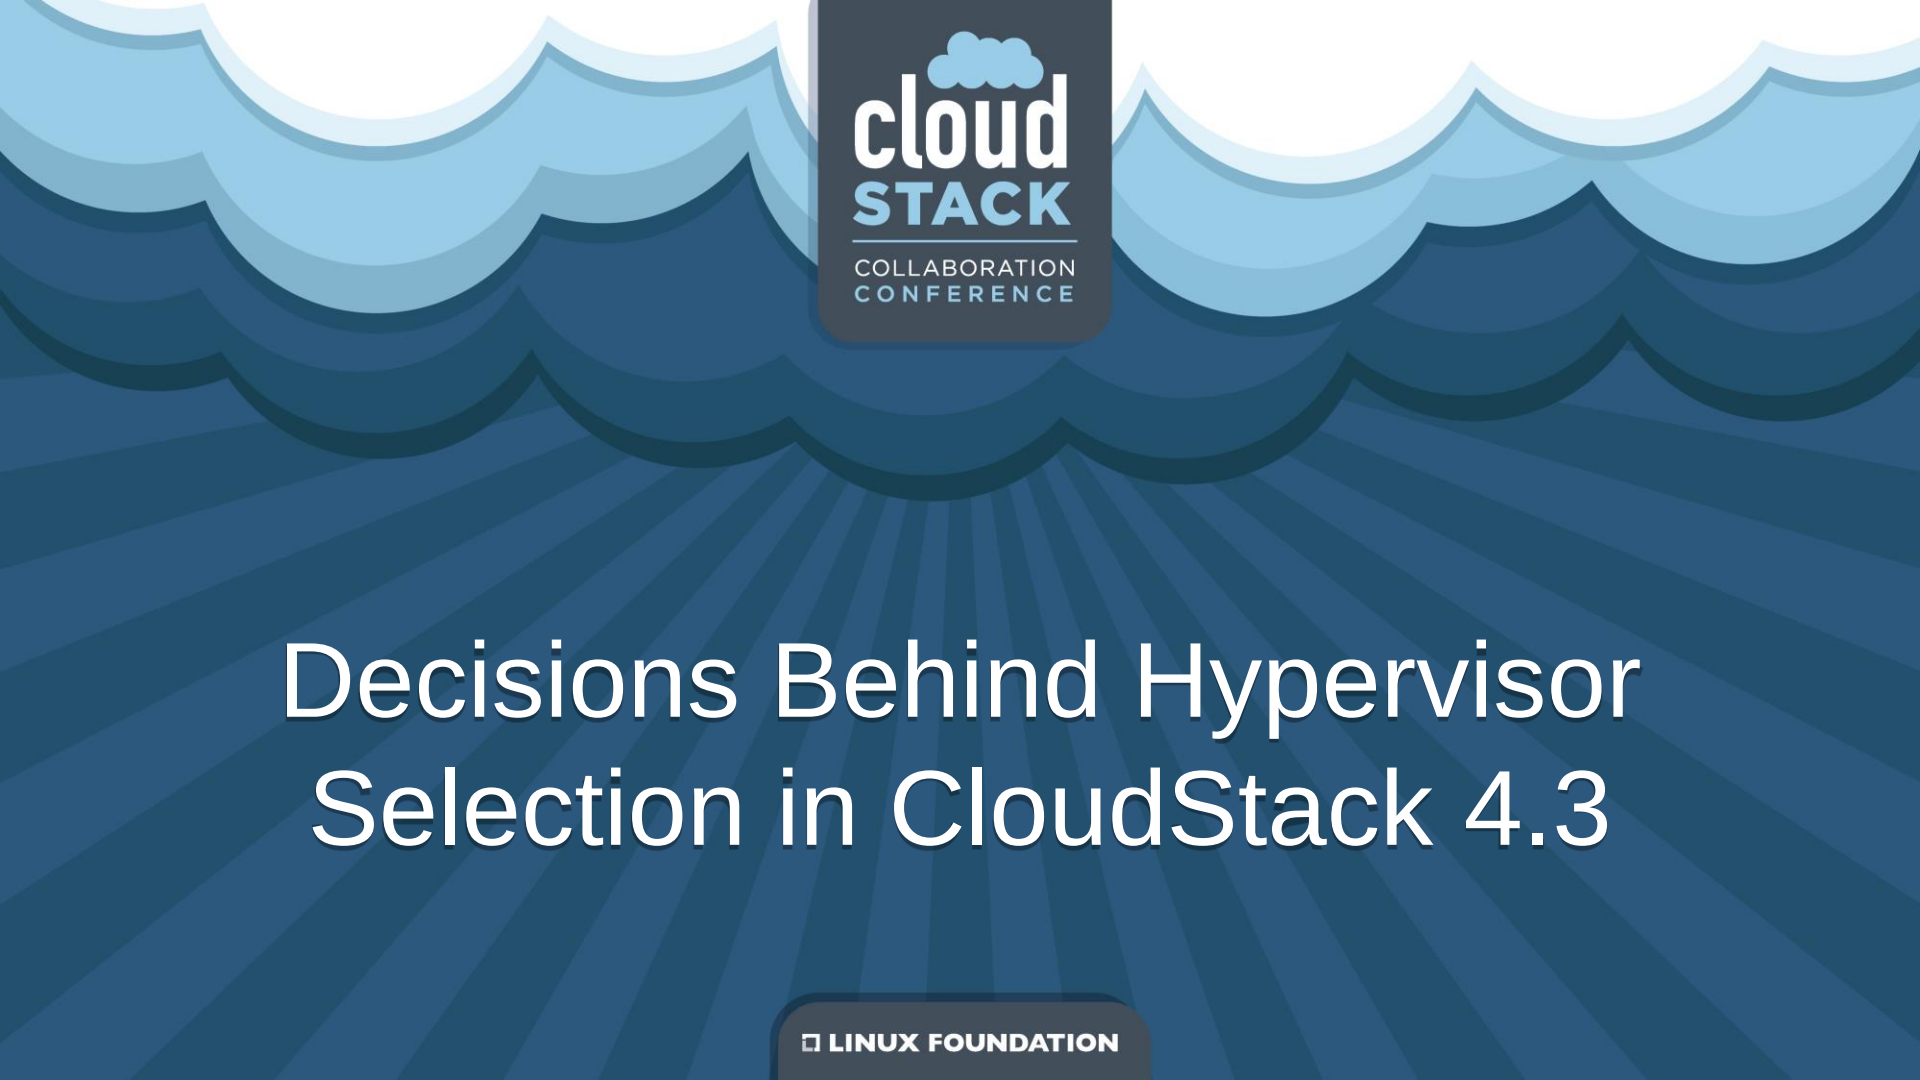
cloud
STACK
COLLABORATION
CONFERENCE
Decisions Behind Hypervisor Selection in CloudStack 4.3
LINUX FOUNDATION

whoami

- Name: Tim Mackey
- Current roles: XenServer Community Manager and Evangelist; occasional coder
- Cool things I've done
  - Designed laser communication systems
  - Early designer of retail self-checkout machines
  - Embedded special relativity algorithms into industrial control system
- Find me
  - Twitter: @XenServerArmy
  - SlideShare: slideshare.net/TimMackey

What are we trying to accomplish?

Service Offerings

- Clearly define what you want to offer
  - What types of applications
  - Who has access, and who owns them
  - What type of access
- Define how templates need to be managed
  - Operating system support
  - Patching requirements
- Define expectations around compliance and availability
  - Who owns backup and monitoring

Define Tenancy Requirements

- Department data local to department
  - Where is the application data stored
- Data and service isolation
  - VM migration and host HA
  - Network services
- Encryption of PII/PCI
  - Where do keys live when data location unknown
  - Need encryption designed for the cloud
- Showback to stakeholders
  - More than just usage, compliance and audits

Virtualization Infrastructure

- Hypervisor defined by service offerings
  - Don’t select hypervisor based on “standards”
  - Understand true costs of virtualization
  - Multiple hypervisors are “OK”
  - Bare metal can be a hypervisor
- To “Pool” resources or not
  - Is there a real requirement for pooled resources
  - Can the cloud management solution do better?
- Primary storage defined by hypervisor
- Template storage defined by solution
  - Typically low cost options like NFS

Defining the network …

Flat Network – Basic Layer 3 Network

Option	XenServer	vSphere	KVM	LXC	Hyper-V
Security Groups	Yes- bridge	No	Yes	Yes	Yes
IPv6	No	No	Yes	Yes	No
Multiple IPs per NIC	Yes	Yes	Yes	Yes	Yes
Nicira NVP	Yes	No	Yes	No	No
BigSwitch VNS	Yes	No	Yes	No	No

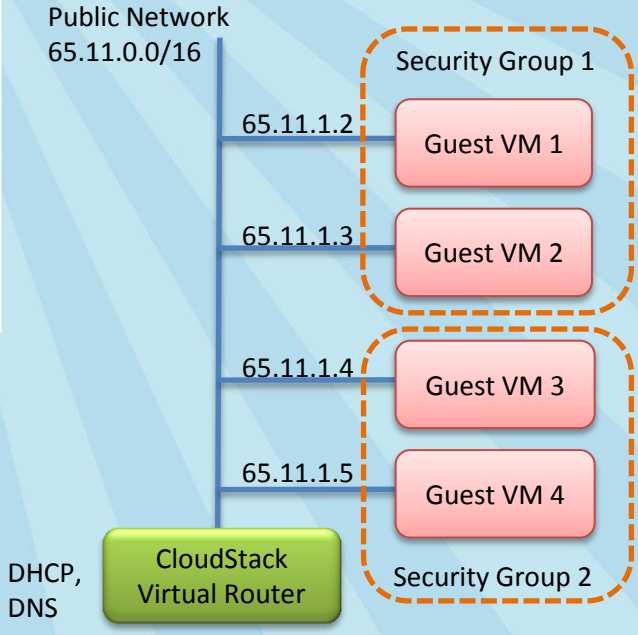

VLANs for Private Cloud

Option	XenServer	vSphere	KVM	LXC	Hyper-V
Max VLANs	800	254	1024	1024	4094
IPv6	No	No	Yes	Yes	No
Multiple IPs per NIC	Yes	Yes	Yes	Yes	Yes
Nicira NVP	Yes	No	Yes	No	No
BigSwitch VNS	Yes	No	Yes	No	No
MidoKura	No	No	Yes	No	No
VPC	Yes	Yes	Yes	No	Yes
NetScaler	Yes	Yes	Yes	No	Yes
F5 BigIP	Yes	Yes	Yes	No	Yes
Juniper SRX	No	Yes	Yes	No	Yes
Cisco VNMC	No	Yes	No	No	No

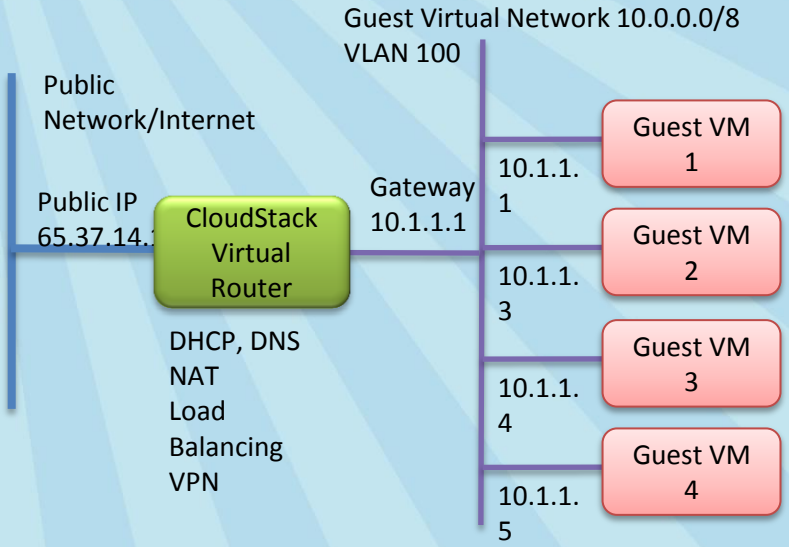

Beyond the VLAN – Network Virtualization

Option	XenServer	vSphere	KVM	LXC	Hyper-V
OVS GRE tunnels	Yes	No	No	No	No
Nicira STT tunnel	Yes	Yes	Yes	No	No
MidoNet	No	No	Yes	No	No
VXLAN	No	Yes	Yes	No	No
NVGRE	No	No	No	No	No
Nexus 1000v	No	Yes	No	No	No
Juniper Contrail	Yes	No	No	No	No
Palo Alto	Yes	Yes	Yes	No	No

OPEN NETWORKING
FOUNDATION

Virtual Private Cloud and nTier Applications

Feature	XenServer	vSphere	KVM	LXC	Hyper-V
PVLAN	Yes - ovs	Yes	ovs	No	No

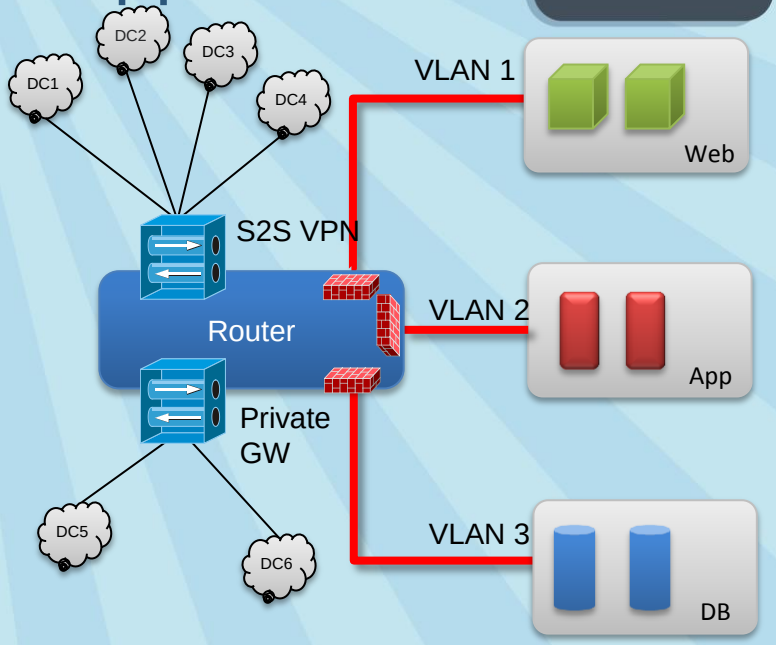

Delivering specific network services

- KVM
  - IPv6
  - Security groups
  - Large quantity of VLANs
- vSphere
  - VXLAN required vSphere Enterprise Plus
  - Cisco Nexus 1000v and ASA 1000v require vSphere Enterprise Plus
- XenServer
  - Security groups
  - Large quantity of VLANs
  - Juniper Contrail

Instances need a home …

Primary Storage Options

Feature	XenServer	vSphere	KVM	LXC	Hyper-V
Local storage	Yes	Yes	Yes	Yes	Yes
NFS	Yes	Yes	Yes	Yes	No
SMB	No	No	No	No	SMB3
Single path iSCSI	Yes	Yes	Yes	No	No
Multipath iSCSI	PreSetup	No	No	No	No
Direct array	No	VAAI	No	No	No
Shared Mount	No	No	Yes	Yes	No
Template format	VHD	OVA	QCOW2	TAR	VHD
SolidFire Plugin	Yes	Yes	Yes	No	No
NetApp Plugin	Yes	Yes	Yes	No	No
Zone wide	No	Yes	Yes	No	No
Ceph RBD	No	No	Yes	No	No
Clustered LVM	No	No	Yes	No	No

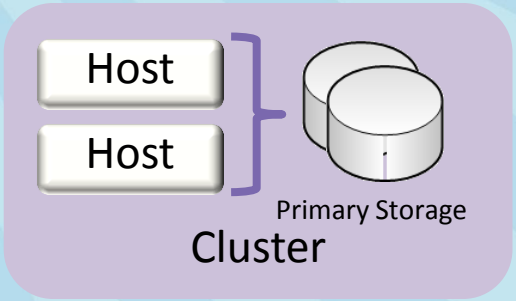

Secondary Storage Options

Option	XenServer	vSphere	KVM	LXC	Hyper-V
NFS	Yes	Yes	Yes	Yes	No
Swift(1)	Yes	Yes	Yes	Yes	No
S3 compatible (2)	Yes	Yes	Yes	Yes	No
SMB	No	No	No	No	Yes

(1) Requires NFS staging area
(2) Can be region wide, but must not have NFS secondary storage in zone

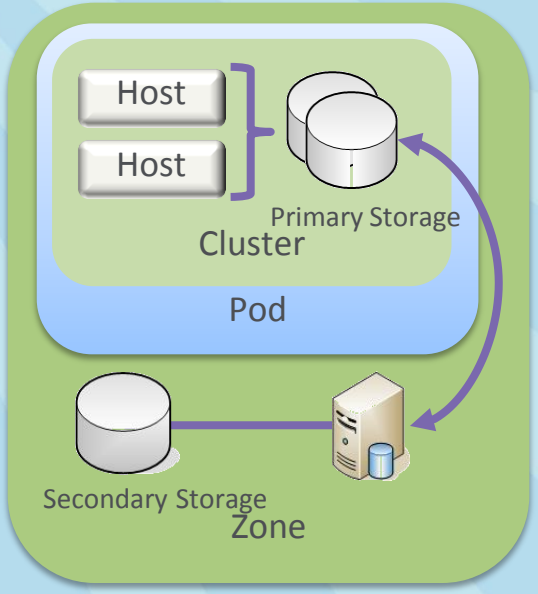

Core limits and features that matter

CloudStack Features

Feature	XenServer	vSphere	KVM	LXC	Hyper-V
Disk IO Statistics	Yes	No	Yes	No	Yes
Memory Overcommit	Yes (4x)	Yes	No	No	No
Dedicated resources	Yes	Not with HA/DRS	Yes	No	Yes
Disk IO throttling	No	No	Yes	Yes	No
Disk snapshot (running)	Yes	Yes	No	No	No
Disk snapshot (pluggable)	Partial	Partial	No	No	No
Disk snapshot (Stopped)	Yes	Yes	Yes	No	Yes
Memory snapshot	Yes	Yes	Yes	No	No
Zone wide primary storage	No	Yes	Yes	Yes	No
Resize disk	Offline	Online Grow	Online	No	No
High availability	CloudStack	Native	CloudStack	No	CloudStack
CPU sockets	6.2 and higher	Yes	Yes	Host count	Yes
Affinity groups	Yes	Yes	Yes	No	Yes

XenServer 6.2

Feature	
Source code model	Open Source (GPLv2)
Maximum VM Density	650 (Linux)
CloudStack VM Density	500
CloudStack integration	Direct XAPI calls
Maximum native cluster size	16
Maximum pRAM	1 TB
Largest VM	16vCPU/128GB
Windows Operating System	All Windows supported by Microsoft
Linux Operating Systems	RHEL, CentOS, Debian, Ubuntu, SLES, OEL
Advanced features supported	ovs, Storage XenMotion, DMC

vSphere 5.5

Feature	
Source code model	Proprietary
Maximum VM Density	512
CloudStack VM Density	128
CloudStack integration	vCenter
Maximum native cluster size	32
Maximum pRAM	4 TB
Largest VM	64 vCPU/1TB
Windows Operating Systems	DOS, All Windows Server/Client
Linux Operating Systems	Most
Advanced features supported	HA, DRS, vDS, Storage vMotion

KVM (RHEL/CentOS 6.5 and Ubuntu 12.04)

Feature	
Source code model	Open Source (GPLv2)
Maximum VM Density	10 times the number of pCores
CloudStack VM Density	50
CloudStack integration	CloudStack Agent (libvirt)
Maximum native cluster size	No native cluster support
Maximum pRAM	2 TB
Largest VM	160 vCPU/2TB
Windows Operating Systems	Windows XP and higher
Linux Operating Systems	Varies
Advanced features supported	None

Linux Containers

Feature	
Source code model	Open Source (GPLv2)
Maximum container Density	6000 (theoretical)
CloudStack container Density	50
CloudStack integration	CloudStack Agent (libvirt), requires KVM for SVMs
Maximum native cluster size	N/A
Maximum pRAM	2 TB
Largest container	2TB
Windows Operating Systems	N/A
Linux Operating Systems	Kernel compatible distros

Microsoft Hyper-V

Feature	
Source code model	Proprietary
Maximum VM Density	1024
CloudStack VM Density	1024
CloudStack integration	CloudStack Agent (C# calling WMI)
Maximum native cluster Size	64
Maximum pRAM	4 TB
Largest VM	64 vCPU/1TB
Windows Operating Systems	All Windows supported by Microsoft
Linux Operating Systems	RHEL, CentOS, Debian, Ubuntu, SLES, OEL
Advanced features supported	None

Picking the “Best One”

KVM

- Primary value proposition:
  - Low cost with available vendor support
  - Familiar administration model
  - Broad CloudStack feature set with active development
- Cloud use cases:
  - Linux centric workloads
  - Dev/test clouds
  - Web hosting
  - Tenant density which dictates SDN options
- Weaknesses:
  - Requires use of an installed CloudStack libvirt agent
  - Limited native storage options
  - No use of advanced native features

Linux Containers

- Primary value proposition:
  - Low cost with available vendor support
  - Familiar administration model
- Cloud use cases:
  - Dev/test clouds
  - Web application hosting
- Weaknesses:
  - Requires use of an installed CloudStack libvirt agent
  - Requires KVM for system VMs
  - No use of advanced native features
  - First introduced in CloudStack 4.2

Microsoft Hyper-V

- Primary value proposition:
  - Unlimited Windows Server VM licenses
  - Familiar Windows management paradigm
- Cloud use cases:
  - Windows and Linux workloads
  - Dev/test clouds
  - .Net application web hosting
  - Desktop as a Service clouds
- Weaknesses:
  - Minimal use of advanced native features
  - First introduced with CloudStack 4.3

vSphere

- Primary value proposition:
  - Broad application and operating system support
  - Readily available pool of vSphere administration talent
  - Large eco-system of vendor partners
  - Many CloudStack features are native implementations
  - Direct feature integration via vCenter
- Cloud use cases:
  - Private enterprise clouds
  - Dev/test clouds
- Weaknesses:
  - vSphere up-front license and ongoing support costs
  - vCenter integration requires redundant designs
  - Single data center per zone model

XenServer

- Primary value proposition:
  - Low cost with available vendor support
  - Broad CloudStack feature set with active development
  - Large CloudStack install base
  - Direct integration via XAPI toolstack
- Cloud use cases:
  - Linux centric workloads
  - Dev/test clouds
  - Web hosting
  - Desktop as a Service clouds
  - Large VM density and secure tenant isolation
- Weaknesses:
  - Minimal use of advanced native features

What About Multiple Hypervisor Support?

- Networking
  - Ensure network labels match
  - Topology is intersect of chosen hypervisors
- Storage
  - For system VMs to specific hypervisor type
  - Zone with primary storage limited
- Operations
  - vSphere Datacenter can not span zones
  - Hyper-V may not be mixed with other hypervisors
  - HA won't migrate between hypervisors
  - Capacity planning at the cluster/pod level more difficult

Tying it all Together

1. Define success criteria
2. Select a topology which works
3. Decide on storage options
4. Define supported configurations
5. Select preferred hypervisor(s)
6. Validate matrix
7. Build your Cloud

Name	Pod	Hypervisor	State
CloudZone145-1	Boston	XenServer	Enabled
CloudZone145-2	Boston	XenServer	Enabled
10.204.136.55/Bedford/vmw02	Boston	VMware	Enabled
10.204.136.55/Bedford/vmw01	Boston	VMware	Enabled
Showcase	Showcase	XenServer	Enabled

Name	Zone	Pod	Cluster	State
10.204.145.47	Zone 145	Boston	10.204.136.55/Bedford/vmw02	Up
mhtpmzone145-36	Zone 145	Boston	CloudZone145-2	Up
mhtpmzone145-32	Zone 145	Boston	CloudZone145-1	Up
mhtpmzone145-34	Zone 145	Boston	CloudZone145-1	Up
10.204.145.46	Zone 145	Boston	10.204.136.55/Bedford/vmw01	Up
mhtpmzone145-30	Zone 145	Boston	CloudZone145-1	Up
10.204.145.45	Zone 145	Boston	10.204.136.55/Bedford/vmw01	Up