

openstack summit



The logo features the words 'openstack' and 'summit' in a white, sans-serif font. A blue circular arrow connects the 'n' in 'openstack' to the 'S' in 'summit', indicating a cycle or relationship between the two. The word 'users' is written in blue above the arrow, and 'devs' is written in blue below it.

users


devs

A report from
Paolo Veronesi

Overview

- Cos'è OpenStack
- Cos'è il Summit di OpenStack
- Qualche numero del summit e qualche use case
- Cosa c'è di nuovo in Havana e cosa ci sarà in Icehouse
- OpenStack e HA
- Sul deployment di OpenStack
- ... e il VOMS?
- Qualche link su hands on

Visione d'insieme



Fonte: OpenStack

I nomi delle release


- In OpenStack ogni *major release* del software ha un nome in codice.
 - Il ciclo di rilascio delle major release è attualmente di 6 mesi (<http://goo.gl/gMRhb>).
 - La versione attuale, rilasciata a ottobre 2013, è chiamata **Havana**. Questa è la ottava release di OpenStack (<http://goo.gl/MIPbu>).
 - La prossima versione, chiamata *Icehouse*, è prevista uscire all'inizio di aprile 2014.



- **AsiaWorld-Expo, November 5-8, 2013**
openstack.org/summit
- What is the OpenStack Summit?
 - Your chance to hear directly from enterprise users and CTOs about the future of Cloud Computing
 - Where the OpenStack roadmap is determined
 - The largest gathering of the global OpenStack community!

The summit

OpenStack



Over 200 Speakers, including tech industry leaders:



•Mark Shuttleworth, **Ubuntu** Founder



Brian Stevens, CTO of **Red Hat**



•Lew Tucker, Cloud CTO at **Cisco**



Martin Casado, CTO of Networking at **VMware**



•Mark Potts, CTO of **HP Cloud**



Daniel Sabbah, CTO & GM from **IBM**



•Ling Shao, CTO of Research at **IBM**



Haiying Wang, CTO of Cloud at **Huawei**

And many more:


- Intel**
- Rackspace**
- Dell**
- Brocade**
- RightScale**
- Juniper**
- Citrix**

Slide from
Openstack Summit Hong Kong (<http://www.slideshare.net/openstack/openstack-summit-hong-kong>)

OpenStack users and enterprises from all over the world:



Ryan Lane,
Technical Leader
at **Wikimedia**



Carmine Rimi,
Director of Cloud
Engineering,
Workday



Subbu Allamaraju
Chief Engineer,
Cloud at **eBay**



Dan Wilson, Sr
Manager at
Concur



Chris Fisher, VP
Technology
Operations at
Shutterstock

And many more:

- **CERN**
- **PayPal**
- **MercadoLibre**
- **Liveperson**
- **Cyberport**
- **Yahoo!**
- **Yahoo! Japan**
- **Ctrip**
- **Qiyi**
- **Qihoo 360**
- **Alibaba**


Slide from
Openstack Summit Hong Kong
(<http://www.slideshare.net/openstack/openstack-summit-hong-kong>)

And Thousands of Stackers from over 40 countries

- **Technical Attendees** including Developers, Architects, Sys-Admins, DevOps
- **Business Leaders** including Product Managers and Development Managers
- **Executives** including CEOs, CTOs, CIOs
- **And more than 400 different organizations**

Slide from

Openstack Summit Hong Kong (<http://www.slideshare.net/openstack/openstack-summit-hong-kong>)



The summit numbers

several parallel sessions


~3k attendees

65 sponsors

250 talks

*room filling rate:
87%*

50 Countries



Alcuni case studies (1/2)

- **Mercadolibre** - ~eBay per l'America Latina
 - Avevano il problema di gestire database con circa 10^9 immagini, con metadata, etc.
 - Le soluzioni precedenti che avevano non scalavano, il caching layer era inusabile causa re-hashing, nessun supporto per formati multipli, etc.
 - → passati a Swift, che attualmente gestisce 1.2PB in replica 3 (416 TB per copia). Questo set-up attualmente gestisce un upload di 13 milioni di immagini al giorno.
- **Liveperson** – sistema di chat collaborativo usato da imprese per connettersi a clienti
 - 4 data centers gestiti da OpenStack con un totale di 1400 VM.
 - Utilizzano Ubuntu, **Puppet** per la configurazione, **Foreman** per l'installazione, **Mcollective** per l'orchestrazione del cluster e la gestione dei data center (esecuzione di azioni programmatiche su cluster di server, reportistica)
 - Utilizzano il filter scheduler di nova per gestire il provisioning ad-hoc delle VM con OpenStack
- **Shutterstock**: memorizzazione di foto, illustrazioni, immagini vettoriali.
 - Usano MogileFS (<https://code.google.com/p/mogilefs/>) per gestire diversi PB di dati, raccolgono circa 1TB di logfile al giorno.
 - Usano Swift per gestire le repliche, interfacciato con **GeoDNS** (<https://github.com/abh/geodns>) per selezionare la locazione più prossima.

Alcuni case studies (2/2)

- **iQIYI** – è una specie di YouTube cinese, forniscono online television e movie portal.
 - Hanno un numero imprecisato di data centers gestiti da OpenStack
 - Circa 8 milioni di unique visitors al giorno
- **www.360.cn** - una Internet security company
 - Hanno 460 milioni di utenti, con 560 milioni di click al giorno.
 - Hanno 20 data centers gestiti con OpenStack da 2 FTE.
 - Identificano attraverso un anti-virus engine dai 13 ai 15 milioni di files al giorno.
- **Ctrip** – la più grande agenzia di viaggi in Cina
 - 160.000 voli prenotati ogni giorno. 40% di tutti gli ordini di hotel e 28% dei voli sono prenotati via mobile platforms.
 - Hanno da 200.000 a 300.000 chiamate al loro call center al giorno. Per risparmiare sui costi dei desktop, li stanno sostituendo con **thin clients** (basati su ARM con 1 o 2 core, 512MB o 1GB di RAM) e virtual desktop basati su OpenStack, SPICE, SaltStack (simil-Puppet): nel 2013 hanno installato 1.000 virtual desktops, **nel 2014 prevedono di installarne altri 13.000.**
 - I server delle loro farm sono commodity hardware con SSD Multi-level-cell, 4 x 200GB HD in RAID10, 2x8-core CPUs, 160GB RAM.
 - Hanno sviluppato in-house un cloud manager per gestire sia bare metal sia VM in affiancamento a OpenStack (simile a WNoDeS?). **Dipartimento di sviluppo con 15 ingegneri.**
- **La città al mondo con il maggior numero di sviluppatori di OpenStack è Pechino.** E il numero di utenti Internet in Cina è attualmente due volte il numero di utenti Internet in USA.

CERN


- Hanno eseguito 6.288 VMs sulla private Cloud di Rackspace.
- Il 1/10/2013 è partito un progetto congiunto Rackspace/CERN, con 1 FTE permanentemente al CERN per lavorare su:
 - Sviluppo di una reference architecture per la federazione di cloud
 - Creazione di blueprints e contributi al codice di OpenStack
 - Creazione di presentazioni e whitepapers.
- Sostanzialmente vogliono investigare come realizzare federazioni in OpenStack tenendo conto di AuthN, gestione immagini, networking e metering. Vogliono dimostrare che è possibile fare il burst di workloads es. sulla cloud pubblica di Rackspace.
- Le priorità sono:
 - **Identità:** utilizzare le stesse credenziali su Rackspace e al CERN
 - **Aggregazione di servizi:**
 - Definizione di un catalogo (all'interno di Keystone)
 - Gestione delle immagini e loro portabilità (build in one place, carry to other clouds)
 - Aree future:
 - Miglioramenti ai compute services (?)
 - Uso efficiente
 - Regole, policy, business logic per supportare un workload management smart (questa è la parte di accodamento delle richieste Cloud che a noi interessa molto, evidentemente a loro molto meno).
- Esplicitamente, solo le prime due priorità sono in-scope per il progetto di 12 mesi tra CERN e Rackspace.
- Mia impressione: lavoro molto basic per il momento.

Deep Dive into the CERN Cloud Infrastructure

<http://sched.co/16WYvGM>

youtube: <http://www.youtube.com/watch?v=QJll5nBclh4>

Main Topics



Havana Release Notes

<https://wiki.openstack.org/wiki/ReleaseNotes/Havana>

•OpenStack Compute (Nova)

- Added Hypervisor support for containers created and managed using Docker
- Nova has a new feature that allows you to "**shelve**" an instance. This allows instances that are stopped for an extended period of time to be moved off of the hypervisors to free up resources
- Support for the QEMU guest agent (**qemu-guest-agent**) has been added for guests created with the **hw_qemu_guest_agent** property set to **yes**
- Support for **passthrough of PCI devices** from the physical compute node to virtualized guests has been added to Nova. Currently only the libvirt driver provides a working implementation
- Added support for extracting QoS parameters from Cinder and rate limiting disk access based on them when using libvirt-based hypervisors
- The Compute service's **EC2 API has been updated** to use error codes that are more consistent with those of the official EC2 API
- Quotas may now be defined on a **per-user basis**
- Attached Cinder volumes can now be encrypted. Data is decrypted as needed at read and write time while presenting instances with a normal block storage device
- Added native GlusterFS support. If **qemu_allowed_storage_drivers** is set to **gluster** in **nova.conf** then QEMU is configured to access the volume directly using **libgfapi** instead of via **fuse**
- Support for Windows Server / Hyper-V Server 2012 R2
 - VHDX format support
 - Dynamic memory support
 - Ephemeral storage support
 - Compute metrics support for Ceilometer integration

Does the hypervisor matter?

- 90%+ di richieste per KVM, 95%+ di deployment
- 8-10% di richieste per ESXi, 5% di deployment, tipicamente con vCenter.
- Richieste infrequenti per XEN e Hyper-V. Qualcuno chiede lxc per web hosting (ma c'è poco supporto in OpenStack e c'è molta frammentazione sulla tecnologia, es. lxc, Parallels, Docker).
- Esistono comunque casi d'uso per il supporto di hypervisor multipli:
 - Connessione ad una infrastruttura di virtualizzazione esistente
 - Fornitura di una common API indipendente dagli HV
 - “Hedge bets” per proteggersi da bug e da costi dei venditori (es. VMware + KVM)
 - Feature non supportate da qualche HV
- KVM with OVS: 7-8 Gbps con tuning e jumbo frames
- KVM with Mellanox IB: 19.4 Gbps host-to-host; VM-to-VM su host differenti 23.1 Gbps, VM-to-VM sullo stesso host 13.9 Gbps.
- Uso di KVM associato a OVS DPDK (fork di OVS) ottimale anche con pacchetti piccoli, cf. <https://01.org/packet-processing/intel%C2%AE-ovdk>
 - Intel® DPDK is a set of software libraries and Ethernet drivers (native and virtualized) that run in Linux user space to boost packet processing throughput on Intel® architecture.

Slide: <http://www.slideshare.net/mirantis/does-hypervisor-matter-in-openstack>

The Elephant in the Room Hypervisor Breakouts

<http://sched.co/1fzUOuq>

Video:

<http://www.youtube.com/watch?v=y8L6B6Q5EdI>

- vm compromessa o maliziosa: può attaccare le altre vm se accessibili. la vm può attaccare il livello qemu. Da qemu si può attaccare il linux os e da lì tutto il compute e quindi le altre vm.
- Essendo un nodo compromesso, possono essere compromessi tutti gli altri nodi (magari vm presenti su diversi compute con la stessa vulnerabilità). Da lì possono arrivare al block storage (se device montato o se vm gira su cinder).
- 2008 - cloudburst (vmware full breakout)
- 2011 - xen (full break out con tre attacchi diversi)
- 2011 - virtunoid (kvm . full breakout)
- 2012 - sysret.64 (xen - full breakout)
- 2012 - vmdk (data leakage / loss)
- 2013 - kvm tre attacchi: ioapic, set msr. time. Denial of service, potential breakout (tutti CVE-2013-1798 1797 1796)
- bluetooth abilitato in kvm (con sound, multiple nic, e altro) reduce attack surface
- harden compilation
- isolate, detect and alert on exploitation through MAC
- harden your base OS/Dom0
- **maggiori dettagli in** <http://docs.openstack.org/security-guide/content/>

Havana Release Notes

<https://wiki.openstack.org/wiki/ReleaseNotes/Havana>

OpenStack Network Service (Neutron)

- **Load Balancer (LBaaS)** - This release ships an updated API and with HAProxy driver support.
- **VPN (VPNaaS)** - Site-to-Site IPSec VPNs are now supported via the VPN service plugin.
- **Firewall (FWaaS)** - The firewall service enable tenant to configure security in depth as rules can be applied both at the edge via the firewall API and on the VIF via the security group API. Experimental.
- **Modular Layer 2 (ML2)** (see later)

Modular layer2 in OpenStack neutron

The Modular Layer 2 (ML2) Plugin is a framework allowing OpenStack Neutron to simultaneously utilize the variety of layer 2 networking technologies found in complex real-world data centers.



- Slide:
 - <http://www.slideshare.net/mestery/modular-layer-2-in-openstack-neutron>
- A new Neutron core plugin in Havana
- Modular
 - Drivers for layer 2 network types and mechanisms - interface with agents, hardware, controllers, ...
 - Service plugins and their drivers for layer 3+
- Works with existing L2 agents
 - openvswitch
 - linuxbridge
 - hyperv
- Deprecates existing monolithic plugins
 - openvswitch
 - Linuxbridge
- Supports multi-segment networks


Neutron Server									
ML2 Plugin					API Extensions				
Type Manager					Mechanism Manager				
GRE TypeDriver	VLAN TypeDriver	VXLAN TypeDriver	Arista	Cisco Nexus	Hyper-V	Population L2	Linuxbridge vSwitch	Open vSwitch	Tail-F NCS

SDN

- Opendaylight (<http://www.opendaylight.org/>)
 - Transition from configuration time to software time
 - Progetto iniziato in Aprile 2013. Prima release programmata per il 9/12/2013.
 - Package composto da 14 sotto-progetti.
 - Ci sarà una “base edition” con supporto a un single OpenFlow controller (OpenFlow 1.0 e 1.3).
 - Una “service provider edition” che aggiunge controllo di DDoS, supporto a BGP.
 - E una “virtualization edition” (probabilmente la più interessante) che comprende anche un plugin per Neutron e in generale supporto per OpenStack.

OpenDaylight: an open source SDN for your OpenStack cloud

<http://www.slideshare.net/aashaikh/open-daylight-openstackicehousesummit>



Havana Release Notes

<https://wiki.openstack.org/wiki/ReleaseNotes/Havana>

▪OpenStack Block Storage (Cinder)

- **Volume migration** - admin API to migrate a volume to a different Cinder back-end
- Scheduler hints extension added to V2 API
- Added ability to **extend the size** of an existing volume
- Added ability to transfer volume from one tenant to another
- Added **Ceph driver to backup** service (allowing Ceph as a backup target with differential backups from Ceph to Ceph)
- Dell EqualLogic volume driver
- VMware VMDK cinder driver
- IBM General Parallel File System (**GPFS**)
- Allow Ceph as an option for volume backup
- IBM Tivoli Storage Manager (**TSM**)

▪OpenStack Object Storage (Swift)

▪Global clusters support


- The "region" concept introduced in Swift 1.8.0 has been augmented with support for using a **separate replication network** and configuring read and write affinity.

▪Disk performance

- The object server now can be configured to use threadpools to increase performance and smooth out latency throughout the system. Also, many disk operations were reordered to increase reliability and improve performance.

▪Added support for pooling memcache connections

▪Much faster calculation for choosing handoff nodes



- **Openstack Summit HK - Ceph defacto** = eNovance <http://www.slideshare.net/enovance/ceph-defactoenolza>
- [On the Gluster vs Ceph Benchmarks](#)
- Commodity Storage Back-end Technologies: http://docs.openstack.org/trunk/openstack-ops/content/commodity_storage_backends.html

Havana Release Notes

<https://wiki.openstack.org/wiki/ReleaseNotes/Havana>

•OpenStack Identity (Keystone)

- Authorization data (tenants/projects, roles, role assignments; e.g. SQL) can now be stored in a separate backend, as determined by the "**assignments**" driver, from authentication data (users, groups; e.g. LDAP), as determined by the "**identity**" driver
- Credentials (e.g. ec2 tokens) can now be stored in a separate backend, as determined by the "**credentials**" driver, from authentication data
- Ability to specify more granular RBAC policy rules (for example, based on attributes in the API request / response body)
- Pluggable handling of external authentication using REMOTE_USER
- Token generation**, which is currently either UUID or PKI based, is now pluggable and separated from token persistence.
- Delegated role-based authorization to arbitrary consumers using OAuth 1.0a
- External authentication providers** can now embed a binding reference into tokens such that remote services may optionally validate the identity of the user presenting the token against an presented external authentication mechanism. Currently, only kerberos is supported.

Havana Release Notes

<https://wiki.openstack.org/wiki/ReleaseNotes/Havana>

•OpenStack Image Service (Glance)

•Storage Support (cinder, sheepdog, GridFS)

•**Multiple Image Locations:** Image service images can now be stored in multiple locations. This enables the efficient consumption of image data and the use of backup images for the event of a primary image failure.

•Storage Quota

•Membership Policy

•OpenStack Dashboard (Horizon)

•Heat, Ceilometer (initial), Keystone (Domains, Groups, and More), Trove Databases (initial)

- Editable default quotas, The ability for an administrator to reset the password of a server/instance.
- Availability zone support, Improved region support, Instance resizing, Improved boot-from-volume support, Per-project flavor support
- VPN as a Service, Firewall as a Service, Editable and interactive network topology visualizations, Full security group and quota parity between Neutron and Nova network.
- **Self-Service Password Change**

Havana Release Notes

<https://wiki.openstack.org/wiki/ReleaseNotes/Havana>

▪ **OpenStack Orchestration (Heat)**

- **Ceilometer integration for metrics/monitoring/alarms**
- **Concurrent resource operations**
- **Initial support for native template language (HOT)**

▪ **OpenStack Metering (Ceilometer)**

- **The statistics endpoint can now be used to group samples by some fields using the groupby argument**
- **A new alarm API is now available:** Alarms are a new feature allowing users and operators to trigger actions based on comparing the statistics trend against a threshold over a period of time.
- **Users can now post their own samples and meters through the API**

Looking forward to the next release: Icehouse

- Many projects were Incubated in Havana creating a hotbed of innovation. Next release, Icehouse, is coming in April 2014:

•Deployment (TripleO)

- See <https://wiki.openstack.org/wiki/TripleO>

•Bare metal (Ironic)

- See <https://wiki.openstack.org/wiki/Ironic>

•Hadoop (Savannah)

- See <https://wiki.openstack.org/wiki/Savanna>

•Relational Databases (Trove)

- See <https://wiki.openstack.org/wiki/Trove>

•Messaging (Marconi)

- See <https://speakerdeck.com/cabrera/introducing-openstack-marconi>

Openstack HA

Tanti componenti in gioco, per ognuno di questi ci sono tanti modi per configurarli in HA

Vengono usate sostanzialmente due modalita':

- Active/passive
- Active/active (ove possibile).

Di base, puo' essere utile un layer di HA comune:



- Nella documentazione di OpenStack si usa pacemaker + corosync

In setup piu' complessi, si usano CELLE, REGIONI, AVAILABILITY ZONE, HOST AGGREGATION (see <http://docs.openstack.org/trunk/openstack-ops/content/scaling.html>)


Alcuni talk dal summit

- Practical Lessons from Building a Highly Available Openstack Private Cloud
<http://www.slideshare.net/Pixelpark/2013-nov-pixelparkpresentationopenstacksummit2013hongkong>
- High Availability OpenStack at PayPal: <http://www.slideshare.net/scottjcarlson/high-availability-openstack-at-paypal-openstack-summit-fall-2013>
- Open stack@ebay: <http://www.slideshare.net/sallamar/open-stackebay>
- Database Cluster as a Service in OpenStack: <http://www.slideshare.net/skysql/webinar-131114d-baas-openstack-for-slidehsare>
- Rally--OpenStack Benchmarking at Scale <http://www.slideshare.net/mirantis/rally-benchmarkingatscale>

Cell VS Region



OpenStack Deployment Frameworks





See: **An Evaluation of OpenStack Deployment Frameworks**

http://www.slideshare.net/shane_gibson/open-stack-hong-kong-summit-2013

Provisioning Evaluation: Network Architecture

Provisioning Evaluation: OpenStack Overview

Provisioning Evaluation: Frameworks Tested

Provisioning Evaluation: The Results

Tool	TTC *	Capabilities	Resiliency	Complexity	Version
Crowbar	★	★	●	★	ver 1.6
Fuel Web	▼	◆	●	◆	ver 3.0.1
Rackspace P.C.	▼	N/A	N/A	★	ver 4.1.0
MaaS/Juju	▼	◆	▼	▼	ver 1.2/0.7
Foreman	◆	★	◆	▼	ver 1.2.0

Basic metal prov.
HW checks
notification/logging

deploy tool HA
multi DC
multi zone


OpenStack HA
hardware/BIOS/RAID
physical network
OpenStack tuning

- ★ good support
- ◆ meets requirements
- ▼ minimum requirements
- missing features

* TTC = time-to-cluster; the time it took deployment to be customized, documented, and repeated by Symantec staff

Scheduling

*HAVANA scheduling:
best effort*



*ICEHOUSE
scheduling:
affinity
anti-affinity
network proximity
host capabilities
resource exclusivity*

External scheduler? Discussion at

<https://etherpad.openstack.org/p/icehouse-external-scheduler>

Da incontro di ieri sera con Enter (www.enter.it): l'idea è unificare i vari scheduler presenti in diversi componenti di OpenStack in un servizio di scheduling unico per tutti i componenti.

Operazione simile fatta per vari componenti (le cosiddette librerie OSLO) in modo tali che i diversi componenti usino librerie comuni.

Non si parla di pluggare scheduler esterni.

E per concludere...

The Virtual Organization Concept for Authorization Management in Federated Clouds

- <http://sched.co/181GiHZ>
- Slide: <https://pandora.infn.it/data/public/24a98f.php>
- use cases:
 - cloud stakeholders for disaster response
 - <http://www.earthobservations.org/index.shtml>
 - cloud-based feature film production for budget film
- Tutti hanno un loro sistema di autenticazione, non hanno bisogno e non vogliono altre credenziali: Federated identity management, trust third parties:
- Ho seguito la presentazione e ci sono due approcci: uno con un voms esterno e uno con le funzionalita' di VOMS integrato in openstack.
- L'impressione, viste anche le numerose imprecisioni nella presentazione, è che non conoscano bene voms (dicono che è uno spof, si preoccupano di chi lo gestisce, etc, tutte cose risolte in grid da tempo direi).
- Non contatti con sviluppatori VOMS

Documentazione sulle presentazioni

- **AGENDA:** <http://www.openstack.org/summit/openstack-summit-hong-kong-2013/agenda/>
- Videos from the Keynote and breakout sessions are being uploaded now:
 - <http://www.openstack.org/summit/openstack-summit-hong-kong-2013/session-videos/>.
- **SLIDESHARE** <http://www.slideshare.net/>

Next Summit:

- April 2014 – Portland
- **Nov 2014 - Paris**

QaaS (QuestionAndMaybeAnswer as a Service)

BACKUP

Hands on

Automate Windows images on OpenStack

Video:

<http://www.openstack.org/summit/openstack-summit-hong-kong-2013/session-videos/presentation/automate-windows-images-on-openstack>

- guest initialization con cloudbase-init (apache 2 license) invece di cloud-init (il merge è in todo)
- sysprep opzionale (ma obbligatorio per avere supporto da ms)
- può essere automatizzato con OOB
- disponibile ws2012 r2 valida per 180 giorni per valutazione (non per produzione).
 - Possibilità di eseguire uno script per aggiungere la licenza e a quel punto si può usare in produzione
- su github script unattended per generare immagine windows per openstack (fa anche windows update)
- si parla di ssh equivalente WSMAN/WinRM. anche questo documentato in github
- disponibile configurazione per hyper-v per nova.
- supporto di heat per windows

Hands-on with OpenStack + vSphere

- <http://sched.co/181F7br>
- vova appliance per vmware con openstack che si può scaricare. lab aperti su www.vmware.com/go/openstacklab in particolare è stato fatto Lab: HOL-SDC-1320 - OpenStack on VMware vSphere

OpenStack Networking Hands-on Lab

- <http://sched.co/19y7tZW>
- Neutron Hands on Lab!
Slides: <http://bit.ly/HBjTsg>
- ottimo lab, slide in ppt, un tenant per utente, documentazione ottima