

 Stay ahead.

Innovate2013

The IBM Technical Summit

開發者大會





OpenStack: Linux in Cloud Management Software

Dr. Tzi-cker Chiueh

Innovate2013

The IBM Technical Summit

開發者大會





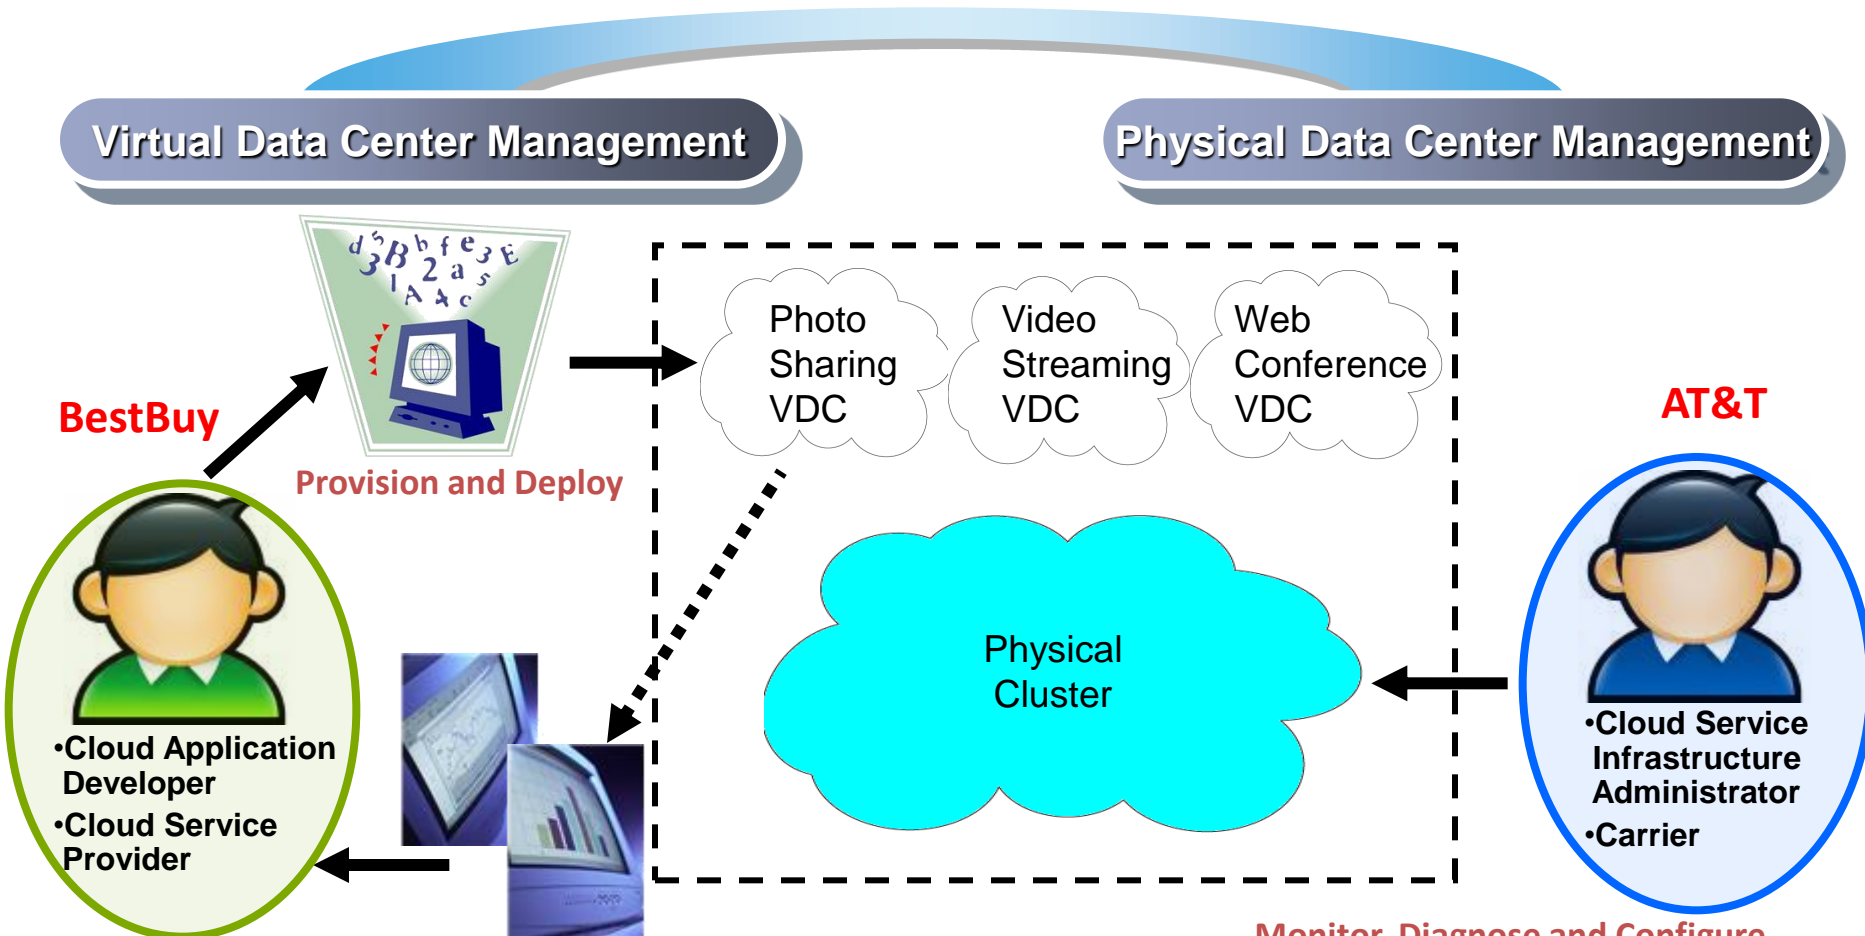
OpenStack Introduction

- An open ecosystem for developing a complete open-source cloud software stack for IaaS
- Target: AWS-EC2 like IaaS
- Consists of (a) an open-source OpenStack core, and (b) a set of APIs that interfaces the core with third-party plugins
- Commercial significance:
 - Provides a platform that enables open competition for third-party component solutions
 - Gives cloud service providers the ability to avoid lock-ins





Data Center Virtualization



Monitor and Configure Virtual Resources

Monitor, Diagnose and Configure Physical Resources





IaaS Service Model

- Multiple virtual data centers or tenants on a single physical data center
- Virtual Data Center specification
 - No. of VM instances each with CPU performance and memory size requirement
 - Per-VM storage space requirement
 - External network bandwidth requirement
 - Security policy
 - Backup policy
 - Traffic shaping policy
 - Load balancing and auto-scaling policy
 - Network configuration: public IP address and private IP address range
 - OS image and application image





OpenStack Technologies

- **Server** virtualization: Nova
- **Storage** virtualization: Cinder
- **Network** virtualization: Quantum
- **Management** virtualization: Horizon and Ceilometer
- **Security**: Keystone
- **Support**: Glance, Swift and Heat





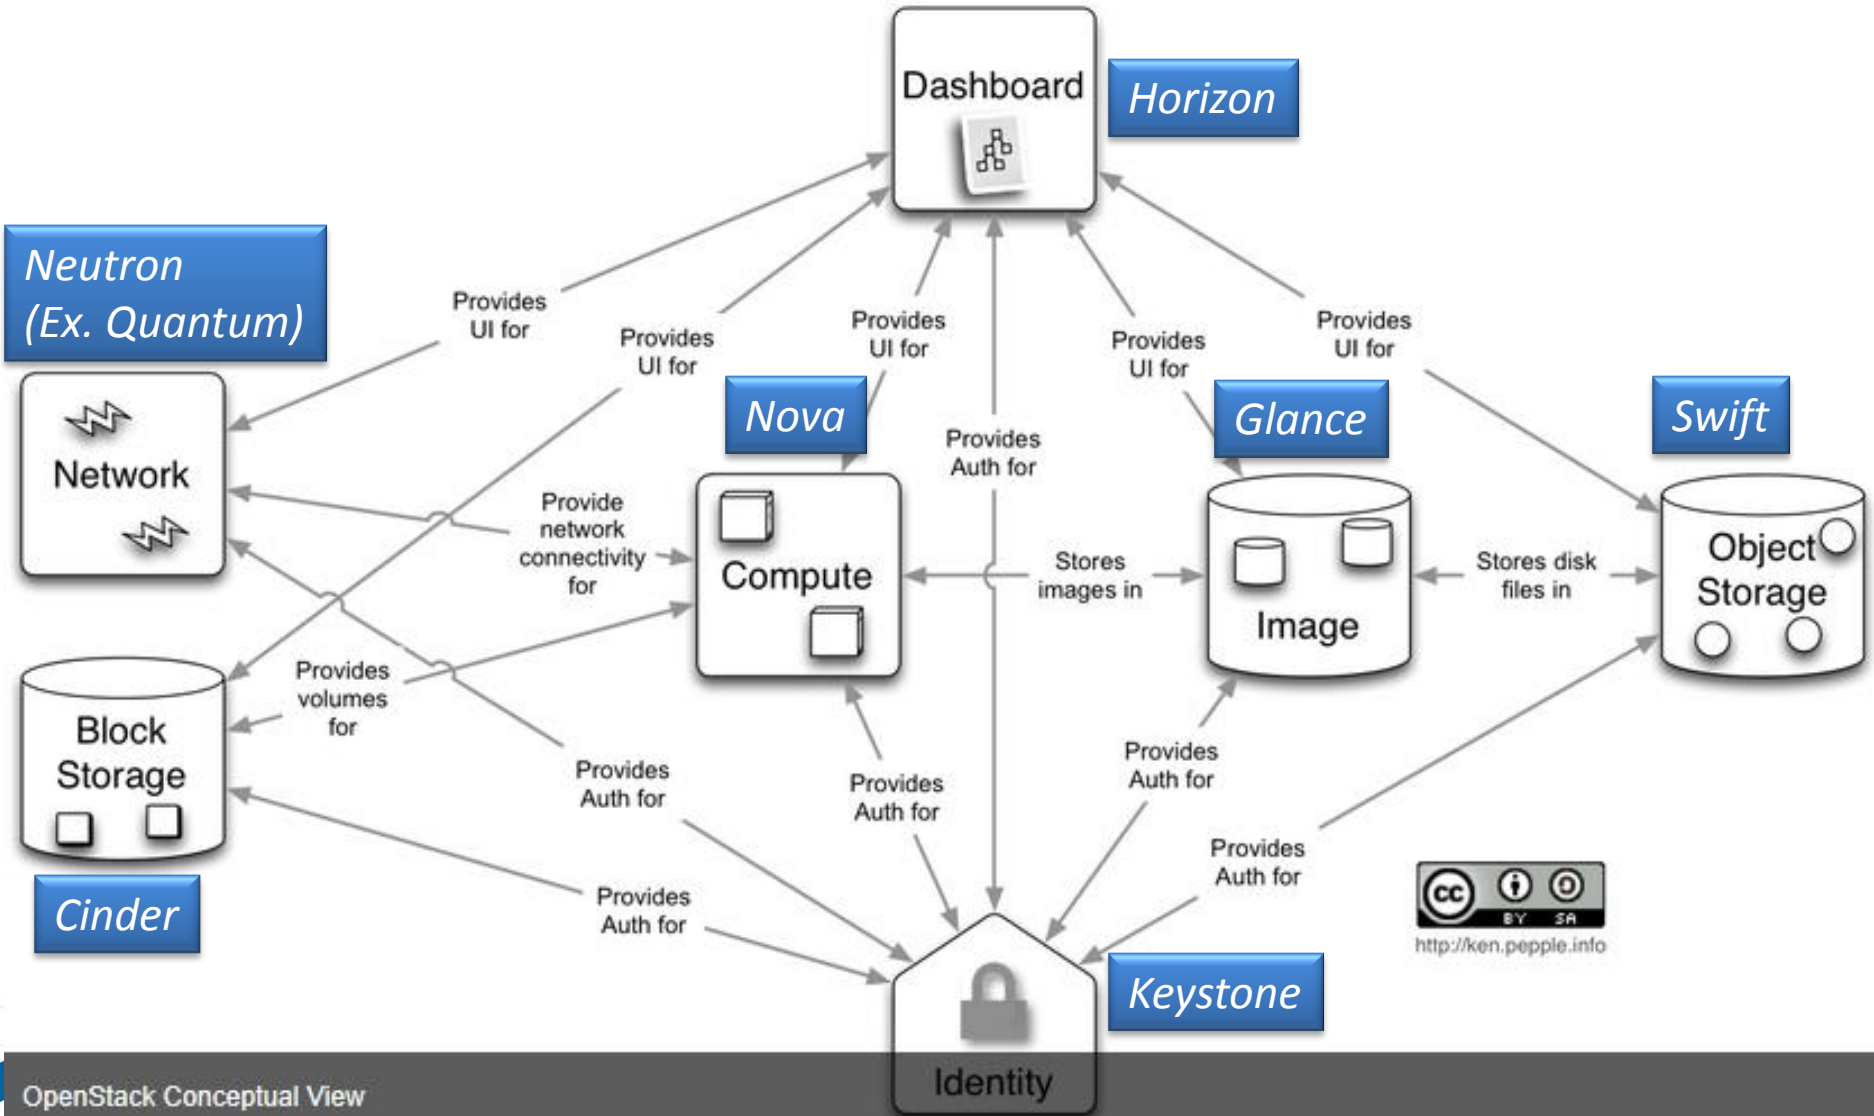
Network Virtualization

- Multiple virtual networks run on a single physical network
- Each VDC has its own virtual network
 - VMs in a VDC could be organized into one or multiple IP subnets
 - A full private IP address space (i.e. 10.x.x.x)
 - A set of public IP addresses as service entry points
 - L3-VPN connects VDCs that share the same IP address space
 - Per-VC firewall, server load balancing and traffic shaping policy
 - Inter-VDC connections through public IP addresses but local routers
- Support for private IP address reuse: a private IP address such as 10.1.2.5 could be used in multiple VDCs simultaneously





OpenStack Conceptual View



OpenStack Is Fast Evolving



	Austin	Bexar	Cactus	Diablo	Essex	Folsom	Greezly	Havana
	Nov-10	Feb-11	Apr-11	Oct-11	Apr-12	Oct-12	Apr-13	Oct-13
OpenStack Compute (Nova)								
OpenStack Object Storage (Swift)								
OpenStack Image Service (Glance)								
OpenStack Identity (Keystone)								
OpenStack Dashboard (Horizon)								
OpenStack Networking								
OpenStack Block Storage (Cinder)								
Metering (Ceilometer)								
Basic Cloud Orchestration & Service Definition (Heat)								
Python libraries (Oslo)								





OpenStack Foundation – Platinum

Platinum Members



AT&T



Canonical



HP



IBM



Nebula



Rackspace



Red Hat, Inc.



SUSE





OpenStack Foundation – Gold

Gold Members



CCAT



Cisco



Cloudscaling



Dell



DreamHost



eNovance



Ericsson



Intel



Juniper Networks



Mirantis



Morphlabs



NEC



NetApp



Piston Cloud



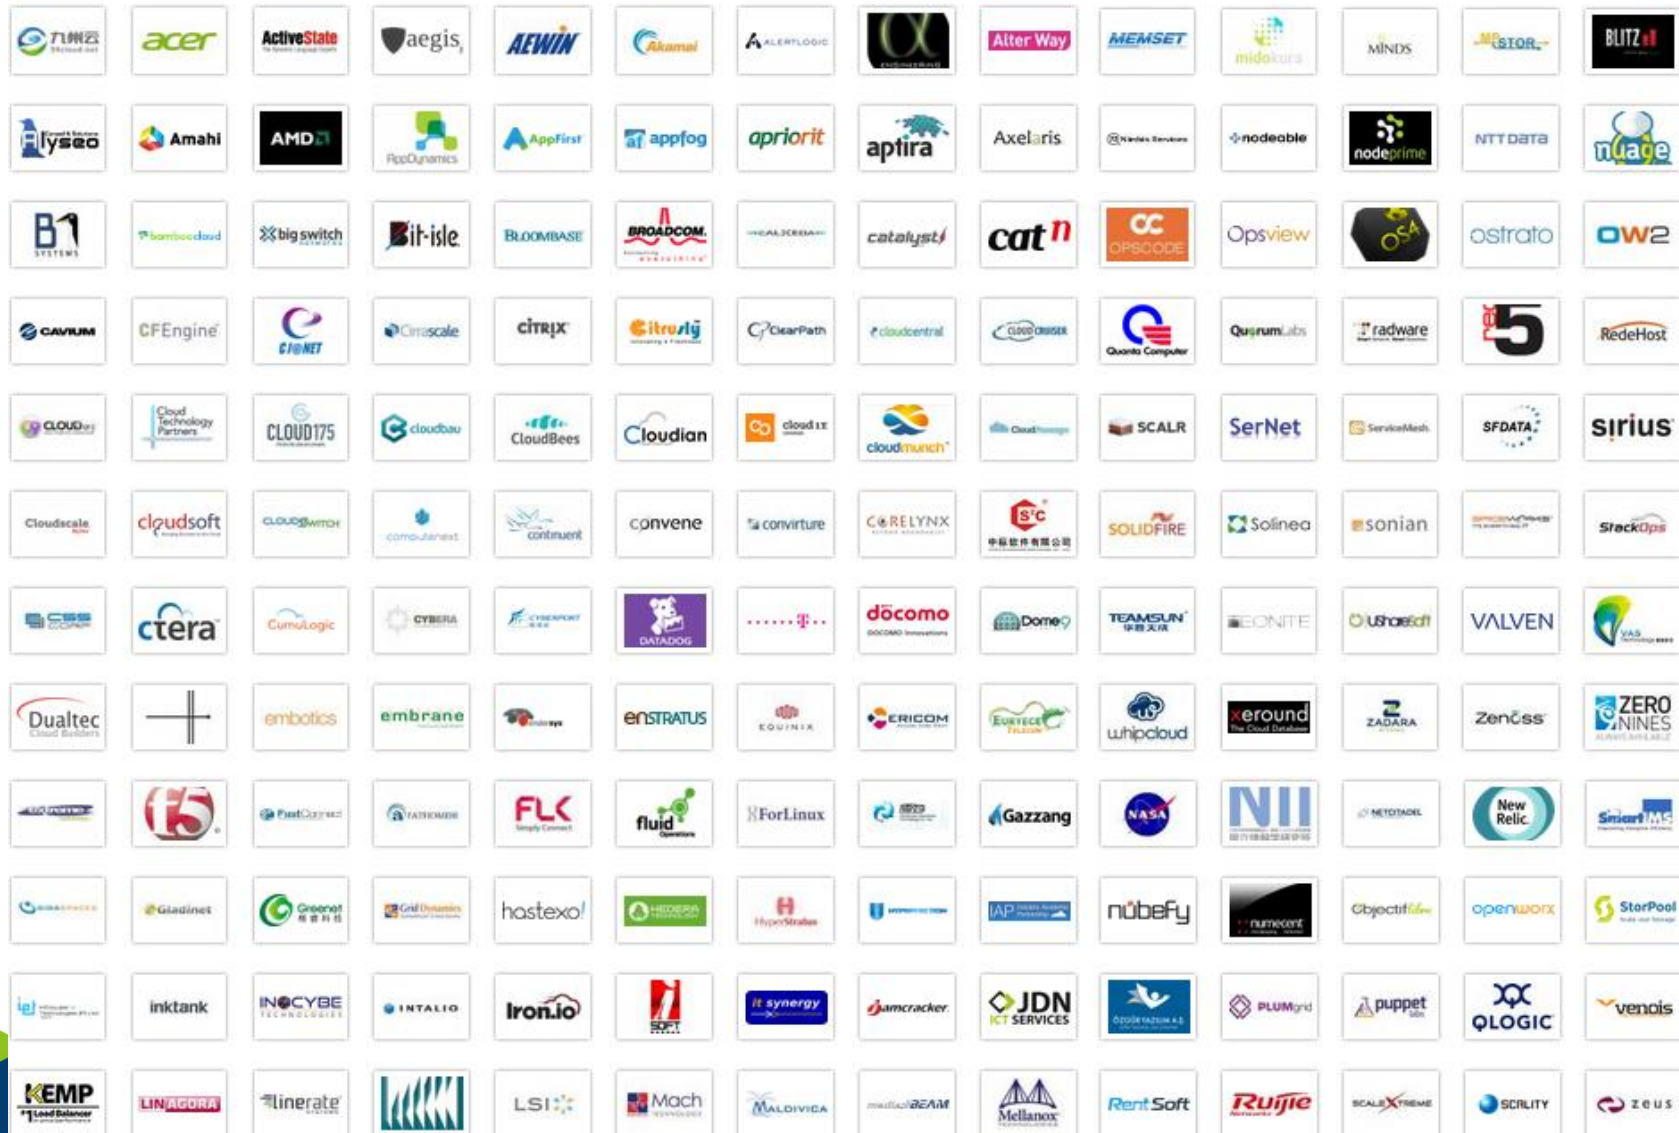
VMware



Yahoo!



OpenStack Foundation – more...



The IBM Technical Summit





OpenStack Ecosystem

- **OpenStack core + API:** OpenStack Foundation
- **OpenStack plugin provider**
 - Cinder: ITRI, Zadara
 - Neutron: ITRI/Pica8, Midokura, bigSwitch
 - Object storage: Swiftstack
- **OpenStack solution provider**
 - ITRI, IBM, HP, Ubuntu, Redhat, Rackspace
- **OpenStack system integrator**
 - Cloudscaling, MorphLabs, Dell, HP, IBM, Piston computing
- **OpenStack user:** Rackspace, KT, AT&T, HP, Metacloud





Current Status

- > 10,000 members over 125 countries
- 860+ code contributors
- 50 summit sponsors
- 2500+ summit attendees
- OpenStack Summit at Hong Kong: Nov. 5-8, 2013
- Taiwan OpenStack user group (TWOSUG)
 - Founded in December 2011 in ITRI
 - 7 meetups, 1 workshop and +500 participants





To-do Items for OpenStack

- More commercially ready network virtualization (Neutron) and storage virtualization (Cinder) plugin
- Management virtualization
- Disaster recovery support
- High availability support
- Scalability support
 - Dynamic optimizations for server virtualization





ITRI Cloud OS

- **Deployment:**
 - Bare metal asset discovery, and installation of hypervisor and Cloud OS
- **Server virtualization:** built on **Nova**
 - **Static provisioning:** virtual machine, virtual disk, virtual network and security
 - **Dynamic resource management:** power consolidation, inter-PM load balancing
- **Storage virtualization:** based on **Cinder API**
 - Comprehensive data protection with low performance overhead
 - Scalable metadata management
- **Network virtualization:** based on **Quantum API**
 - Private IP address space reuse with multiple IP subnets
 - VDC-specific auto-scaling, server load balancing, traffic shaping, VPN and NAT
- **Management virtualization:**
 - Separation of **physical** from **virtual** data center administration
 - Application performance reporting
 - Real-time CPU/memory/disk/network resource usage metering
- **Security:**
 - **Centralized** L3/L4 Firewall
 - **Distributed** L4/L7 and Web Application Firewall
 - **Inter-VDC isolation:** VMs in one VDC cannot ping VMs in another VDC



Comprehensive Data Protection

- Storage hardware: **JBOJBOD**, or just a bunch of JBOD servers
 - Centralized **metadata** node with HA support and distributed **data** nodes
- **SW/HW RAID**: disk failure
- **Inter-data-node replication**: disk, controller, data node, and network switch failures
- **Local data backup**: manual error
 - Zero data copying
 - Scalable Disk de-duplication
 - Scalable garbage collection
- **Wide-area data backup**: site failure
- **HA support** for metadata node: metadata node failure



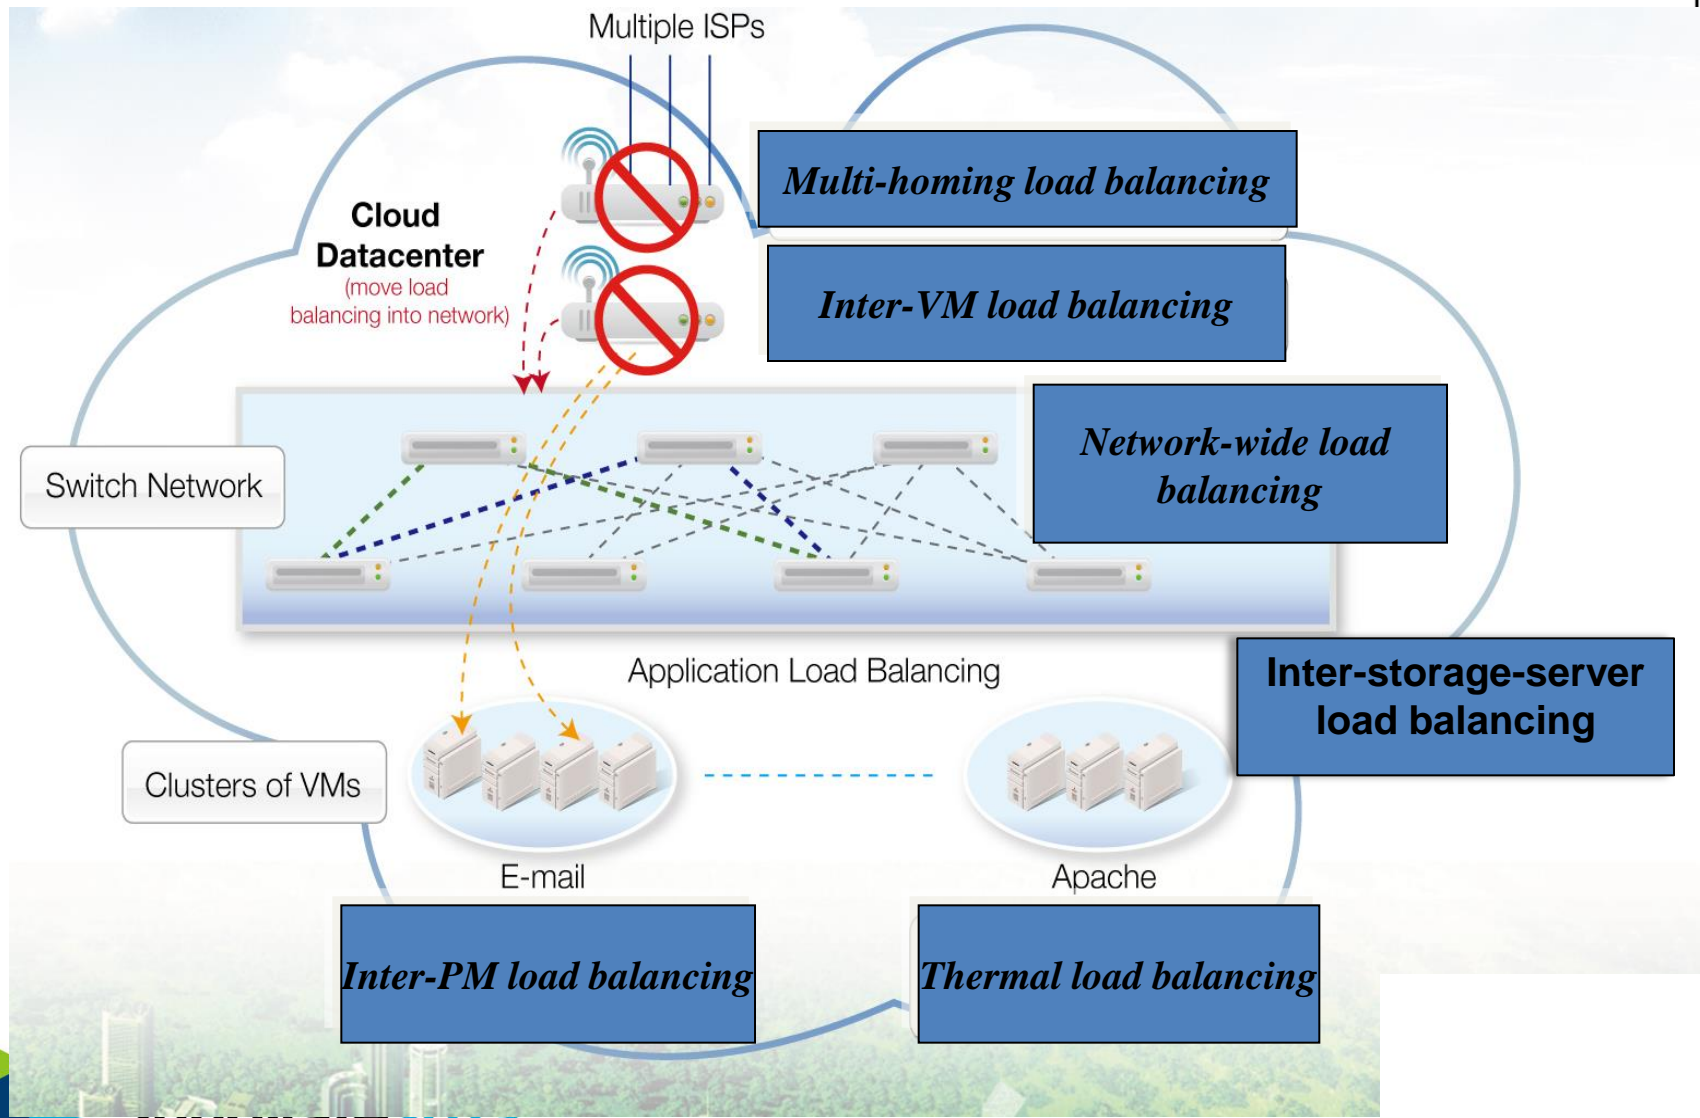


High Availability Support

- Cloud OS components
 - Active-passive
 - Logging + Log replication + DRBD + Heartbeat
 - Active-active:
 - Clustering
- User VMs
 - Disk state-preserving fail-over for application VMs running inside VDCs
 - Shared persistent state + VM restart + take-over
 - Memory state-preserving fail-over for application VMs running inside VDCs
 - Shared persistent state + Continuous VM state migration + take-over



Multi-Dimensional Load Balancing





System Management/Administration

- Separation of **physical** from **virtual** data center administration
 - Multi-tenancy aware
- Comprehensive monitoring
 - Health and usage data for servers, switches, and disks
 - Health of cloud management software components
- **Multi-tenancy-aware virtual to physical resource mapping**
 - Virtual Machines → Physical Machines
 - Virtual Volumes → Physical Disks
 - Virtual Network Links → Physical Network Links
- **Unified** log collection, filtering and access
- Integration with **trouble ticketing**





Summary

- OpenStack is coming! OpenStack is coming!
- OpenStack is not just another open-source cloud management software distribution, but an open ecosystem that allows everybody to play
- Existing open-source OpenStack distribution is not complete and may take another 12 to 18 months to mature
- ITRI Cloud OS is three years in the making and is going to be compatible with the Grizzly distribution by the end of 2013
- The current version of ITRI Cloud OS is already more functional than the current OpenStack roadmap





Thank You!

Questions and Comments?

tcc@itri.org.tw

Please check out ITRI Cloud OS Demo in OpenStack Booth



Innovate2013
The IBM Technical Summit