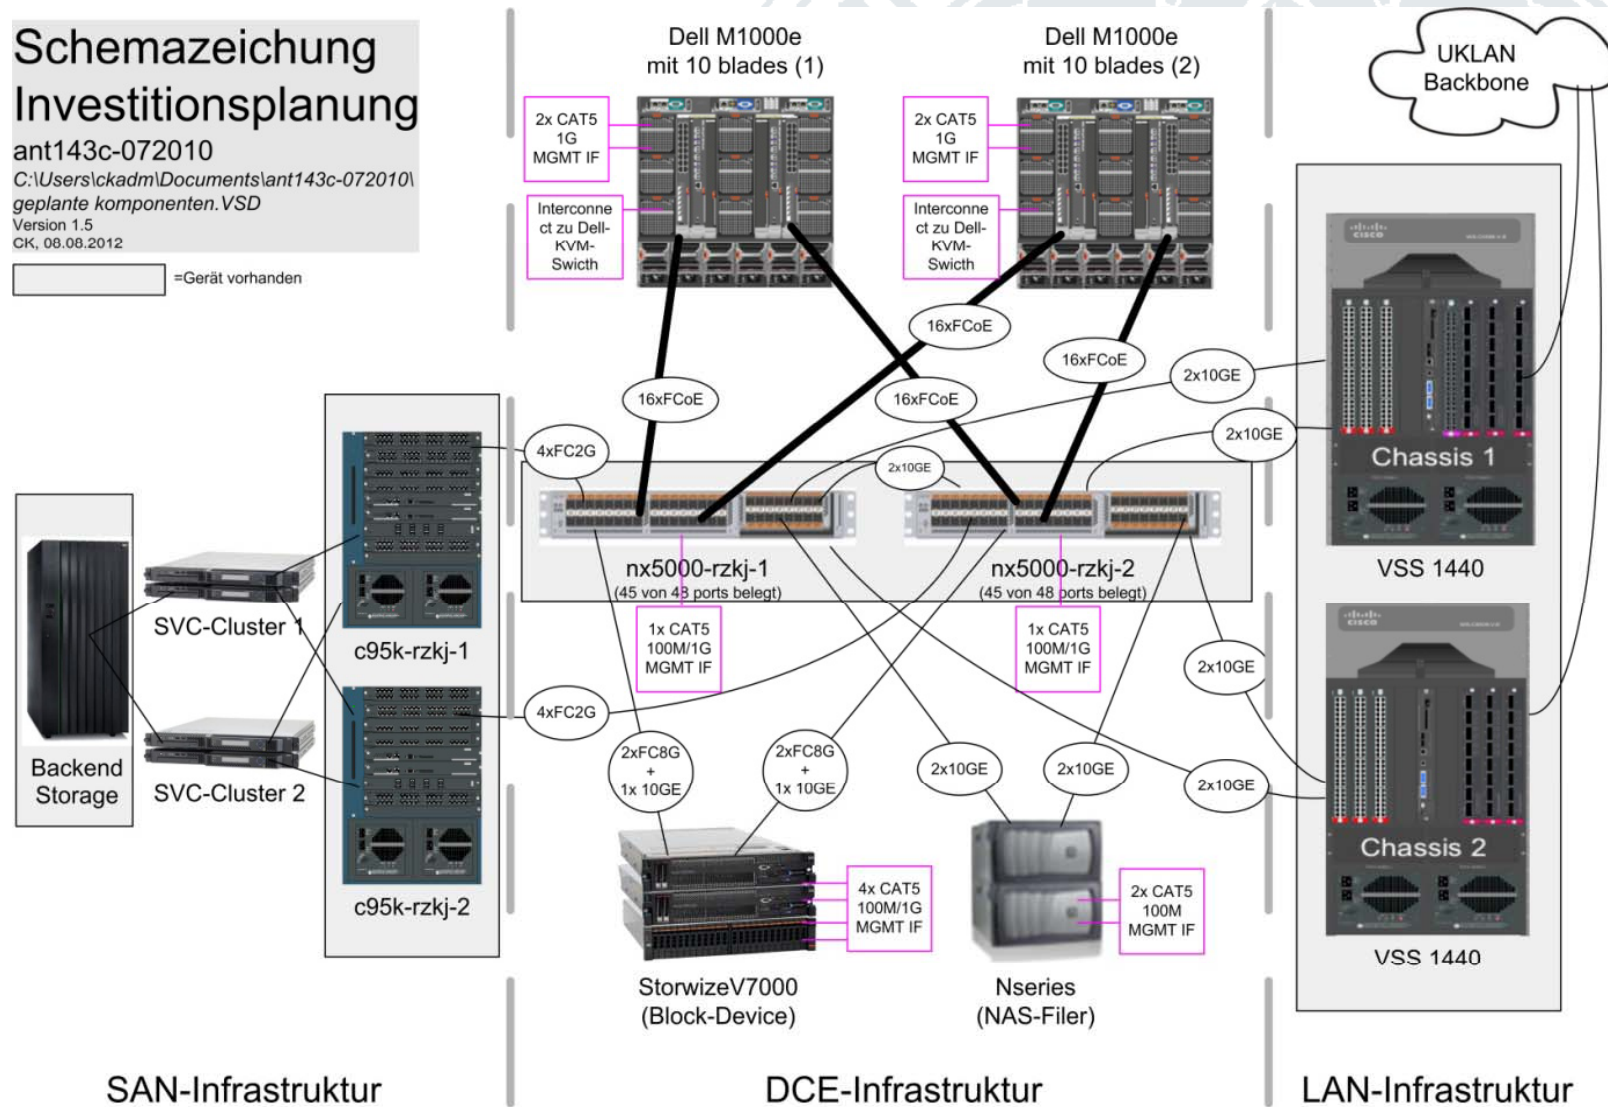
Schemazeichnung Investitionsplanung

ant143c-072010

C:\Users\ckadm\Documents\ant143c-072010\geplante komponenten.VSD

Version 1.5
CK, 08.08.2012

=Gerät vorhanden



Cloudstack technical Features

- **Guest Migration („vMotion“), Guest Console**
- **Snapshots (cold)**
- **Templates (the >Menu<)**
- **Multi-Hypervisor (KVM, XEN, VMWARE, HYPER-V, ...)**
- **Multistorage (primary - nfs, rbd, clvm, Gluster, secondary – nfs, smb/cifs, S3, swift), ?VSAN?**
- **HA-feature**
- **Firewall, Load-Balancer appliances**
- **Private Networking (VPN)**
- **Appliances (F5, Netscaler, ...)**
- **Plugins (openvswitch, ...), SDN, isolated+shared networks**



Cloudstack marketing features

- **Multi-tenant**
- **Multi-Hypervisor**
- **Well documented, release process**
- **API (special and Amazon compatible)**
- **Dashboard**
- **Roadmaps available**
- **Support available (in US and GB)**




General CloudStack Best Practices (!)

- A staging system that models the production environment is strongly advised. It is critical if customizations have been applied to CloudStack.
- Allow adequate time for installation, Proof-of-Concept (POC), and learning the product. Allow at least 4-8 weeks for a POC to work through all of the integration issues. It takes months to gain confidence with CloudStack and related technologies.



Cloudstack Network Principle

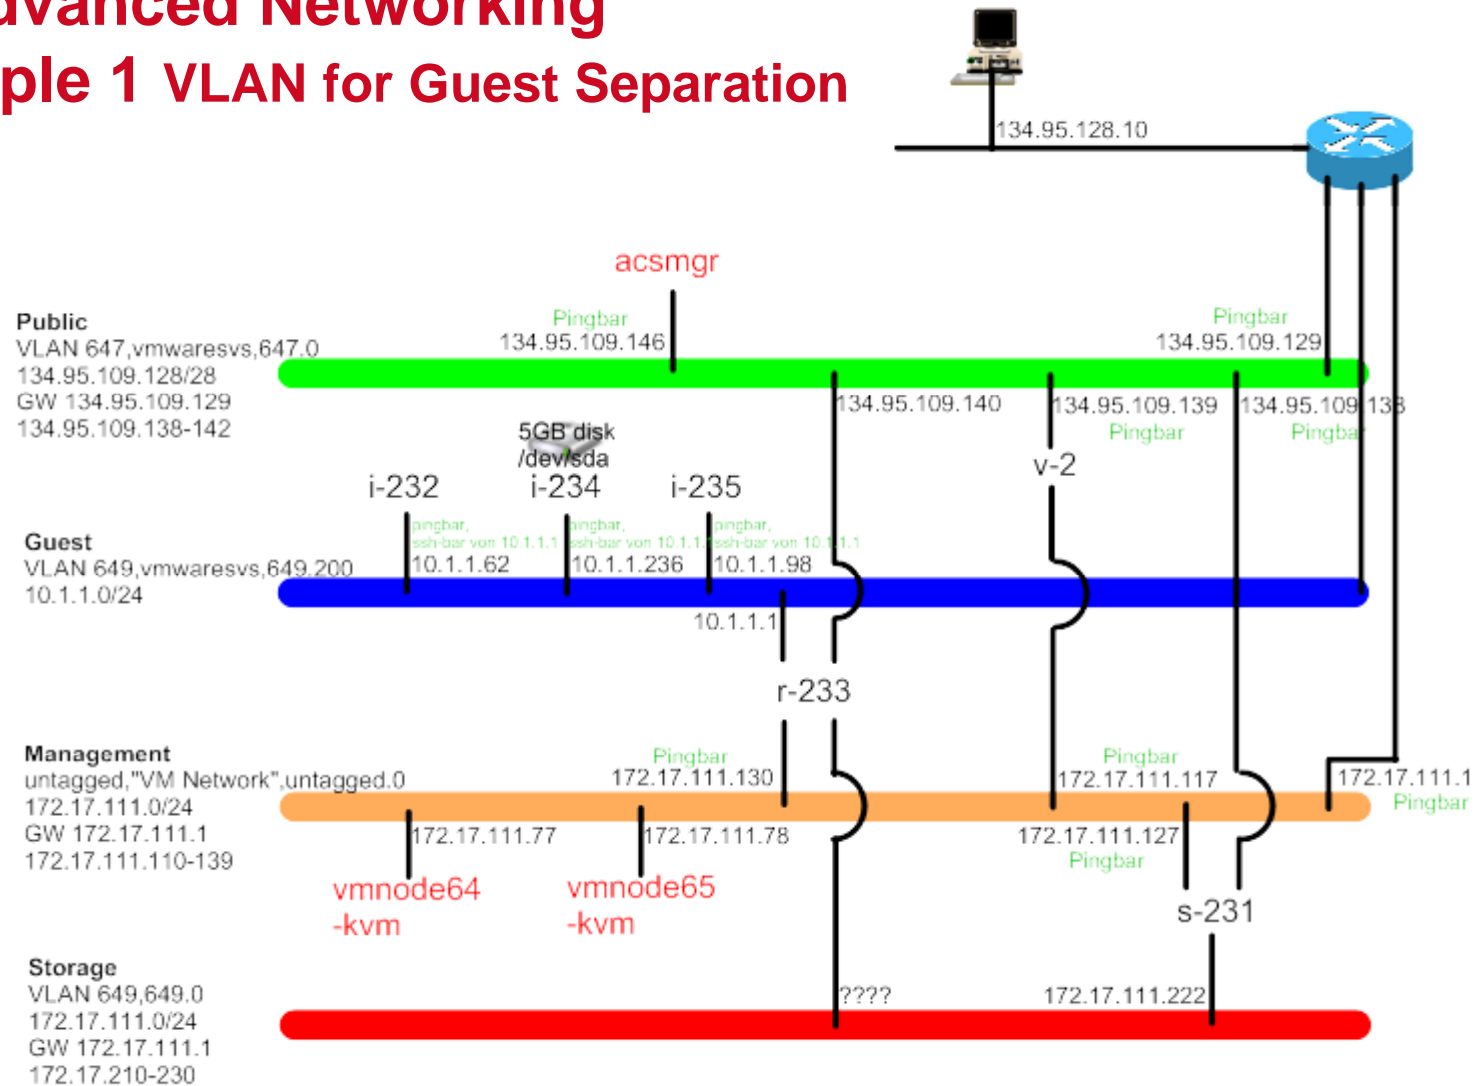


Cloudstack Networking

- **Basic: use ebttables to seperate guest traffic on same broadcast domain**
- **Advanced: use VLANs or VNIs to seperate guest traffic**
 - VLANs are rare&precious (network dept.)
 - VNIs are disposable good (>4096)
 - VNI Usage requires tunnels (GRE, VXLAN)

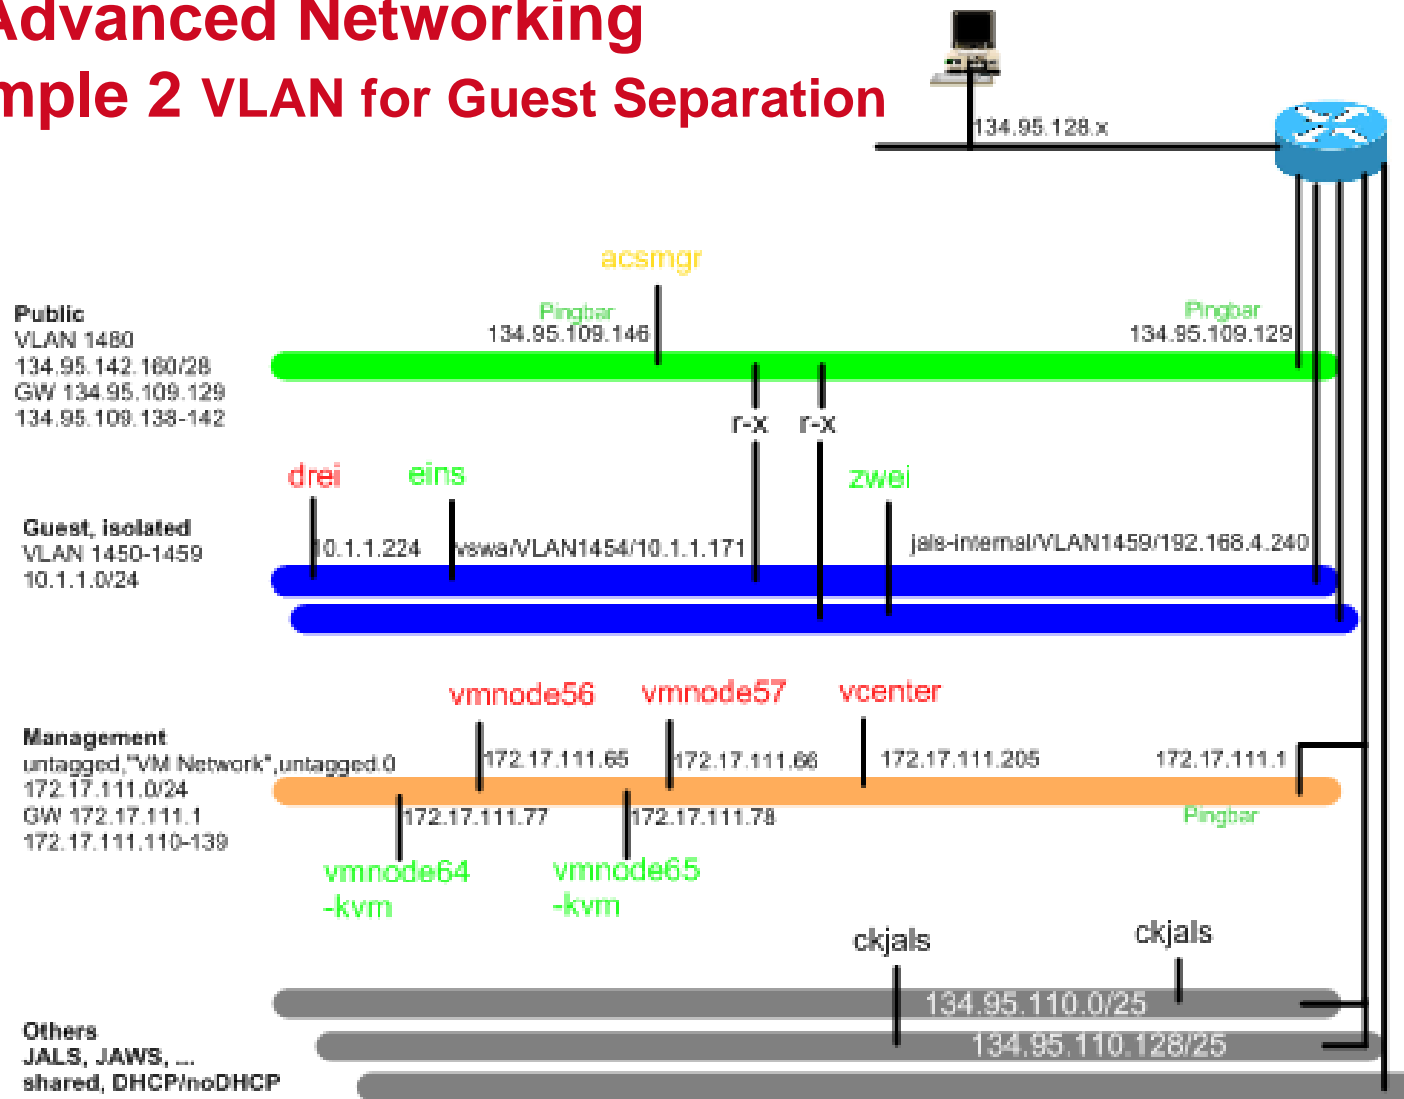
CS Advanced Networking

Example 1 VLAN for Guest Separation



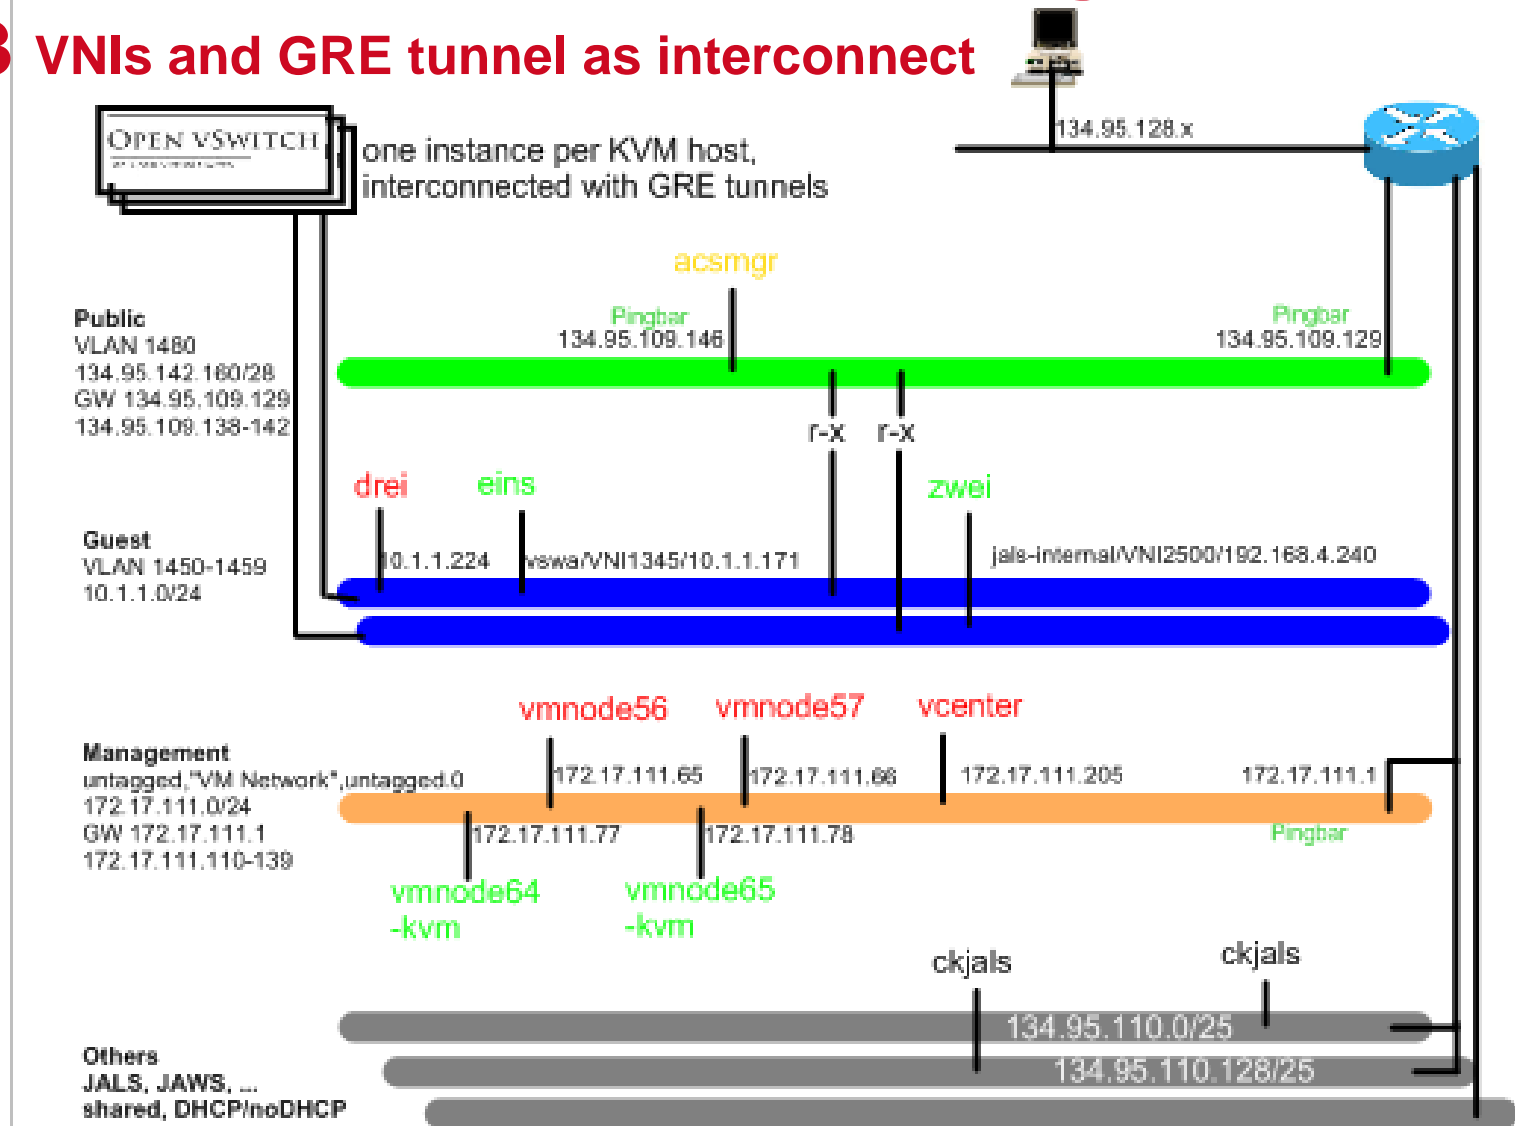
CS Advanced Networking

Example 2 VLAN for Guest Separation



Cloudstack Software Defined Networking (SDN)

Example 3 VNIs and GRE tunnel as interconnect



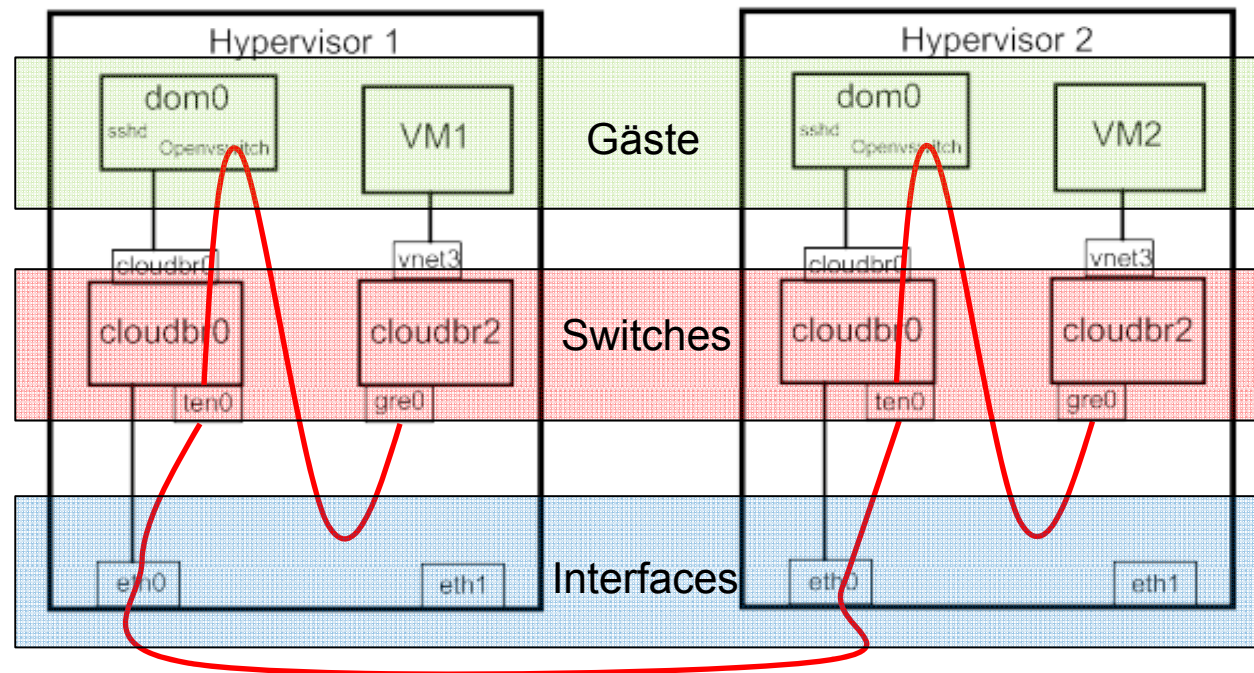
Break

- **Life Demo: How about our guests?**



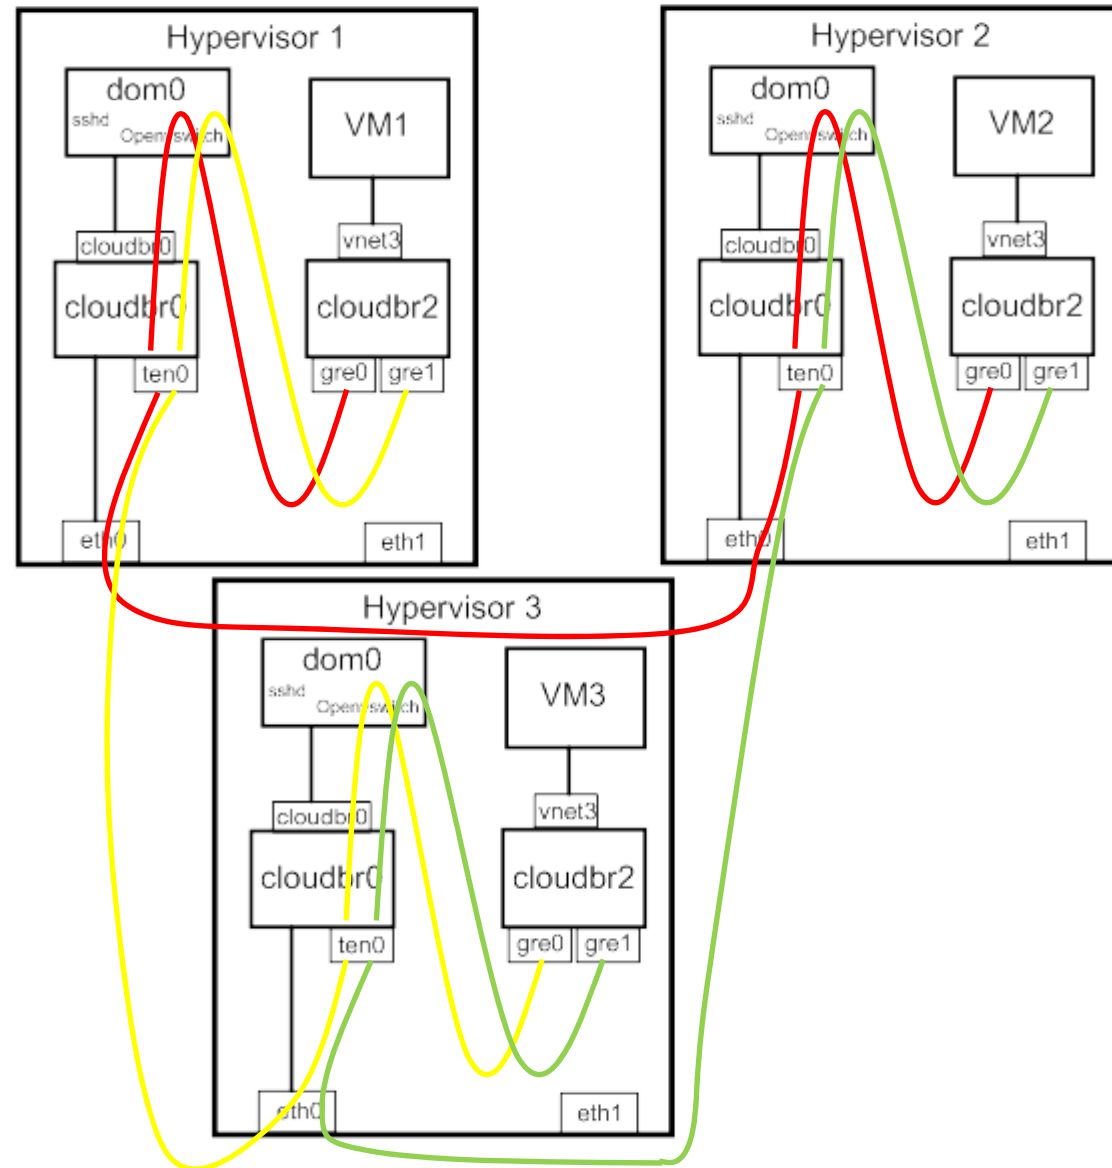
KVM Host Setup for GRE Tunnel

- As of Cloudstack 4.3 (April 2014)



GRE with 3 hosts

- **$N \times (N-1)$** gre interfaces
- **Automatic** in Cloudstack4.4



Agenda

- Cloud
- VMware
- OpenSource: Apache Cloudstack
 - Functions
 - Network-Options
- **Functional Comparison**



NIST Cloud „essential characteristics“

- ***On-demand self-service***
- ***Broad network access***
- ***Resource pooling***
- ***Rapid elasticity***
- ***Measured service***

Quelle: NIST Special Publication 800-145, *NIST Definition of Cloud Computing*, September 2011



Funktioneeller Vergleich

*** 3 NIST-Points ***

VMware

- vMotion
- HA
- Heterogene Gäste
- Import exist. VLAN-Struktur
- Distributed vSwitch
- NSX: virtual Networking
- LDAP-Integration
- Templates
- **MISSING:**
**easy-to-use Self-Service
Function, Accounting**


***** 5 NIST-Points *****

Apache Cloudstack /kvm

- Life migration
- HA
- Heterogene Gäste
- Import exist. VLAN-Struktur
- Standard virt. bridge support
- Openvswitch mit GRE-Tunneln
- LDAP-Integration
- Templates
- **Self-Service Interface**
- **Load-Balancer, Firewall
Appliances**
- **Accounting**



Cloudstack Bibliografie



[Larger Cover](#)

60 Recipes for Apache CloudStack Using the CloudStack Ecosystem

By [Sébastien Goasguen](#)

Publisher: O'Reilly Media

Final Release Date: October 2014

Pages: 156



[Read 3 Reviews](#) | [Write a Review](#)

Planning to deploy and maintain a public, private, or hybrid cloud service? This cookbook's handy how-to recipes help you quickly learn and install Apache CloudStack, along with several API clients, API wrappers, data architectures, and configuration management technologies that work as part of CloudStack's ecosystem.

[Full description](#)

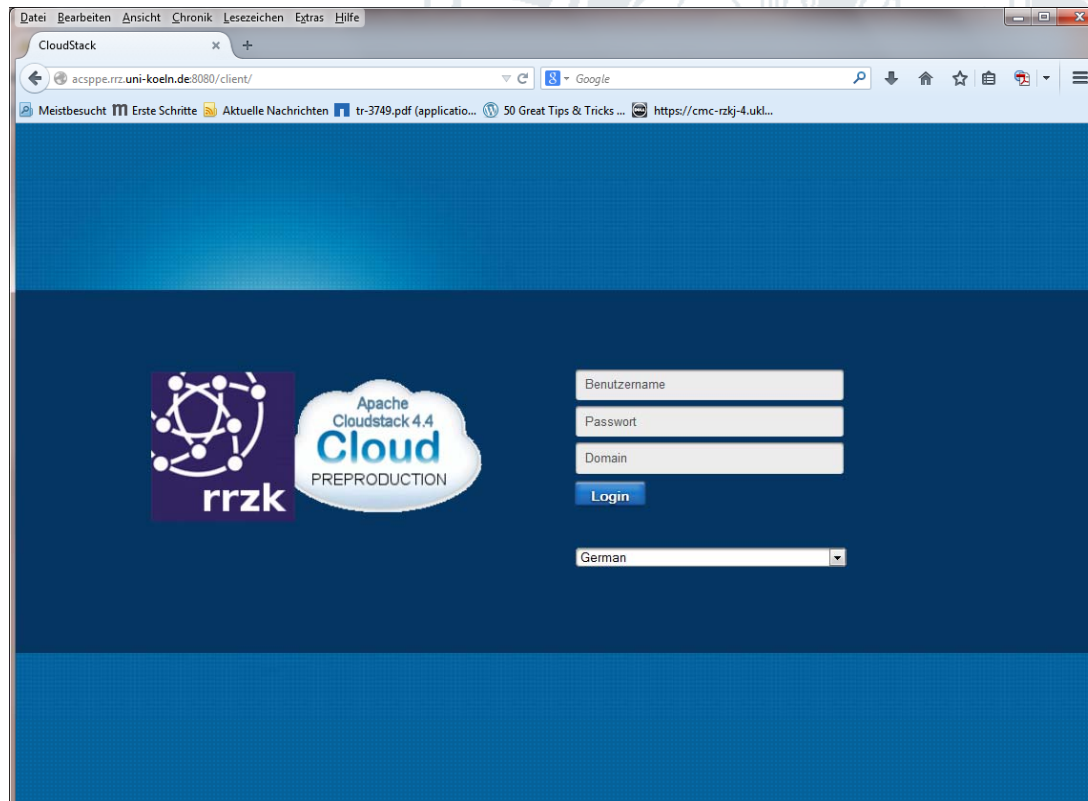
Questions?

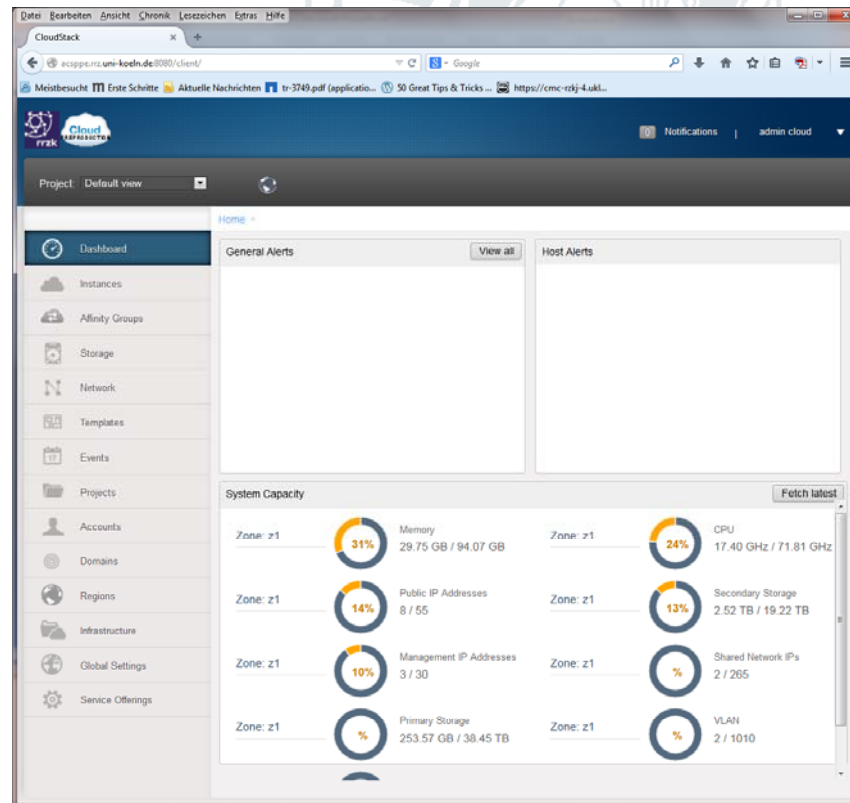


Apache CloudStack

Cloudstack Logos published by Erdösi Péter <fazy@niif.hu>







The screenshot shows the CloudStack web interface. The browser address bar indicates the URL is `acspp.rz.uni-koeln.de:8080/client/`. The page title is "CloudStack". The navigation sidebar on the left includes options like Dashboard, Instances, Affinity Groups, Storage, Network, Templates, Events, Projects, Accounts, Domains, Regions, Infrastructure, Global Settings, and Service Offerings. The "Instances" section is active, displaying a table of virtual machines.

Name	Internal name	Display name	Account Name	Host	State	Quickview
w12-sc1	i-2-32-VM	w12-sc1	admin	vmnode13.rz...	Running	+
w12dc-1	i-2-31-VM	w12dc-1	admin	vmnode12.rz...	Running	+
ck1	i-2-41-VM	ck1	admin	vmnode13.rz...	Running	+
plagsniffervm	i-2-19-VM	plagsniffervm	mboenigk	vmnode12.rz...	Running	+
ck-filer	i-2-18-VM	ck-filer	admin	vmnode12.rz...	Running	+
freeradius3-1	i-9-34-VM	freeradius3-1	a0620	vmnode14.rz...	Running	+
w2k8neu	i-2-23-VM	w2k8neu	admin	vmnode14.rz...	Running	+
ck65	i-2-33-VM	ck65	admin	vmnode13.rz...	Running	+



