de.NBI

Cloud computing

General terms and concepts

Jan Krüger
14.03.2018

SPONSORED BY THE
Federal Ministry of Education and Research
Projektträger Jülich Forschungszentrum Jülich

Cloud… ?

Cloud… ?

Cloud… ?

Cloud !

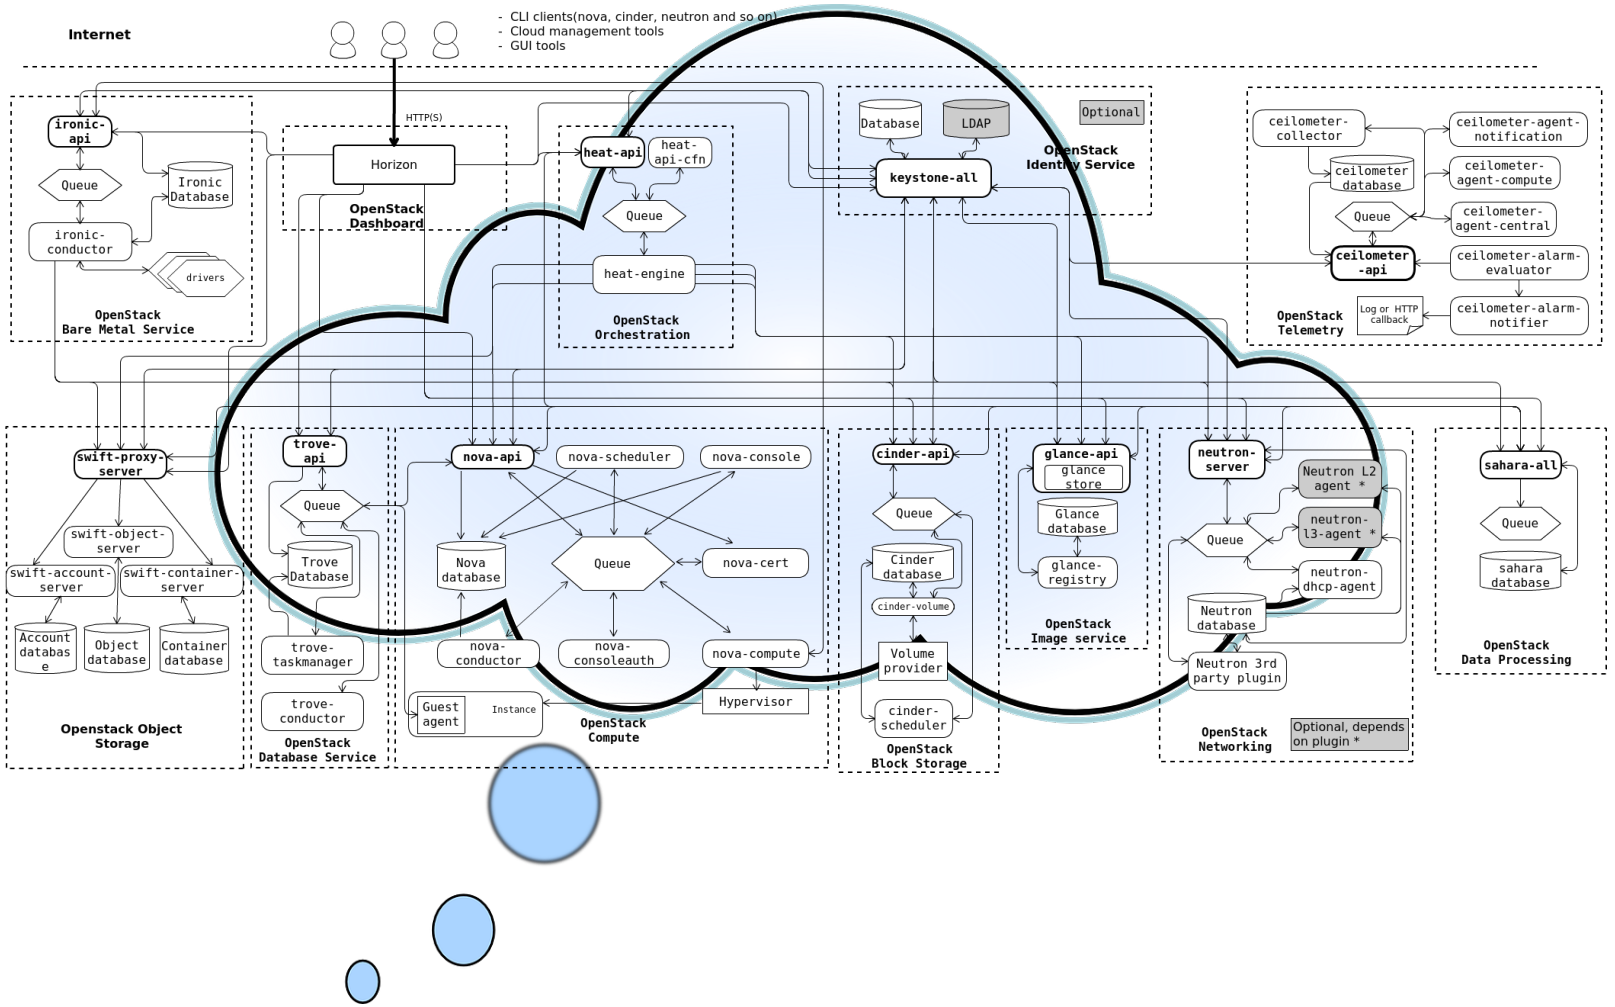
Internet
- CLI clients(nova, cinder, neutron and so on)
- Cloud management tools
- GUI tools
HTTP(S)
Horizon
OpenStack Dashboard
ironic-api
Queue
Ironic Database
ironic-conductor
drivers
OpenStack Bare Metal Service
heat-api
heat-api-cfn
Queue
heat-engine
OpenStack Orchestration
Database
LDAP
Optional
keystone-all
OpenStack Identity Service
ceilometer-collector
ceilometer-agent-notification
ceilometer database
ceilometer-agent-compute
Queue
ceilometer-agent-central
ceilometer-api
ceilometer-alarm-evaluator
Log or HTTP callback
ceilometer-alarm-notifier
OpenStack Telemetry
swift-proxy-server
swift-object-server
swift-account-server
swift-container-server
Account database
Object database
Container database
Openstack Object Storage
trove-api
Queue
Trove Database
trove-taskmanager
trove-conductor
OpenStack Database Service
nova-api
nova-scheduler
nova-console
Nova database
Queue
nova-cert
nova-conductor
nova-consoleauth
nova-compute
Guest agent
Instance
Hypervisor
OpenStack Compute
cinder-api
Queue
Cinder database
cinder-volume
Volume provider
cinder-scheduler
OpenStack Block Storage
glance-api
glance store
Glance database
glance-registry
OpenStack Image service
neutron-server
Neutron L2 agent *
Queue
neutron-l3-agent *
neutron-dhcp-agent
Neutron database
Neutron 3rd party plugin
OpenStack Networking
Optional, depends on plugin *
sahara-all
Queue
sahara database
OpenStack Data Processing

Cloud !

Scalable
(up and out, on-demand, re-deployment)

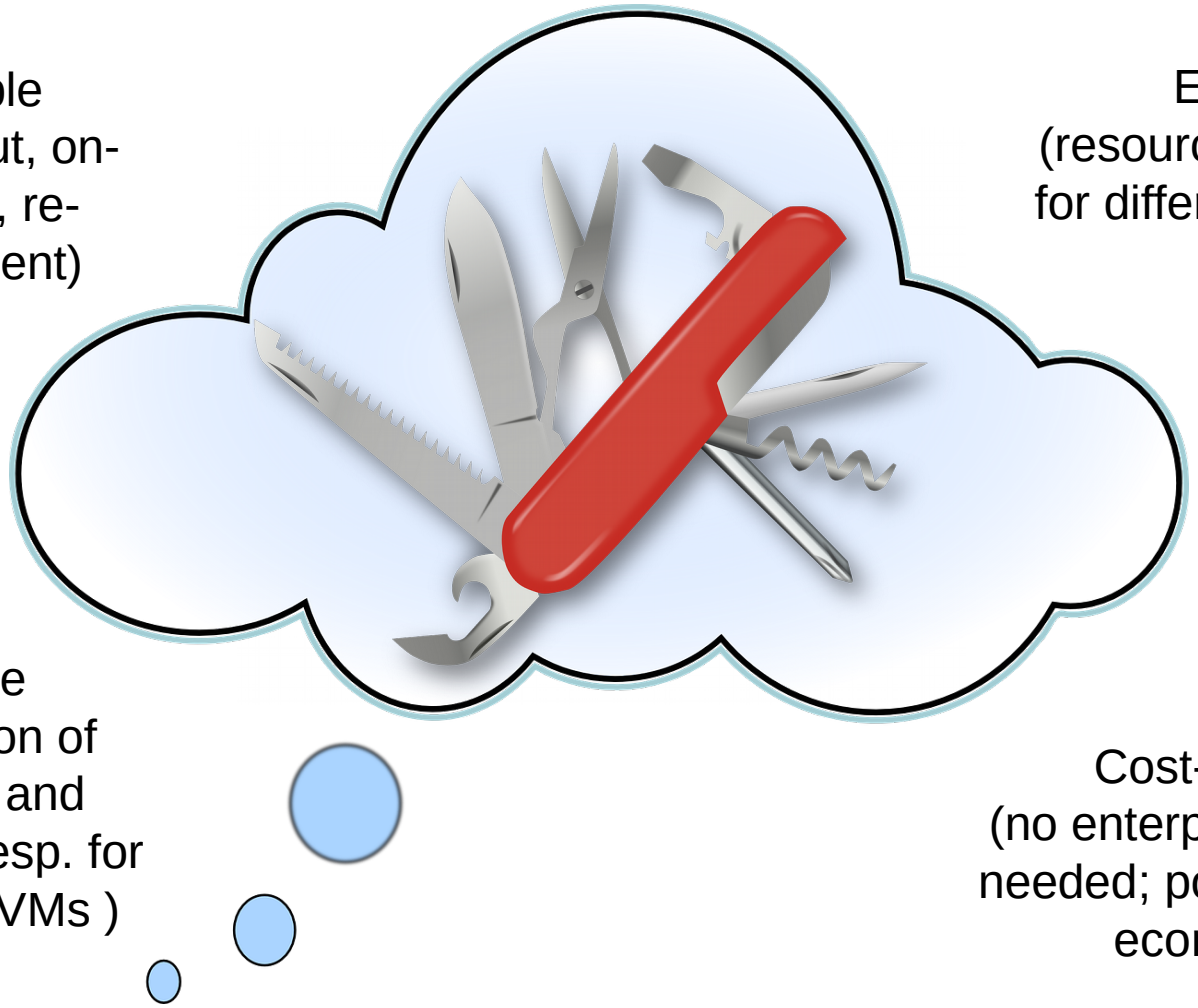

Efficient
(resources re-usable for different services)

Flexible
(separation of services and hardware; esp. for research: VMs)

Cost-Effective
(no enterprise hardware needed; pooled resource economics)

Cloud Computing

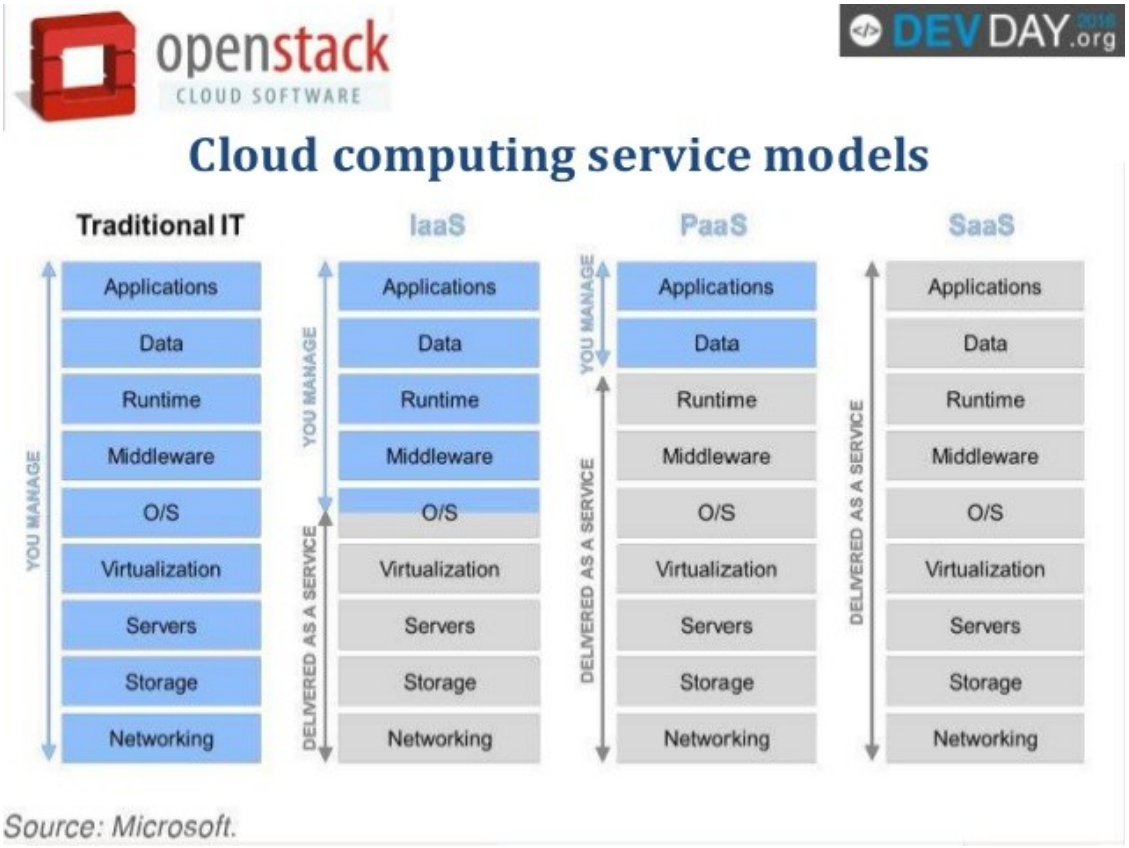

Infrastructure as a Service (IaaS)
→ most basic cloud service
→ provision of VMs, IPs, DBs, Firewalls, ...

Platform as a Service (PaaS)
→ Simplifies Cloud access
→ Provision of a OS, APIs, Web Server, execution & development environments ...

Software as a Service (SaaS)
→ Access via clients on computers, mobile devices, browsers and etc.
→ Provision of applications and services (often installed + operated by the cloud provides)

Cloud Computing

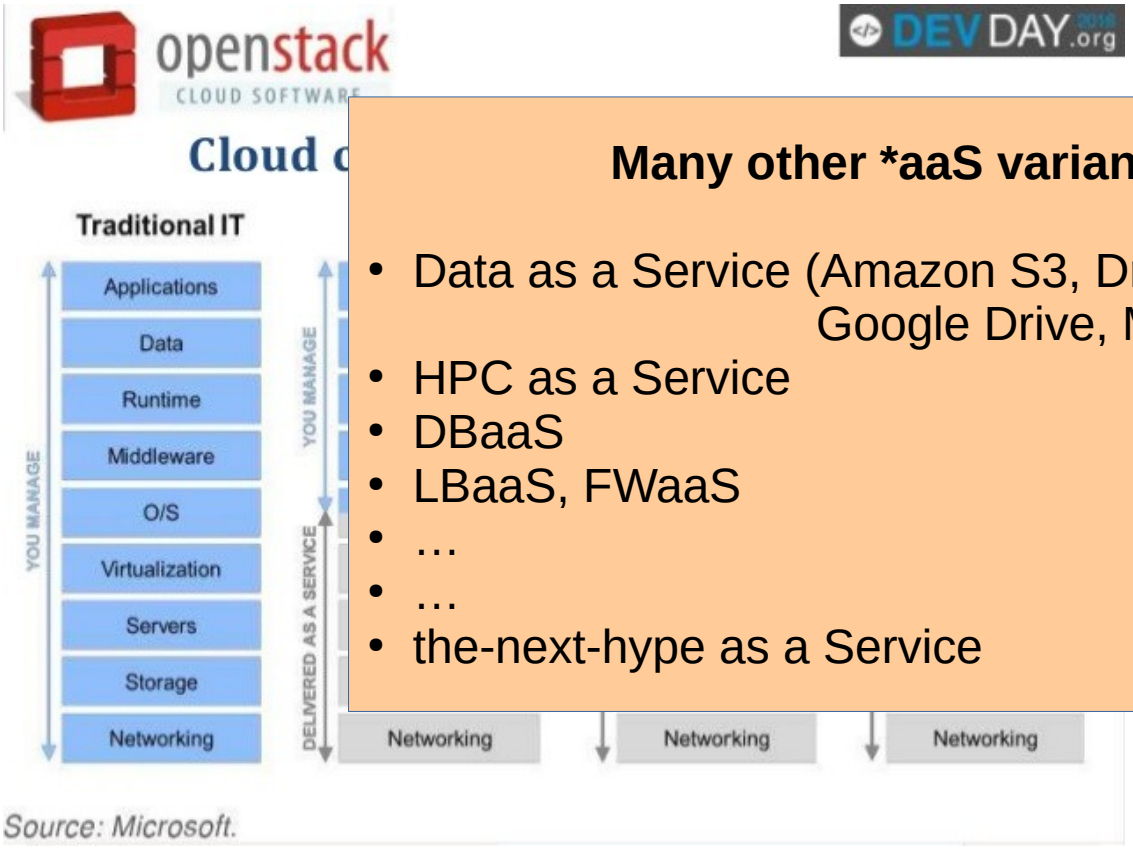

Source: Microsoft.

DEV DAY .org

Infrastructure as a Service (IaaS)
→ most basic cloud service
→ provision of VMs, IPs, DBs,

Many other *aaS variants around:

- Data as a Service (Amazon S3, Dropbox, Google Drive, Microsoft OneDrive)
- HPC as a Service
- DBaaS
- LBaaS, FWaaS
- …
- …
- the-next-hype as a Service

aaS)
ss
s, Web
lopment

aaS)
computers,
s and etc.
→ Provision of applications and services (often installed + operated by the cloud provides)

Virtualization

- Simulating a certain environment / resources
- In case of “virtual machines”:
  - BIOS
  - CPU
  - RAM
  - IO devices (network, disk, serial I/O, display)

- Overhead (especially for I/O)
- Hardware support in modern CPUs

Virtualization

- Hypervisor
  - Software providing virtualization
    - VMware, Xen, HyperV, VirtualBox, qemu/kvm, …
    - Application or complete operation system
  - Machine a hypervisor software runs on
- May also provide virtual networks / storage

Virtualization vs. Cloud

- Virtualization → Cloud:
  - Virtualization is base for IaaS clouds

- Cloud → Virtualization:
  - Cloud software manages hypervisors
  - Configuration of networks, storage, virtual machines
  - Actual work done by hypervisors

Cloud is great…

- Highly flexible
- Scalable on demand
- User friendly
- “Emancipation” of users
- Widely adopted and accepted
- Public clouds / private clouds / hybrid clouds

...but...

Cloud is great, but...

https://apod.nasa.gov/apod/image/1011/thundercell_heavey_big.jpg

Cloud is great, but...

heise Security News Hintergrund Foren Events

Security > News > 7-Tage-News > 2017 > KW 2 > Anzahl entführter MongoDB-

Anzahl entführter MongoDB-Datenbanken st rasant

10.01.2017 10:46 Uhr – Dennis Schirrmacher

(Bild: Bleepingcomputer.com)

Sicherheitsforscher warnen, dass mittlerweile verschiedene Erpressergruppen in größerem Stil offene MongoDB-Datenba und Lösegeld verlangen.

Die Angriffe auf nicht sicher konfigurierte MongoDB-Datenbanken ab und nehmen sogar noch zu, warnen verschiedene Sicherheitsf gegenüber Bleepingcomputer.com. Derzeit sollen über 28.000 Ser worden sein.

Google Security Blog

The latest news and insights from Google on security and safety on the Internet

CVE-2015-7547: glibc getaddrinfo sta

February 16, 2016

Posted by Fermin J. Serna, Staff Security Engineer and Kevin St

Have you ever been deep in the mines of debugging were staring at something far more interesting than alone! Recently a Google engineer noticed that their they tried to connect to a specific host. That engine behavior and after an intense investigation we disco in SSH as we were expecting.

Thanks to this engineer's keen observation, we were result in remote code execution. We immediately be issue to determine whether it could be exploited, an challenge, and after some intense hacking sessions, exploit!

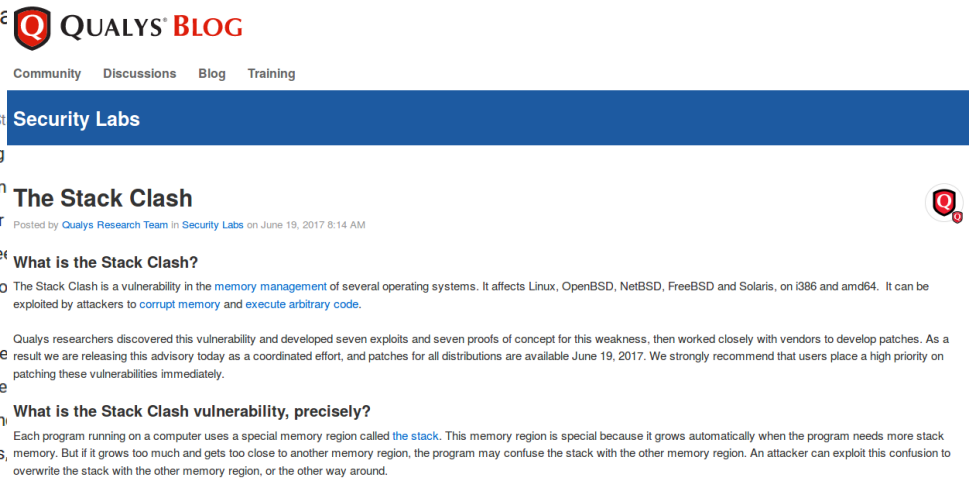

QUALYS BLOG

Community Discussions Blog Training

Security Labs

The Stack Clash

Posted by Qualys Research Team in Security Labs on June 19, 2017 8:14 AM

What is the Stack Clash?

The Stack Clash is a vulnerability in the memory management of several operating systems. It affects Linux, OpenBSD, NetBSD, FreeBSD and Solaris, on i386 and amd64. It can be exploited by attackers to corrupt memory and execute arbitrary code.

Qualys researchers discovered this vulnerability and developed seven exploits and seven proofs of concept for this weakness, then worked closely with vendors to develop patches. As a result we are releasing this advisory today as a coordinated effort, and patches for all distributions are available June 19, 2017. We strongly recommend that users place a high priority on patching these vulnerabilities immediately.

What is the Stack Clash vulnerability, precisely?

Each program running on a computer uses a special memory region called the stack. This memory region is special because it grows automatically when the program needs more stack memory. But if it grows too much and gets too close to another memory region, the program may confuse the stack with the other memory region. An attacker can exploit this confusion to overwrite the stack with the other memory region, or the other way around.

https://www.heise.de/security/meldung/Anzahl-entfuehrter-MongoDB-Datenbanken-steigt-rasant-3592539.html
https://security.googleblog.com/2016/02/cve-2015-7547-glibc-getaddrinfo-stack.html
https://blog.qualys.com/securitylabs/2017/06/19/the-stack-clash

Cloud is great, but…

- Software is not perfect
- New security holes are discovered daily
- Bugs are actively abused
- Systems need to be protected and updated regulary

“Classic” setup:

- Automatic system updates for user machines (desktops)
- System administrator’s duty for servers

How about cloud setups and virtual machines?

Cloud is great, but...

Not the system / cloud administrator's task!

- Not **ALLOWED** to inspect virtual machines
- Inspection technically not feasible
- A zoo of different operating systems, releases, image formats…
- Updates may break applications

...but…

- Monitoring of activities
- Inform user in case of abnormalities
- Terminate virtual machines

(part of the fine print of the cloud's user policy)

Cloud is great, but...

Partly developers / image creators task!

- Users rely on sane images / well developed applications
- Use best practices for security (e.g. no password based login)
- Avoid standard passwords for services
- Ensure passwords are updated by the user

- Requires understanding of infrastructure/operation/administration
  → **DevOps**

- DevOps / automation sessions on tuesday!
- Image hardening on wednesday!

Cloud is great, but...

Also the user's task!

- Users may install their own applications
- Correctly securing a system is difficult
- Usually no automatic updates configured in VM images
- Tutorials often skip security aspect
  - Just get the software up and running
  - Default passwords ("change them later….")
  - Forget about securing...

→ Cloud setup a compromise between

- Freedom (users allowed to do everything / use every application)
- Security (users are restricted)

Cloud is great, but...

de.NBI way:

- Recommended official supported images (Ubuntu LTS), but also allow users to upload & use their own images
- Restrict external network access
  - SSH
  - HTTP/HTTPS
- Work with datacenter on network monitoring

- **EDUCATE users!**


openstack™
CLOUD SOFTWARE

What is OpenStack?

- Open source cloud computing platform (Rackspace and NASA founded)
- Main directives
  - Software to provide virtual machines on standard hardware at massive scale
  - Open source software to build public and private clouds
  - Software to reliably store billions of objects distributed across standard hardware
- Collection of different open source components
- Global software community of developers
- Used by corporations, service providers, researchers, and data centers

OpenStack Components

- Keystone – identity / user management
- Nova – compute & virtualization
- Glance – image repository
- Cinder – block storage / volumes
- Neutron – networking
- Swift – object storage
- Horizon – dashboard / user interface

(openstack.org currently lists 46 components / sub projects…..)

OpenStack at one glance

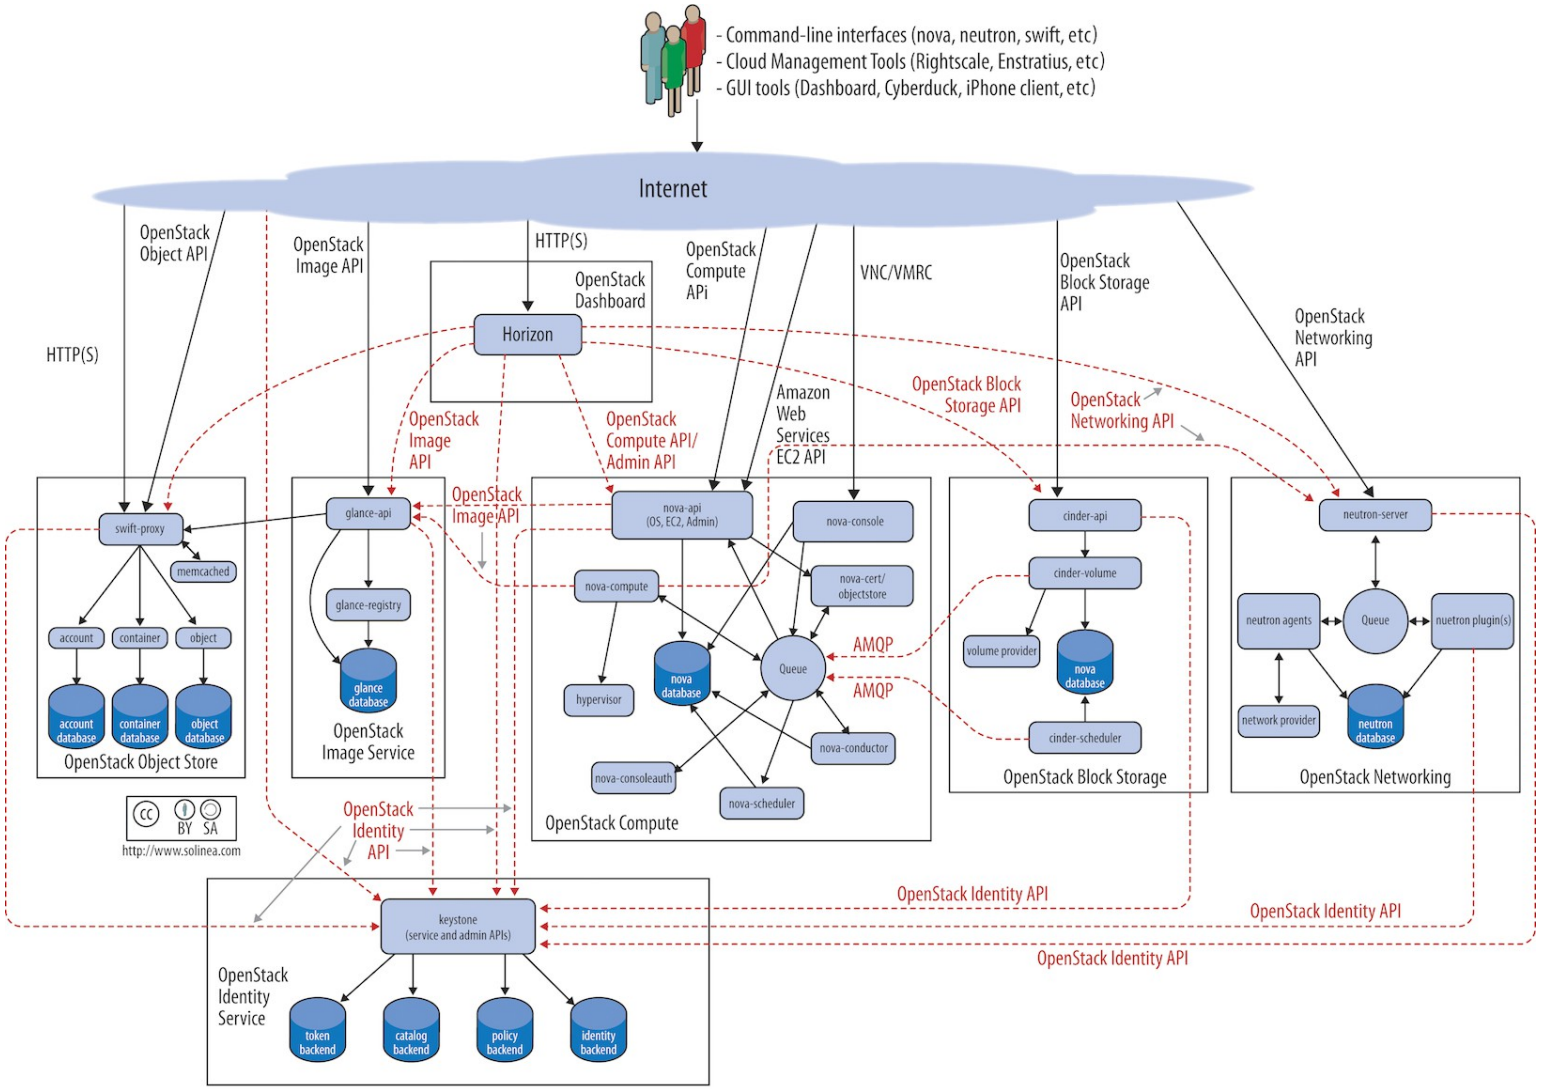

OpenStack: Nova

- Execution of compute workloads (compute controller)
  - Scheduler places instances on compute hosts
    (respecting filter rules, e.g. CPU, RAM, Disk, ...)
- No virtualization by itself (libvirt API interaction with hypervisors)
  - Hypervisor agnostic
  - Support for Libvirt (KVM, QEMU, Xen, LXC), XenAPI, Hyper-V,
    VMware ESX, PowerVM, Docker, Bare-metal, ...
- Overcommit of RAM and CPUs

OpenStack: Glance

- Provides registration & delivery service for virtual disk images
- Highly available (write-once, read-many storage)
- Possible to run independent as IaaS (Image as a Service)
- Image copied on use by Nova
- Format agnostic
  - e.g. raw, qcow2, VHD, AMI, VDI
  - Image verification (integrity)
- Metadata properties
  - e.g. specify virtual hardware preferences

OpenStack: Neutron

- Network infrastructure management
- Allows configuration and deployment of
  - Networks
  - Routers
  - Sub-nets
  - Ports
- Each instance bound to one security group
- Security Group: Collection of network access rules
- Access rules
  - Control network traffic from/to to all VM instances in this group
  - Can be modified at any time
  - Automatically enforced to all running instances

OpenStack: Cinder

- Storage in instances is non-persistent
  - lost when instance terminated
- Provides persistent block storage
  - API compatible with Amazon’s Elastic Block Store
- Volumes are accessed via iSCSI
- Attached uniquely (only one instance)
- Supports multiple storage backends
  - LVM, Gluster, Ceph, ...

OpenStack: Horizon

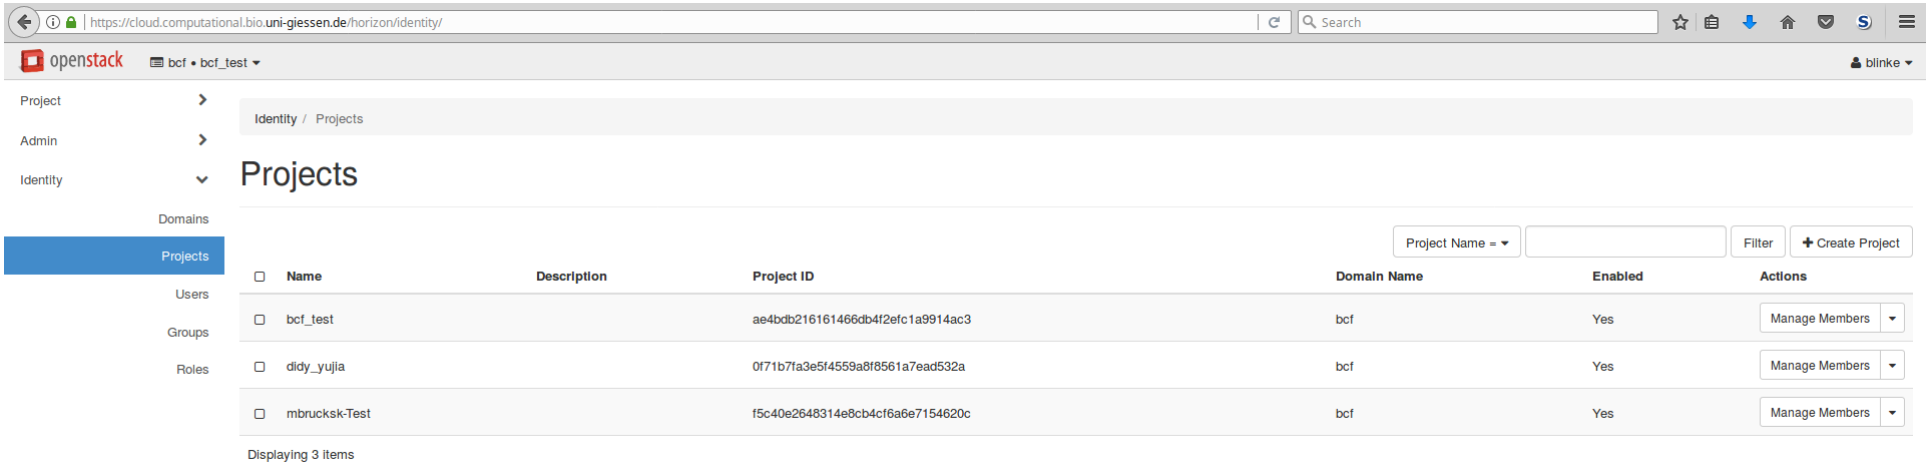

- Web based user interface
- Suitable for most simple task
- Complex task via command line interface or REST API

OpenStack: Horizon

Thanks for your attention!

Questions?