z/VM Systems Management Overview

The Strategies and Options

September 6, 2015 v2.16

Bill Bitner
z/VM Customer Focus and Care
bitnerb@us.ibm.com

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

BladeCenter*	FICON*	Performance Toolkit for VM	Storwize*	System z10*	zSecure
DB2*	GDPS*	Power*	System Storage*	Tivoli*	z/VM*
DS6000*	HiperSockets	PowerVM	System x*	zEnterprise*	
DS8000*	HyperSwap	PR/SM	System z*	z/OS*	
ECKD	OMEGAMON*	RACF*	System z9*		

* Registered trademarks of IBM Corporation

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.
ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.
Java and all Java based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.
Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
OpenStack is a trademark of OpenStack LLC. The OpenStack trademark policy is available on the OpenStack website.
TEALEAF is a registered trademark of Tealeaf, an IBM Company.
Windows Server and the Windows logo are trademarks of the Microsoft group of countries.
Worklight is a trademark or registered trademark of Worklight, an IBM Company.
UNIX is a registered trademark of The Open Group in the United States and other countries.
* Other product and service names might be trademarks of IBM or other companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.

This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

This information provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g., zIIPs, zAAPs, and IFLs) ("SEs"). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the “Authorized Use Table for IBM Machines” provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html (“AUT”). No other workload processing is authorized for execution on an SE. IBM offers SE at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

Notice Regarding Specialty Engines (e.g., zIIPs, zAAPs and IFLs):

Any information contained in this document regarding Specialty Engines ("SEs") and SE eligible workloads provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g., zIIPs, zAAPs, and IFLs). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the “Authorized Use Table for IBM Machines” provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html (“AUT”).

No other workload processing is authorized for execution on an SE.

IBM offers SEs at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

Agenda

- Overview
  - Definition of “Systems Management”
  - Systems Management Disciplines
- Functional Options
  - z/VM Systems Management Related Interfaces
  - xCAT
  - OpenStack Enablement
  - IBM Wave for z/VM
- Cloud Strategy
- Product and Offering Survey, organized by Discipline

Overview

Definition of Systems Management

- Means different things to different people
  - What do you mean by “system”?
    z/VM? Linux? Application?
  - What do you mean by “manage”?
    Install? Configure? Change? Clone? Provision?
    Software/Code Management?
    Operational Monitoring or Performance Monitoring?

- “Any product or process that is not involved in actually using the system for the purpose it was purchased.” – Phil Smith III

- What it does **not** cover, at least in this presentation:
  - Initial installation & initial configuration of disks, memory, devices

- Systems Management needs may mature/expand over time.

Three Dimensions of Systems Management

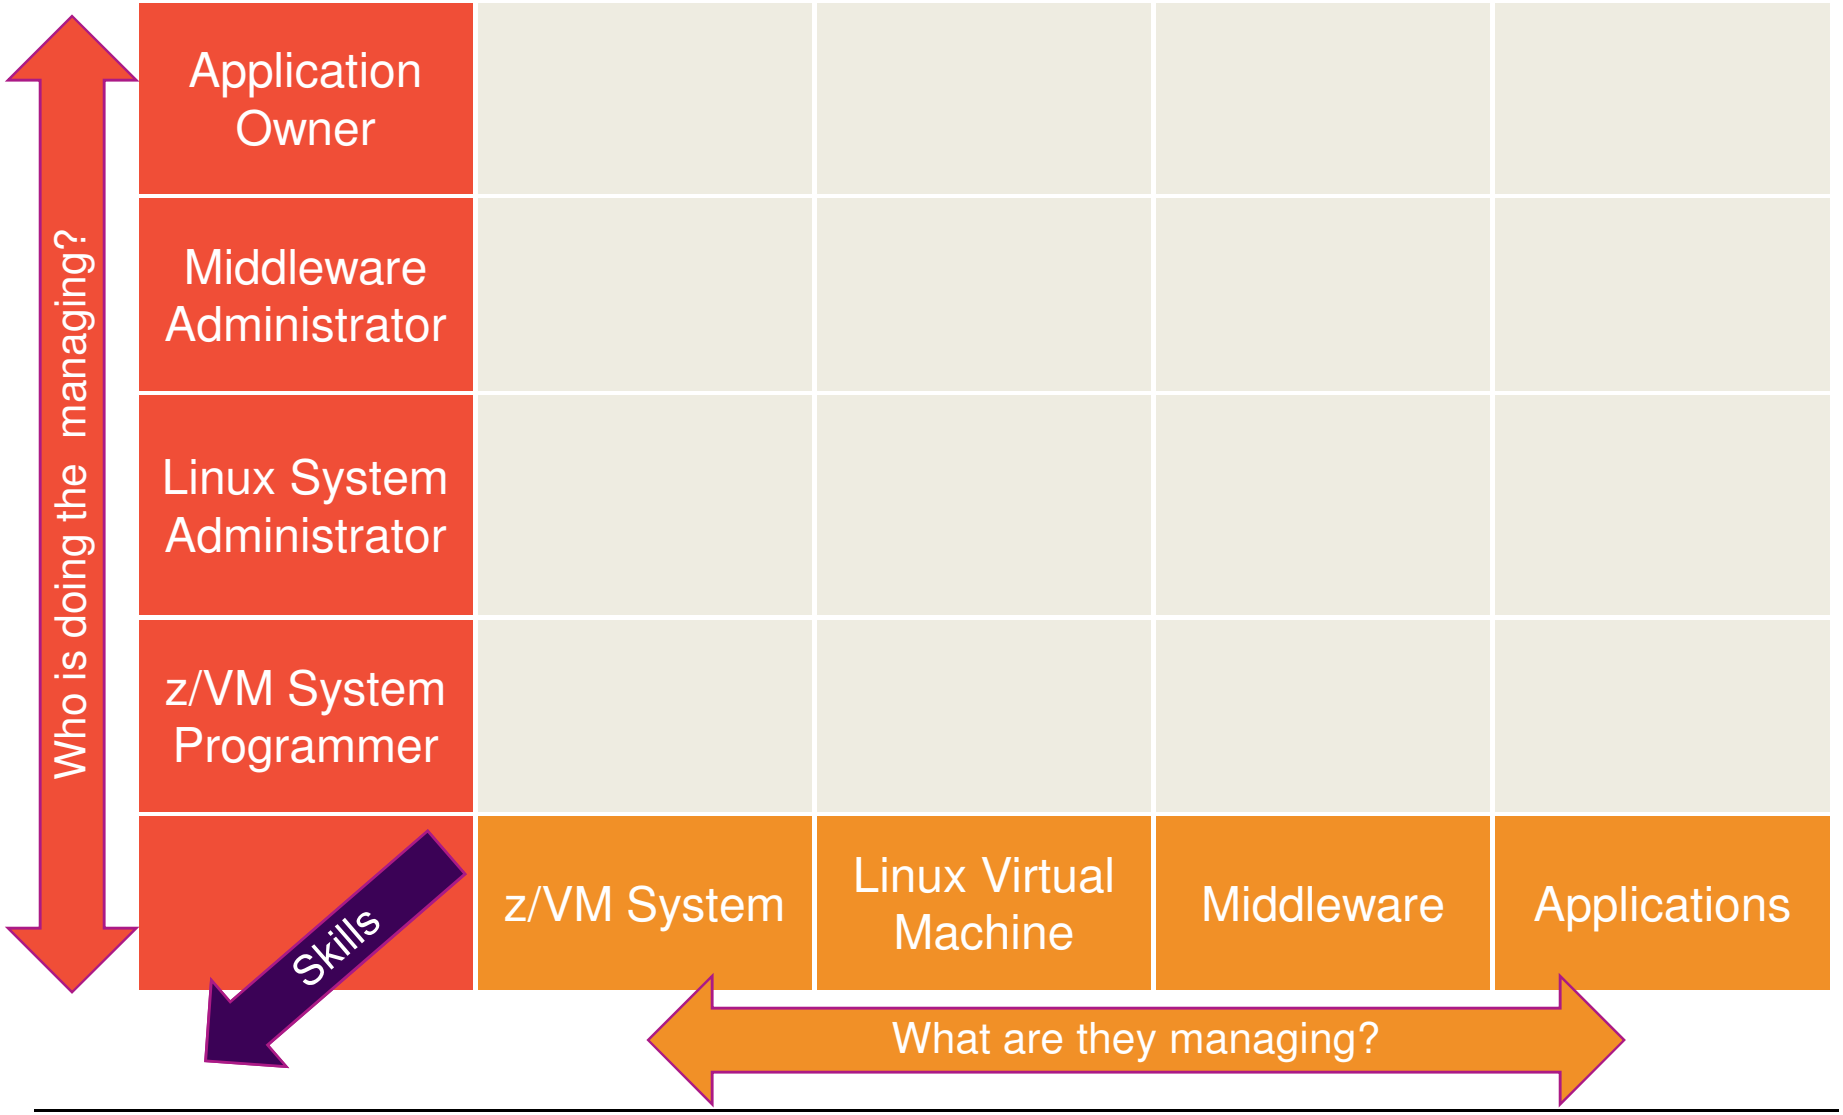

Your Investment Portfolio

Buy

- Solutions
- Integration
- Support

Do It Yourself

- Develop Solutions
- Integrate Solutions
- On going support

Disciplines of Systems Management

Automation	• Scheduling and running common events and processes
Monitoring	• Monitor and acting on events and conditions of the system
Business Continuity	• Disaster Recovery and Backups/Restores
Data Management	• Tape and storage management
Accounting	• Charge back and billing; Inventory
Security	• Enforcing permissions, tracking, and auditing.
Performance	• Performance Management and Capacity Planning
Provisioning	• Creating and changing the resources and configurations of virtual machines and perhaps guest systems

Interface Layers

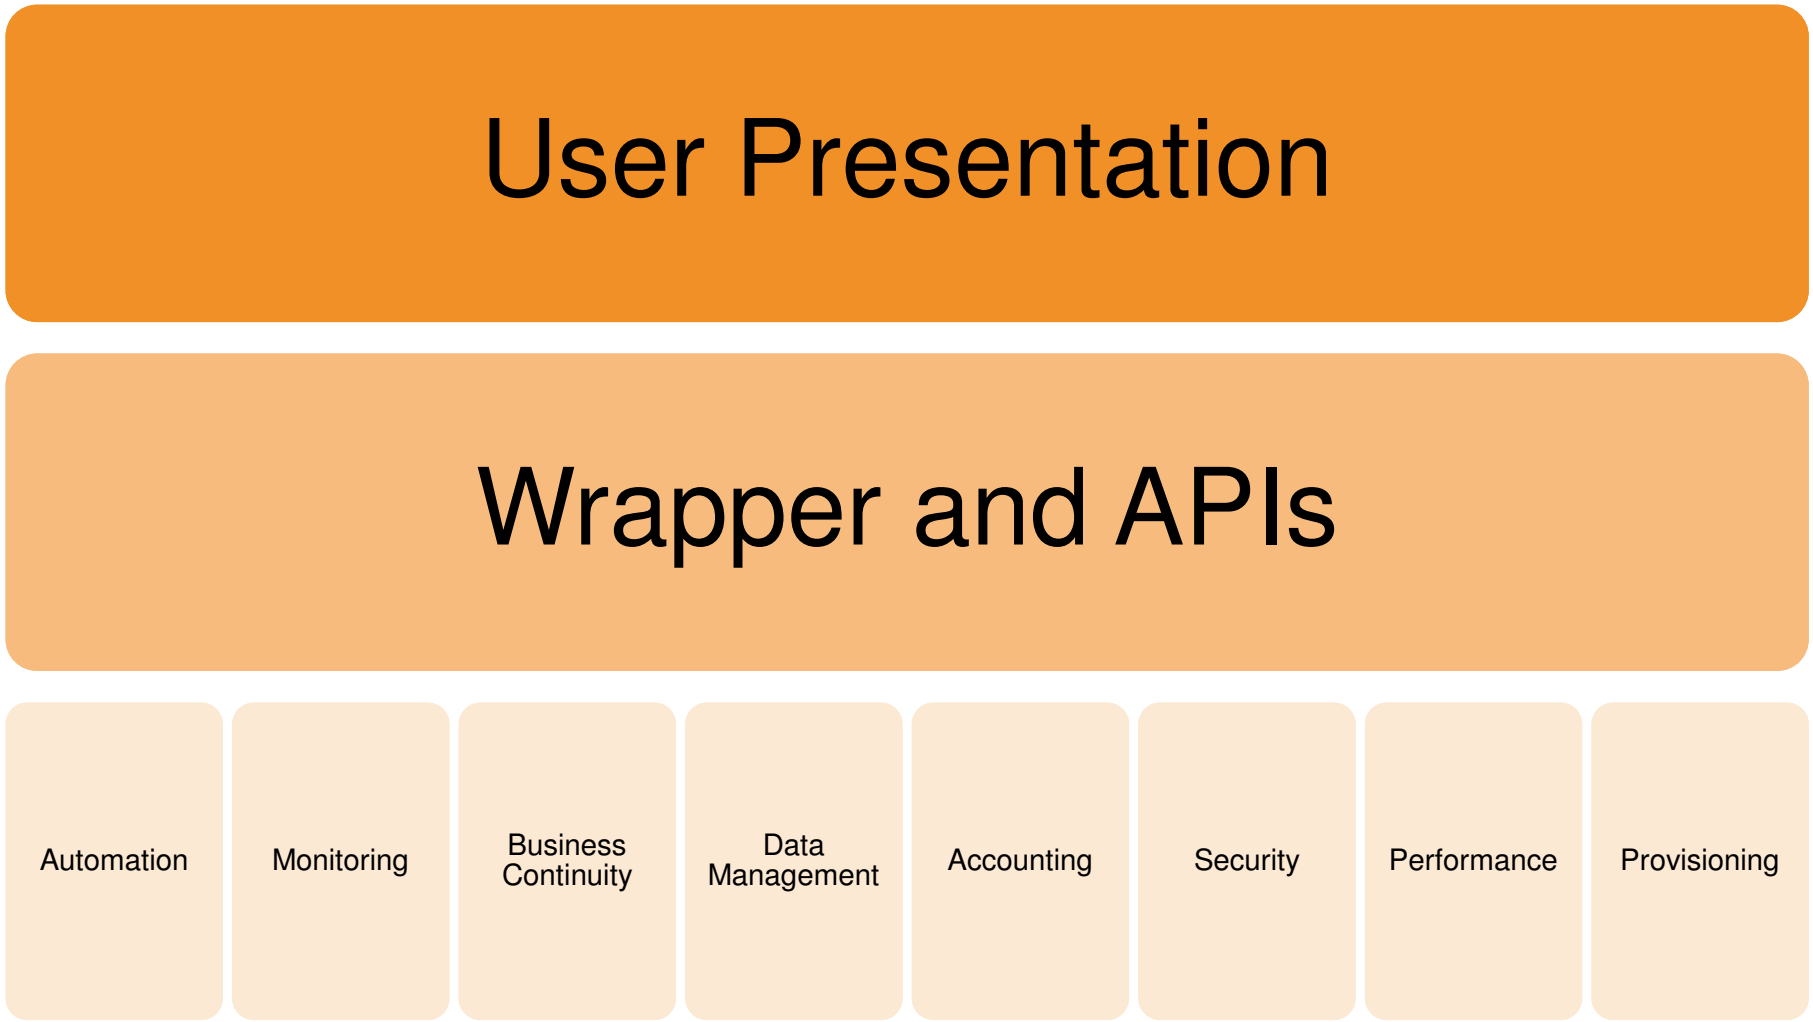

Functional Options

Systems Management Related APIs and Interfaces

- CP Exit Support
- CP System Services
  - *ACCOUNT
  - *MONITOR
  - *RPI
  - *VMEVENT
- Systems Management APIs (SMAPI)
- xCAT REST APIs
- OpenStack Enablement

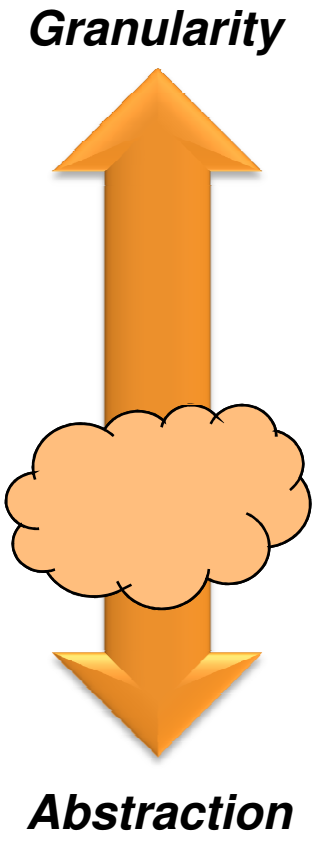

CP Exit Support

- CP Exit Support (aka “CP Exits”)
- Long History of customer and ISV extensions to the z/VM Hipervisor
- Typically used by ISVs and a small subset of customers with very specific modifications to the control program
- Can add or delete:
  - CP commands
  - Diagnose codes
  - Locally-developed CP message repositories
  - CP exit routines
- Dynamically loaded code
- 23 IBM defined CP Exit Points
  - Pre- and Post-processing Logon command
  - Shutdown command screening
  - Separator Page processing
  - Live Guest Relocation processing
- Defining and overriding commands
- z/VM CP Exit Customization SC24-6176

CP System Services

- Allow programmatic interface with the z/VM Control Program
- IUCV (Inter User Communication Vehicle) connection to CP
- Most common used by various systems management tools and products:
  - *ACCOUNT
    - Collects resource usage information for virtual machines and certain other audit like functions.
  - *MONITOR
    - Collects thousands of data points about performance on a system, device, and virtual machine basis
    - Uses memory shared between CP and the virtual machine application to share the actual data
  - *RPI – Access Verification System Service
    - Allows interaction between CP and an External Security Manager (ESM) such as Resource Access Control Facility (RACF)
  - *VMEVENT
    - Allows notification to an application of significant system events, such as virtual machines logging on/off and relocation
- Most are documented in the z/VM CP Programming Services Manual SC24-6179
  - *MONITOR is in z/VM Performance Manual SC24-6208

z/VM Systems Management APIs

- Intended for use by ISVs and IBM program products for managing the z/VM System
- Socket-based server replaced Remote Procedure Call server in z/VM 5.3
  - RPC Server support dropped in z/VM 6.1
- Three types of API functions are supported for managing the z/VM system and virtual machines (images):
  - IBM-supplied directory manager functions
  - IBM-supplied non-directory manager functions
  - Customer-defined functions.

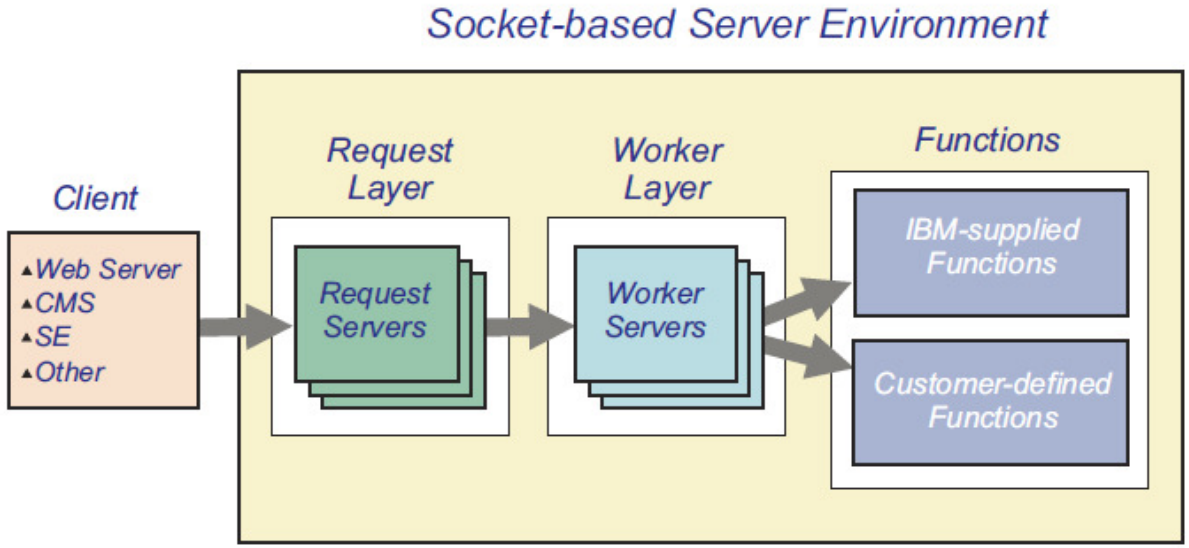

Categories of SMAPI functions

- ABEND Dump Management
- Authorization
- Directory Manager Control
- Directory Manager Local Tag and Scan
- Directory Parsing
- Directory Update Subscription
- Disk Management
- Event Management
- Image Characteristics
- Image Connectivity
- Image CPUs
- Image Devices
- Image IPL Management
- Image Operations
- Image Volume Management
- List-Directed IPL
- Name List
- Profile Management
- Prototype Management
- Response Recovery
- Server
- Shared Memory Management
- Single System Image (SSI) Cluster Management
- System Configuration
- VMRM Configuration Update

xCAT – Extreme Cloud Administration Toolkit

- Open Source tool to manage, provision, and monitor physical and virtual machines on all IBM architectures.
- Made available as a download and customer install in July 2012
- Integrated into the base of z/VM 6.3 in July 2013
- xCAT has four different interfaces
  - REST APIs (Representational State Transfer) – used by OpenStack Solutions
  - Browser based Graphical User Interface
  - Command Line Interface (CLI)
  - XML

Where does xCAT fit in?

- It provides the tools you need to administer your cloud
  - Automate redundant steps, e.g. Linux installation
  - Manage virtual resources

- xCAT is not a cloud operating system
  - You still need to tell xCAT what to do
  - You still need a z/VM system administrator

- xCAT is a toolkit
  - Equivalent to virt-manager or virsh for KVM
  - Tool to manage, provision, and monitor physical and virtual machines on IBM System x®, Power Systems®, and IBM z Systems
  - Provides rudimentary cloud functionality

- Value
  - Low overhead and lightweight (download ~100 MB)
  - Open source (EPL)
  - Cross platform

xCAT History

- Developed by IBM starting in 1999 to manage clusters of Linux and AIX systems.
- Released to the Open Source community under the Eclipse Public License
- Used to manage clusters such as
  - Roadrunner, the fastest computer in the world in 2008
  - Watson, the IBM machine that competed and won on Jeopardy
- Today xCAT can manage physical or virtual machines, such as: RHEL, CentOS, Fedora, SLES, AIX, Windows, VMWare, KVM, PowerVM, z/VM.

Source: http://sourceforge.net/apps/mediawiki/xcat/index.php?title=Main_Page

Examples of Supported xCAT Features

- Automation
  - Power on/off virtual machine

- Provisioning virtual machine and host
  - Create/change/delete a virtual machine
  - Add/delete virtual network devices to a virtual machine
  - Provisioning diskless virtual machine using an NFS read-only root filesystem
  - Add/delete ECKD/SCSI disk to a virtual machine
  - Add/delete ECKD/SCSI disk to the system disk pool
  - Create/change/delete layer 2 or 3 QDIO Guest LAN or Vswitch
  - Cloning virtual machines
  - Plain installation of Linux via AutoYast or Kickstart

Examples of Supported xCAT Features

- Accounting
  - Software and hardware inventory of a virtual machine
- Performance
  - System resource inventory (e.g. disks, networks)
- Business Continuity
  - Live Guest Relocation

z/VM 6.3 Pre-Installed Systems Management

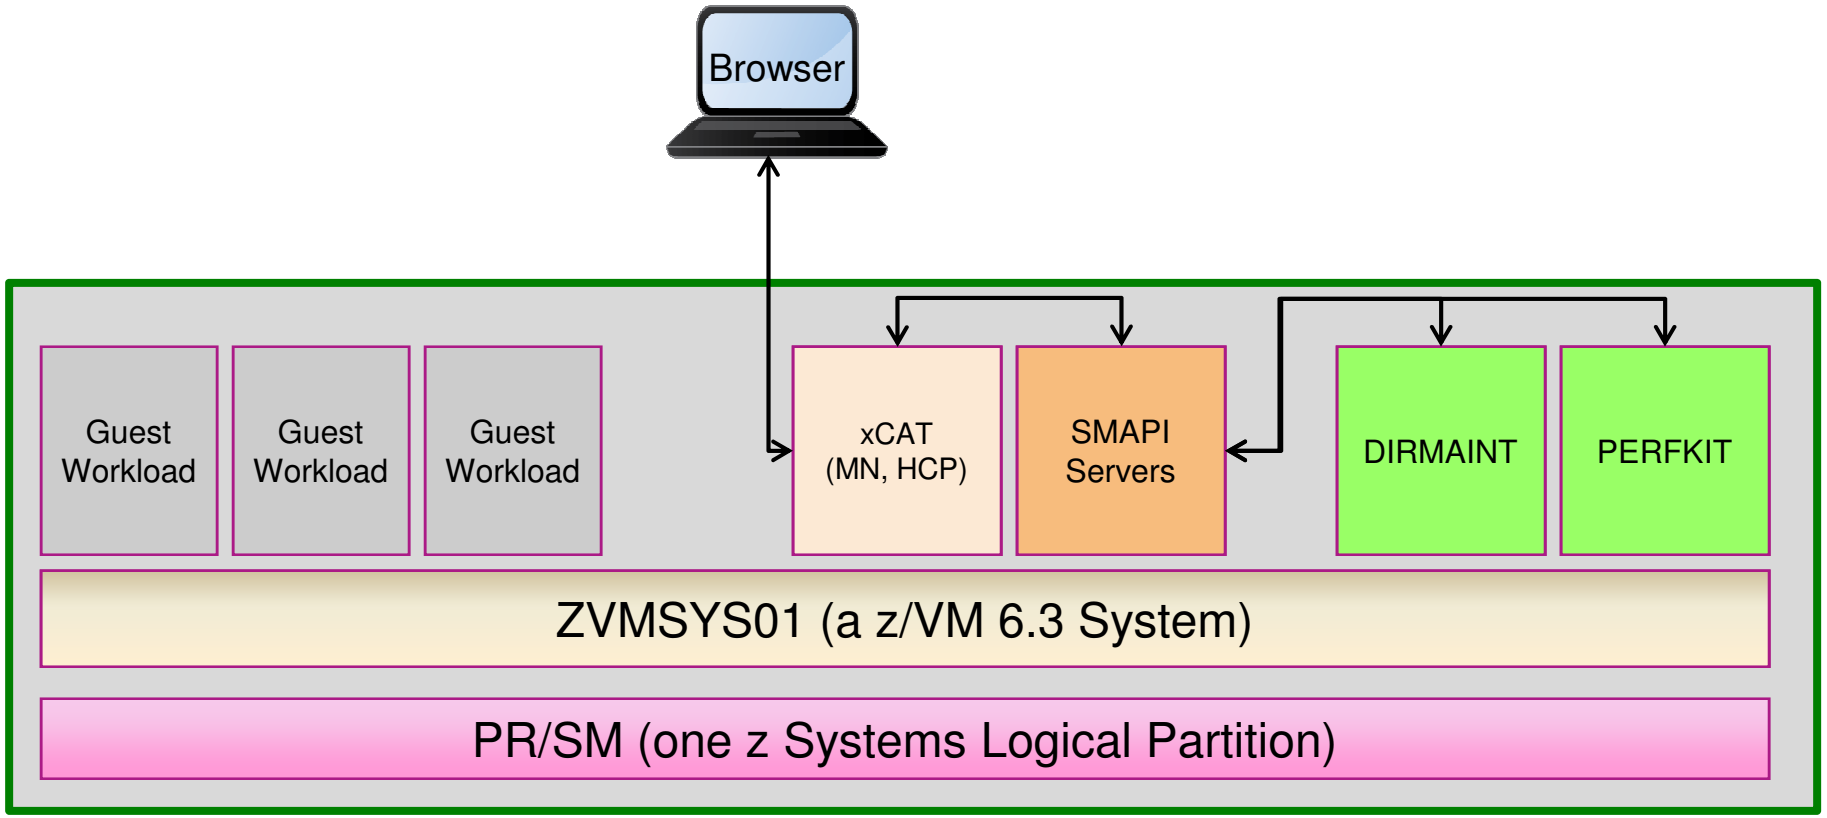

- Everything inside the z/VM LPAR is shipped with z/VM 6.3, up to and including the xCAT interfaces
- Versions of DirMaint and Performance Toolkit are included free of charge with the SMAPI server, but these versions of the products only communicate with SMAPI, there is no way to interact with them directly. It is generally recommended to use the full function versions of DirMaint and Performance Toolkit.

How xCAT Manages z/VM

zHardware Control Point: Manages other VMs via Systems Management APIs and CP Commands.
Each z/VM system needs to have a zHCP

xCAT Maintenance Node: Central management server. Only one MN is needed for multiple systems.

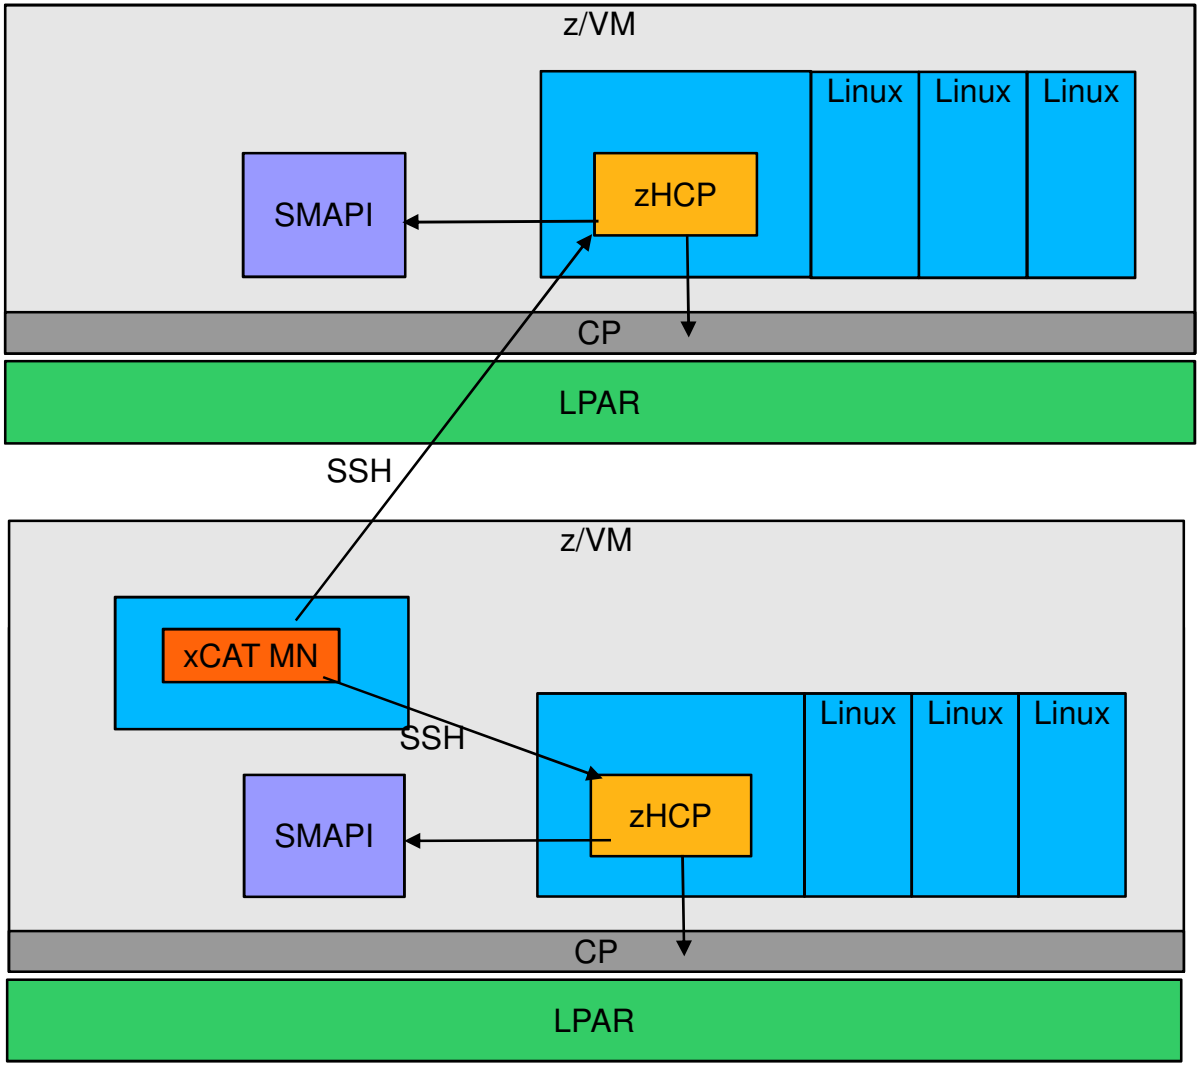

xCAT Resources

- xCAT wiki for z/VM – Information on using xCAT
  - http://sourceforge.net/apps/mediawiki/xcat/index.php?title=Main_Page
- xCAT YouTube channel: Various Demonstrations
  - http://www.youtube.com/user/xcatuser
- xCAT user mailing list
  - xcat-user@lists.sourceforge.net
- z/VM Home Page
  - http://www.vm.ibm.com/sysman/
- z/VM xCAT Maintenance
  - http://www.vm.ibm.com/sysman/xcmntlvl.html
- xCAT Downloadable References
  - http://sourceforge.net/apps/mediawiki/xcat/index.php?title=Download_xCAT
  - http://sourceforge.net/apps/mediawiki/xcat/index.php?title=XCAT_Documentation
  - http://xcat.sourceforge.net/pdf/xCAT-UI-Tutorial.pdf
  - http://sourceforge.net/apps/mediawiki/xcat/index.php?title=XCAT_zVM_Setup

IBM Wave for z/VM V1.2

Helps Simplify and Automate Virtualization Management

For z/VM and Linux virtual servers

- Automate, simplify management and monitor virtual servers and resources-all from a single dashboard
- Perform complex virtualization tasks in a fraction of the time compared to manual execution
- Provision virtual resources (Servers, Network, Storage) to accelerate the transformation to cloud infrastructure
- Supports advanced z/VM® management capabilities such as Live Guest Relocation with a few clicks
- Delegate responsibility and provide more self service capabilities to the appropriate teams

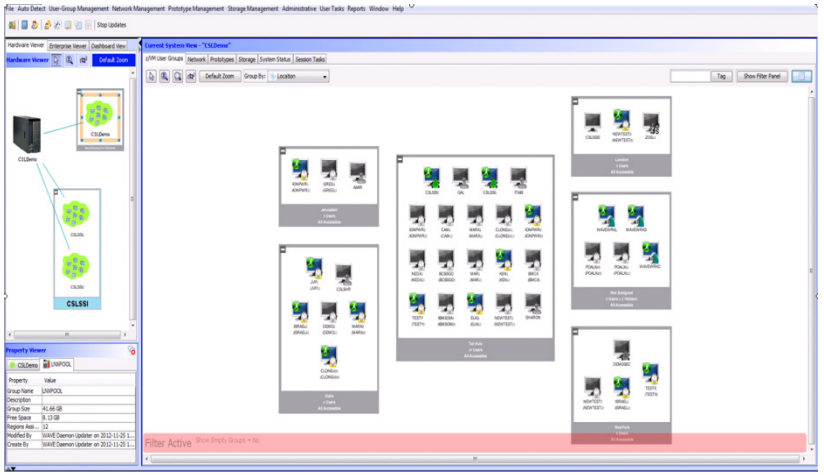

A simple, intuitive virtualization management tool providing management, provisioning, and automation for a z/VM environment supporting Linux® virtual servers

Extend the Reach of Skills with IBM Wave for z/VM

Intelligent Visualization

- Shorten the learning curve needed to manage complex environments
- Organize and simplify management of z/VM and virtual Linux servers
- View servers and storage utilization graphically; understand the status of system resources with Intelligent icons
- Reduce unnecessary steps using highly customizable views
- Graphical or tabular displays with layered drill down

Simplified Monitoring

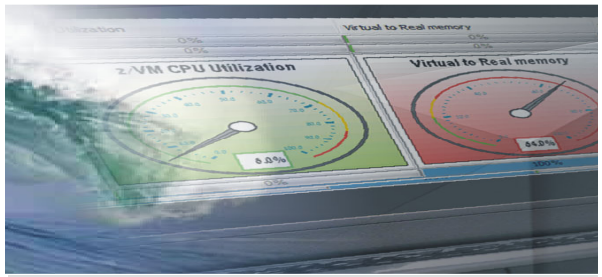

- Monitor the status of z/VM systems through an innovative interface
- Monitor performance of CPU, paging devices, spool disks and more;
- Use agentless discovery to detect an accurate view of your environment
- Use advanced filters, tagging, layout and layer selection to make monitoring and management more meaningful
- Complements IBM OMEGAMON® XE used for in-depth performance monitoring

Unified Management

- Manage your system from a single point of control
- Assign and delegate administrative access with role based assignments
- Provision, clone, and activate virtual resources . Define and control virtual network and storage devices
- Perform management tasks such as live guest relocation
- Annotate resources for additional policy based management
- Execute complex scripts with a single mouse click

IBM Wave for z/VM vs Operations Manager for z/VM

- IBM Wave for z/VM provides an interactive GUI interface for:
  - Provisioning of Linux guests
  - Basic performance information
  - Monitoring of virtual server resources
- Operations Manager for z/VM provides operational monitoring & automation
  - In the background
    - Monitoring of console messages for z/VM service machines and Linux guests
    - Monitoring “state” information for z/VM service machines and Linux guests
    - Monitoring spool and page space on the z/VM system
    - Automated responses to these monitors when they are triggered
      - Email
      - SNMP alerts
      - Integration with IBM Tivoli Netcool/OMNIbus enterprise alert system
      - Actions that address the problem immediately in addition to or instead of alert notificat
  - Interactive when needed
    - View and interact with live service machine and Linux guest consoles
    - View and manage spool files
- Complementary solutions
  - Use Operations Manager to monitor the IBM Wave service machines
  - Use Operations Manager to automatically initiate tasks in Wave via the Wave CLI

A Different Better vs. A Common Good

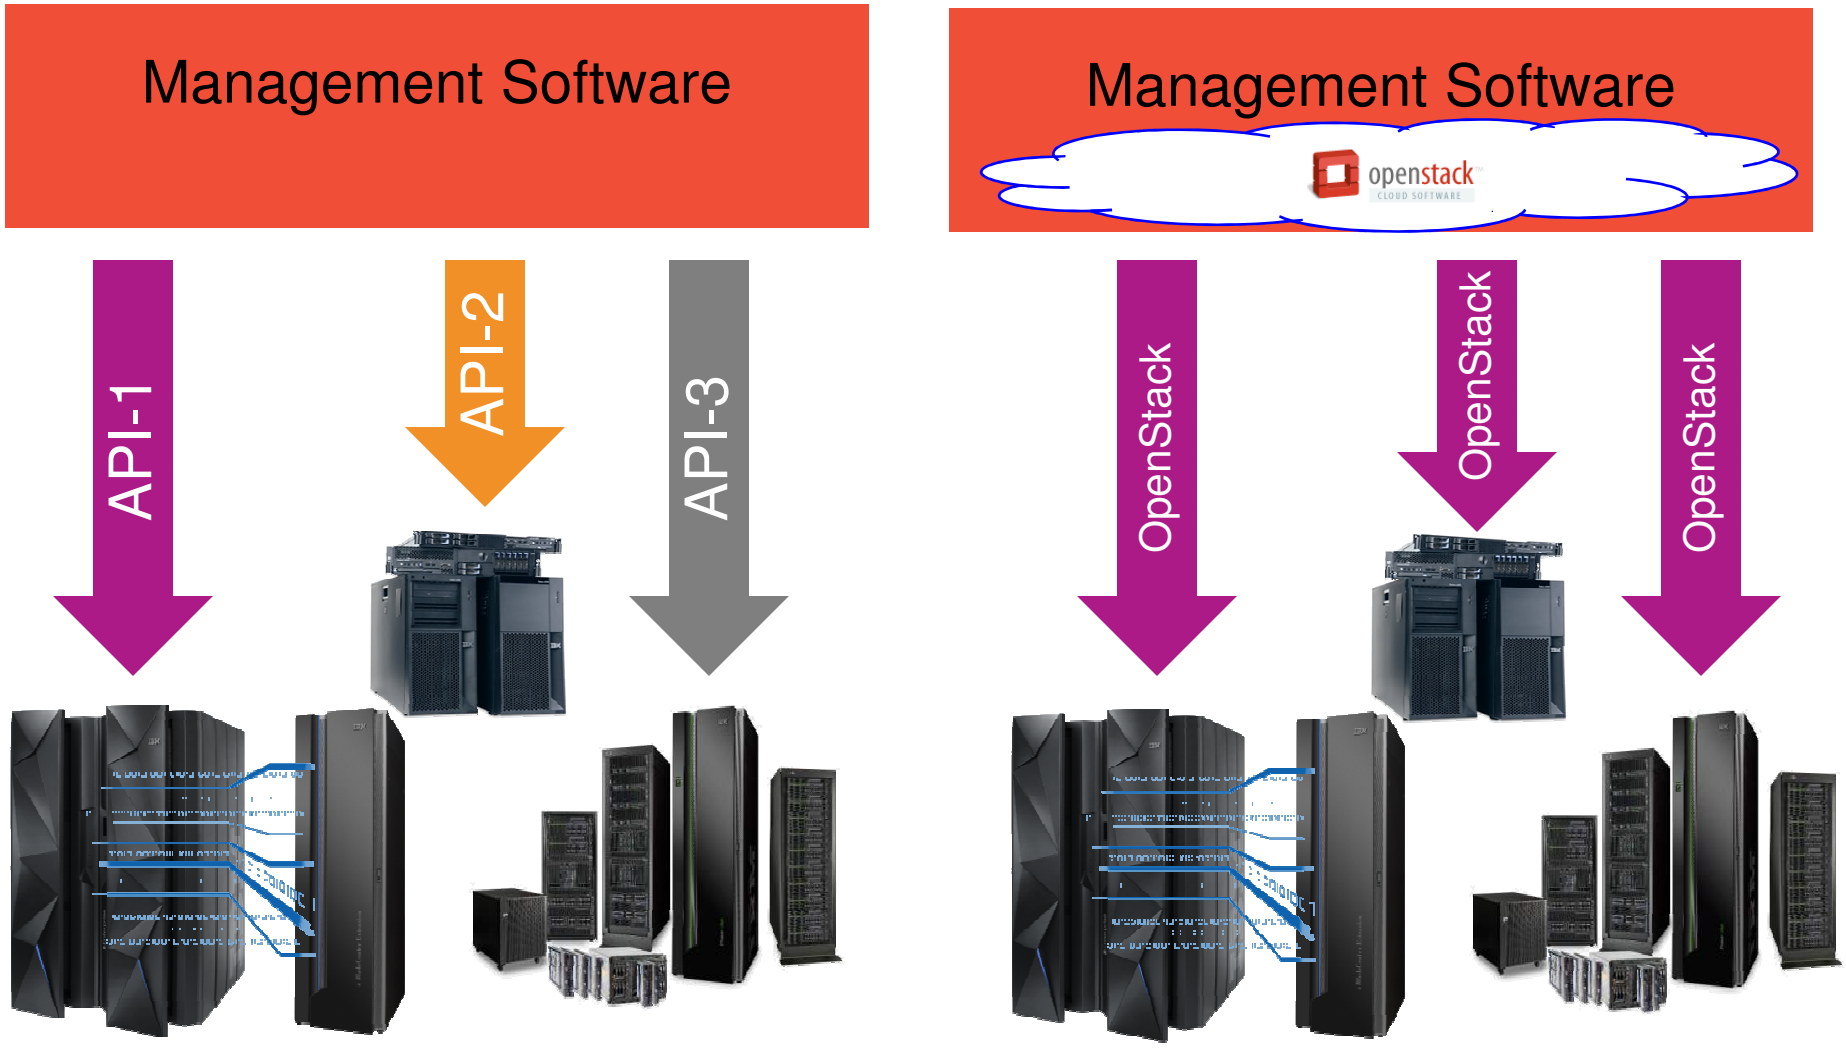

The OpenStack Food Chain

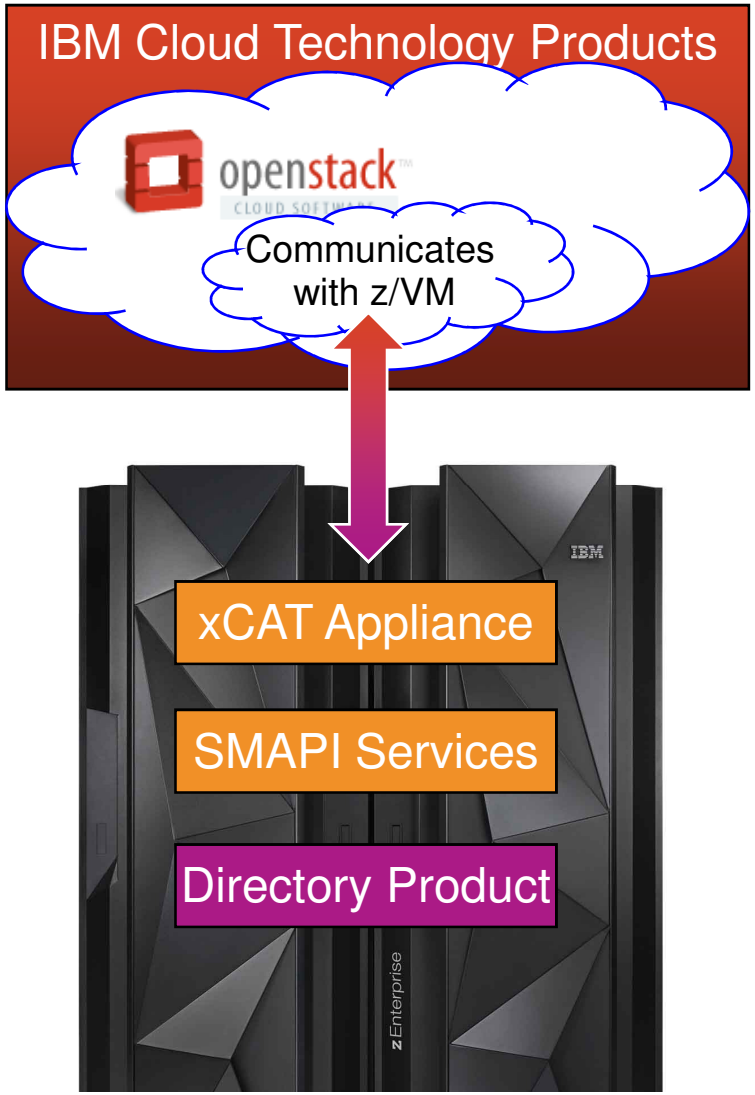

- **Top Half of the Solution:**
  - An IBM Cloud Technology product or other vendor product will include the OpenStack support.
  - Portions of that OpenStack support will know z/VM (i.e. code that connects and understands how to talk to z/VM).
- **Bottom Half of the Solution:**
  - Rest APIs are used to communicate with the OpenStack code from the top half.
  - The xCAT Appliance utilizes new and existing Systems Management APIs (SMAPI) to interact with the z/VM system
  - SMAPI can interact with additional products or features (e.g. a directory manager).

Product with OpenStack Support

z/VM 6.3 Product

Additional Product or Feature

z/VM 6.3 Pre-Installed Systems Management

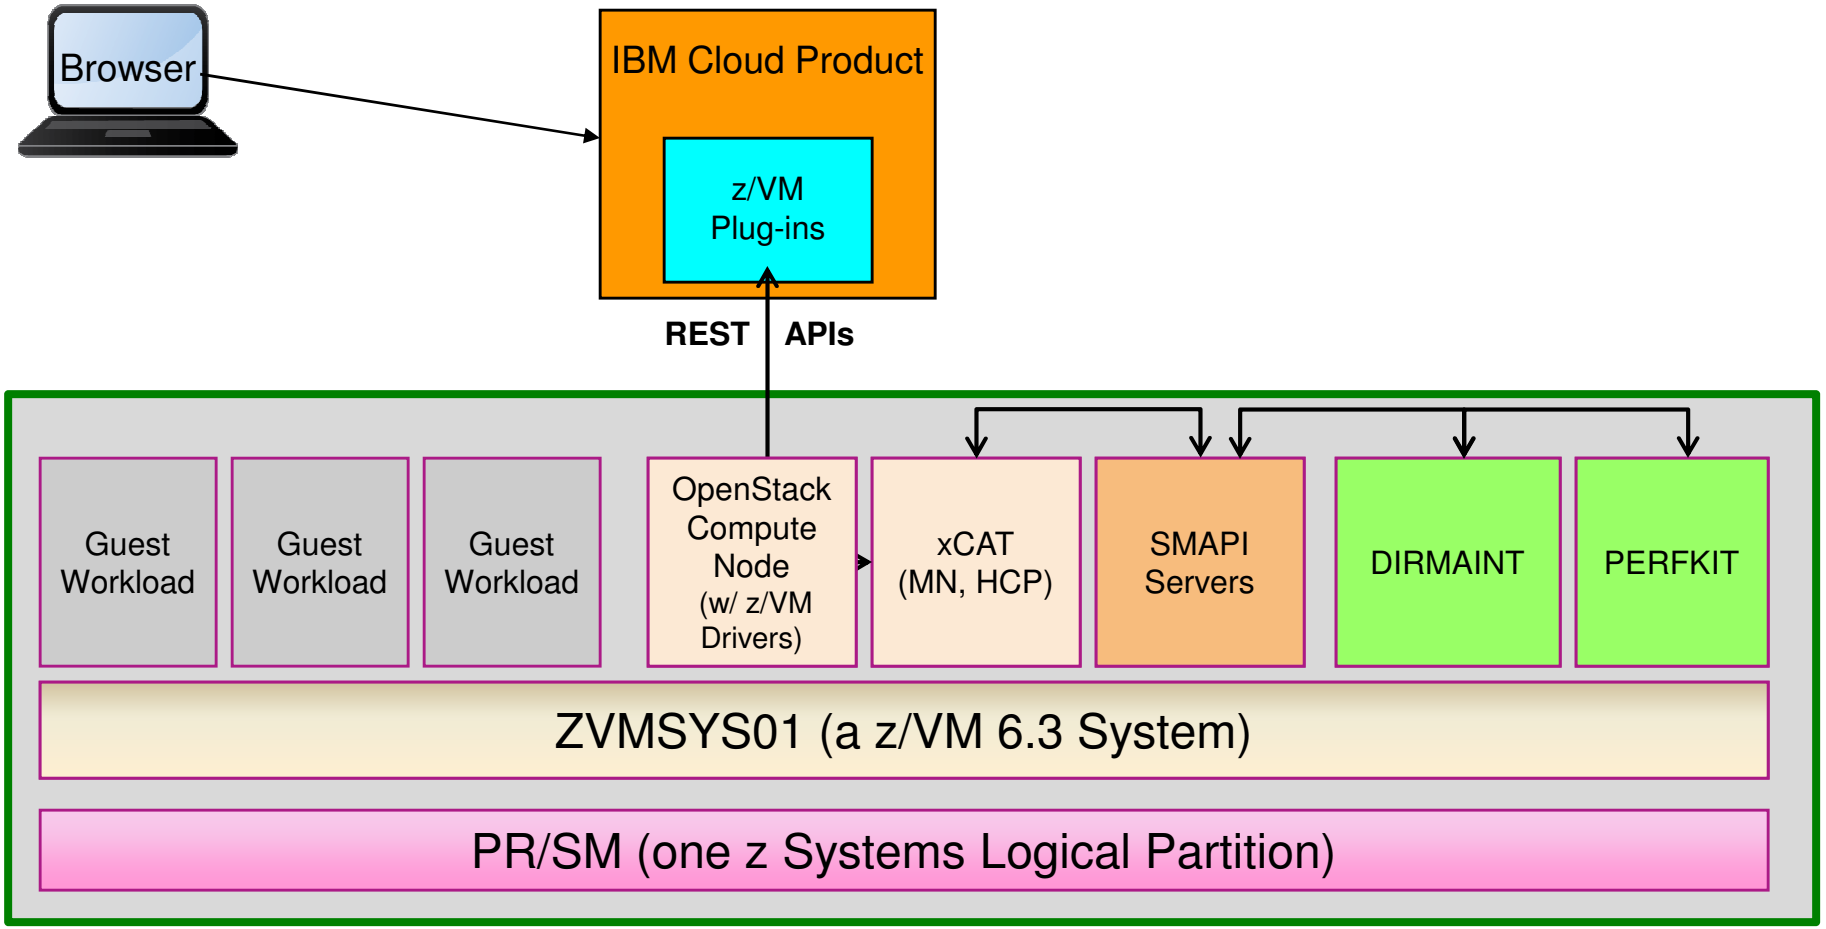

- Everything inside the z/VM LPAR is shipped with z/VM 6.3, up to and including the xCAT interfaces
- Versions of DirMaint and PerfKit are included free of charge with the SMAPI server, but these versions of the products only communicate with SMAPI, there is no way to interact with them directly.

OpenStack

- Open Source project to provide Infrastructure as a Service.
- Started by NASA and Rackspace in 2010
- Now backed by IBM and many other corporations
- Consists of separate projects to handle different types of resources
- New releases every 6 months

Source: http://sourceforge.net/apps/mediawiki/xcat/index.php?title=Main_Page

OpenStack Community

More than 6000 people and 100 companies

Active online community through mailing lists, IRC, wiki

Bi-yearly design summits

Companies need to donate money AND people that ACTIVELY contribute

Platinum Members

AT&T

Canonical

HP

IBM

Nebula

Rackspace

Red Hat, Inc.

SUSE

and many more …
http://www.OpenStack.org/foundation/companies/

OpenStack Release Names

These codenames are chosen by popular vote. Codenames are cities or counties near where the corresponding OpenStack design summit took place, with some exceptions to the rule.

- **Austin:** The first design summit took place in Austin, TX
- **Bexar:** The second design summit took place in San Antonio, TX
- **Cactus:** Cactus is a city in Texas
- **Diablo:** Diablo is a city in the bay area near Santa Clara, CA
- **Essex:** Essex is a city near Boston, MA
- **Folsom:** Folsom is a city near San Francisco, CA
- **Grizzly:** Grizzly is an element of the state flag of California, summit is in San Diego, CA
- **Havana:** Havana is an unincorporated community in Oregon with design summit nearby
- **Icehouse:** Ice House is a street in Hong Kong
- **Juno:** Juno is a locality in Georgia
- **Kilo:** Short for “kilogram”, the last SI base unit tied to a reference artifact (stored in Sevres)
- **Liberty:** A village in the Canadian province of Saskatchewan

Top Ten OpenStack Project Contributors

Havana September 2013		Icehouse April 2014		Juno September 2014		Kilo April 2015	
Company	**Commits**	**Company**	**Commits**	**Company**	**Commits**	**Company**	**Commits**
Red Hat	1858	Red Hat	3475	HP	3223	HP	5365
HP	1456	HP	2174	Red Hat	2859	Red Hat	4911
Rackspace	1160	IBM	1801	IBM	1407	Mirantis	3040
IBM	972	Rackspace	1730	Mirantis	1371	IBM	2446
Mirantis	362	Mirantis	1478	Rackspace	1280	Rackspace	1966
OpenStack Foundation	295	SUSE	985	*independent	898	*independent	1056
SUSE	279	OpenStack Foundation	716	SUSE	762	Yahoo!	818
eNovance	156	eNovance	704	eNovance	414	VMWare	807
VMware	147	*independent	686	B1 Systems	410	OpenStack Foundation	723
NEC	128	VMware	554	VMware	405	Intel	623

OpenStack Programs

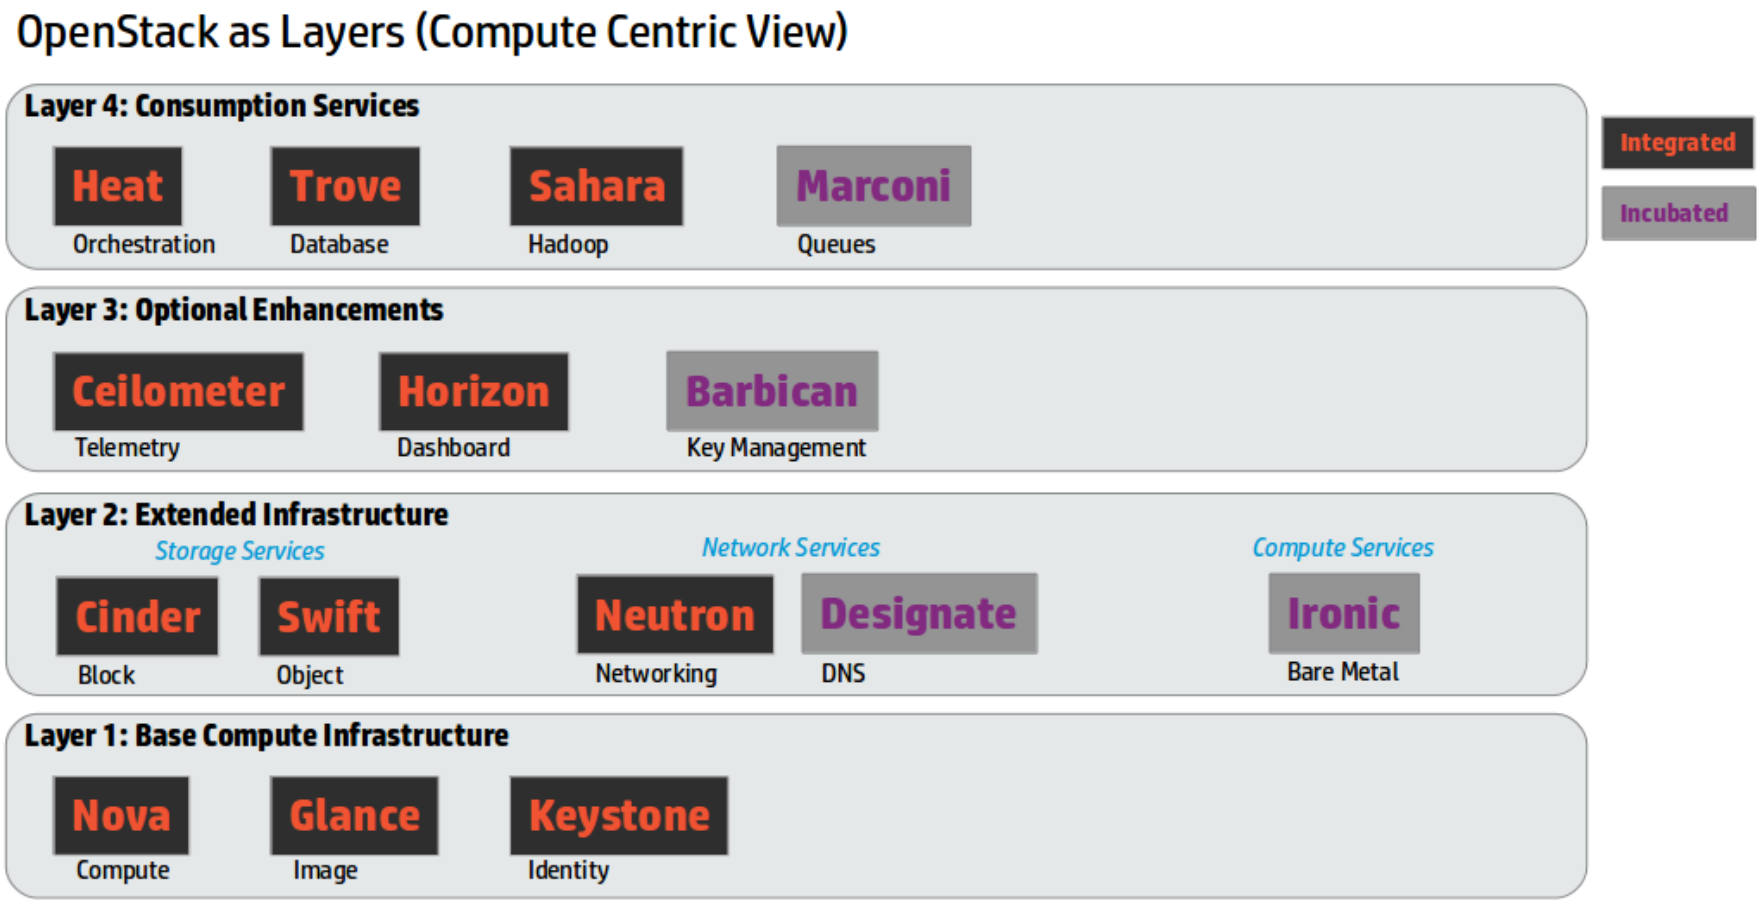

http://hackstack.org/x/blog/2013/09/05/openstack-seven-layer-dip-as-a-service/

OpenStack Programs

- Compute (Nova)
- Block Storage (Cinder)
- Network (Neutron)
  - Provision and manage virtual resources
- Dashboard (Horizon)
  - Self-service portal
- Image (Glance)
  - Catalog and manage server images
- Identity (Keystone)
  - Unified authentication and authorization
- Object Storage (Swift)
  - Petabytes of secure, reliable object storage
- Telemetry (Ceilometer)
  - Data collection
- Orchestration (Heat)
  - Engine to launch cloud applications based on templates
- Database Service (Trove)
  - Cloud Database-as-a-Service
- Data Processing (Sahara)
  - Data processing stack and management

OpenStack Deliverables and z/VM

- OpenStack releases are every 6 months
  - Expect z/VM service and releases to open source outside the normal z/VM release cadence
- The framework for support of OpenStack communication is in base of z/VM 6.3
- Initial z/VM OpenStack enablement (part on the z/VM side) available as service in December 2013
- Original IBM cloud technology for z/VM was based on Havana level by incorporating code z/VM Development has released to open source.

OpenStack Infrastructure Management Software

Simple Console
Built using OS REST API
Basic GUI for OS functions

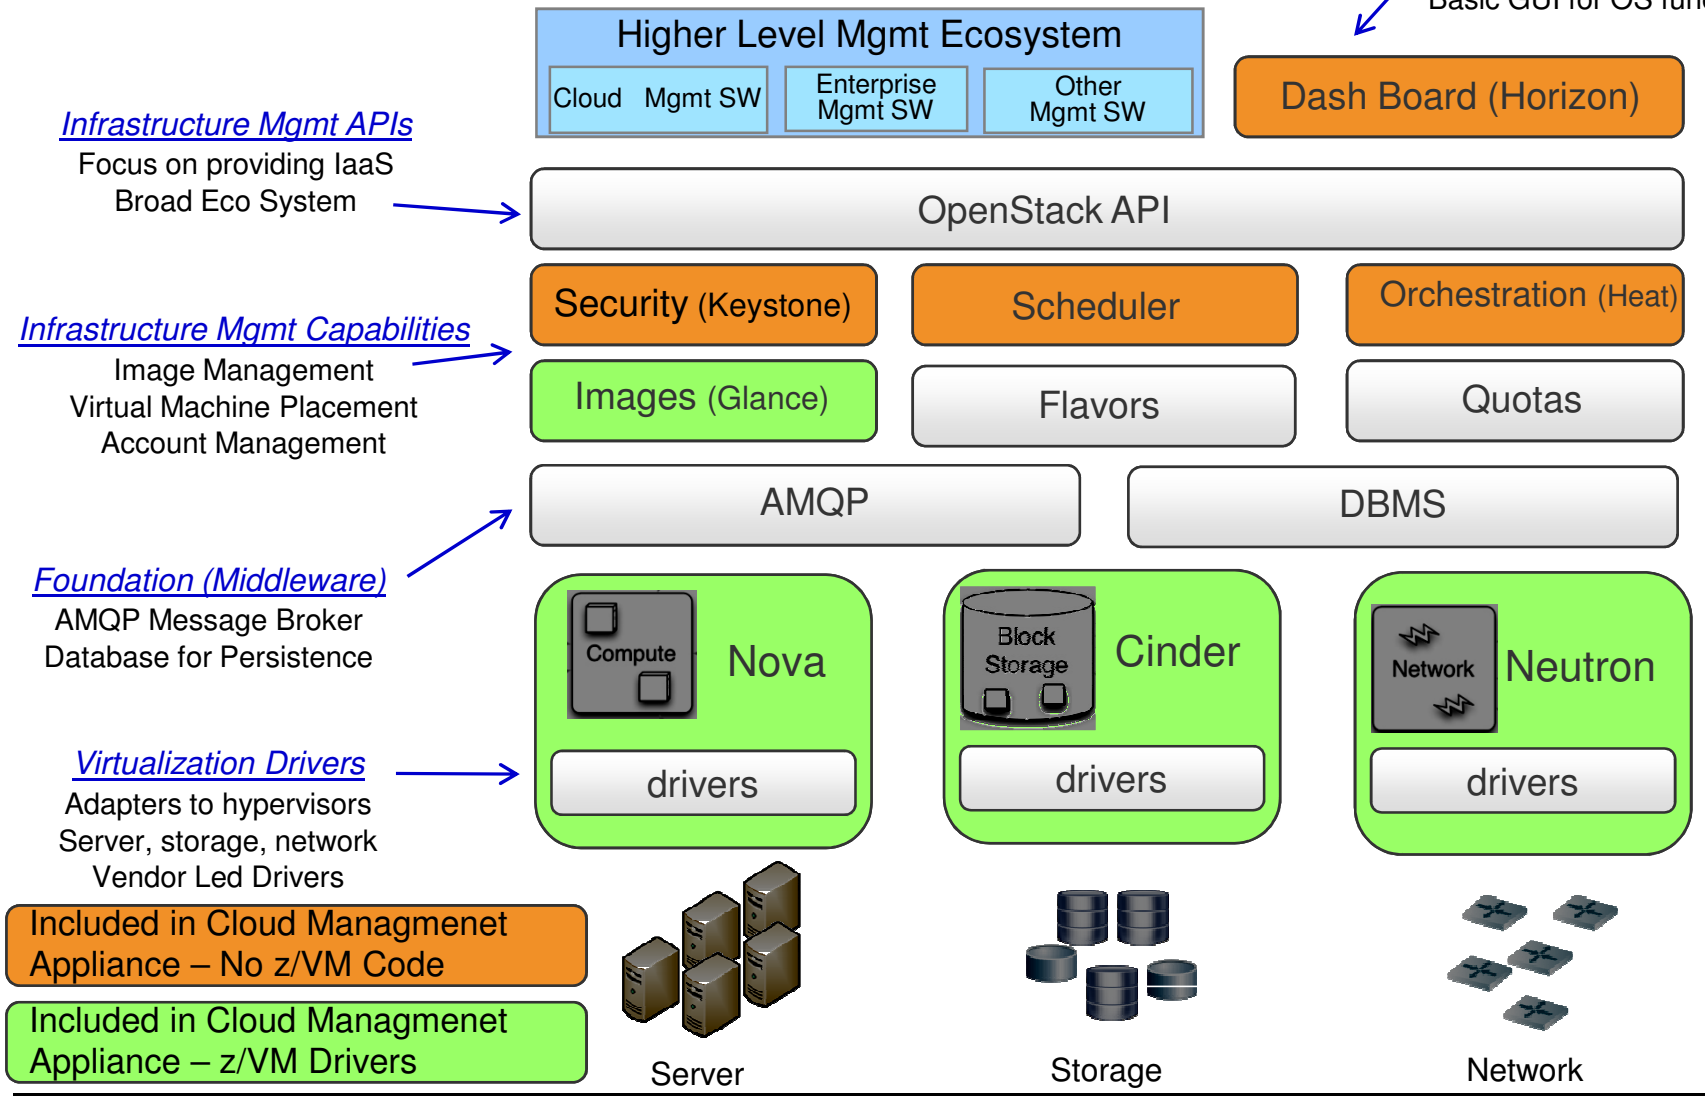

Examples of Supported z/VM OpenStack Features

- Provisioning virtual machines and host
  - Resize virtual machine (memory, CPU) [Nova]
  - Disk (Add SCSI disk to virtual machine) [Nova,Cinder]
  - Support for Open vSwitch [Neutron[1]]

- Automation
  - Start / Stop virtual machine [Nova]
  - Reboot Linux virtual machine [Nova]
  - Pause / Unpause virtual machine [Nova]
  - Capture / Deploy virtual machine [Nova, Glance]
  - Activate Image [Nova]

- Business Continuity
  - Live Guest Relocation [Nova]

[1] – Formerly known as Quantum.

Cloud Strategy

Cloud Computing - Based on Virtualization and Standardization

To position the various technologies in this space, we need to first understand that Cloud computing is a journey beginning with virtualization and consolidation of environments and ending with workload pattern-based deployment of IT services.

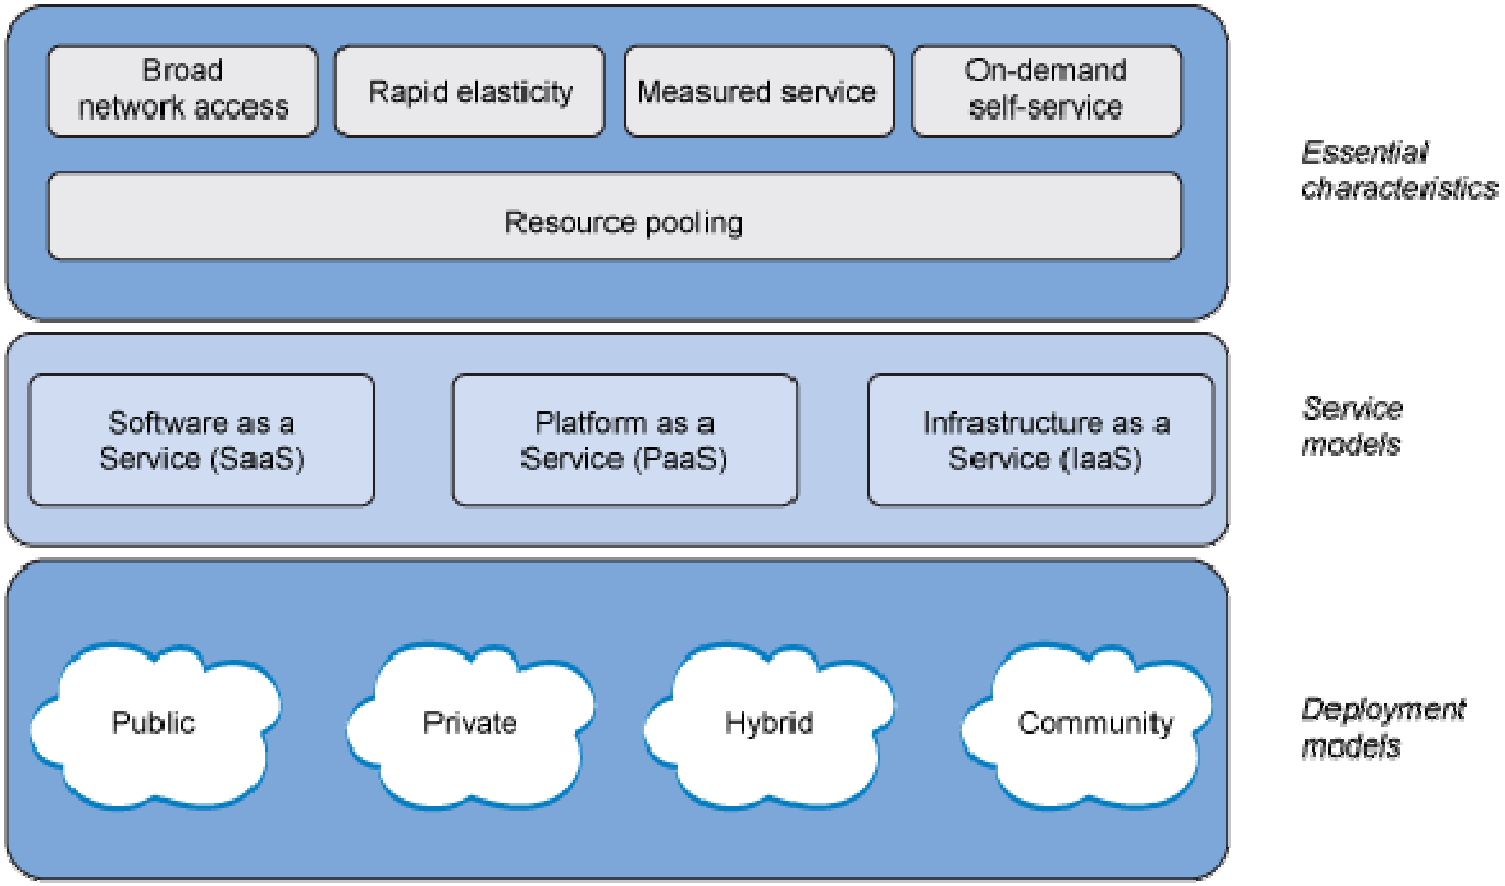

Cloud Computing - Based on Virtualization and Standardization

To position the various technologies in this space, we need to first understand that Cloud computing is a journey beginning with virtualization and consolidation of environments and ending with workload pattern-based deployment of IT services.

This is not always a step-wise progression. Some clients begin by optimizing their virtualization foundation for a workload, then gradually move to cloud.

Others require capabilities from the beginning and may start with entry or advanced cloud solutions.

Not all customers will want/need to get to the final step!

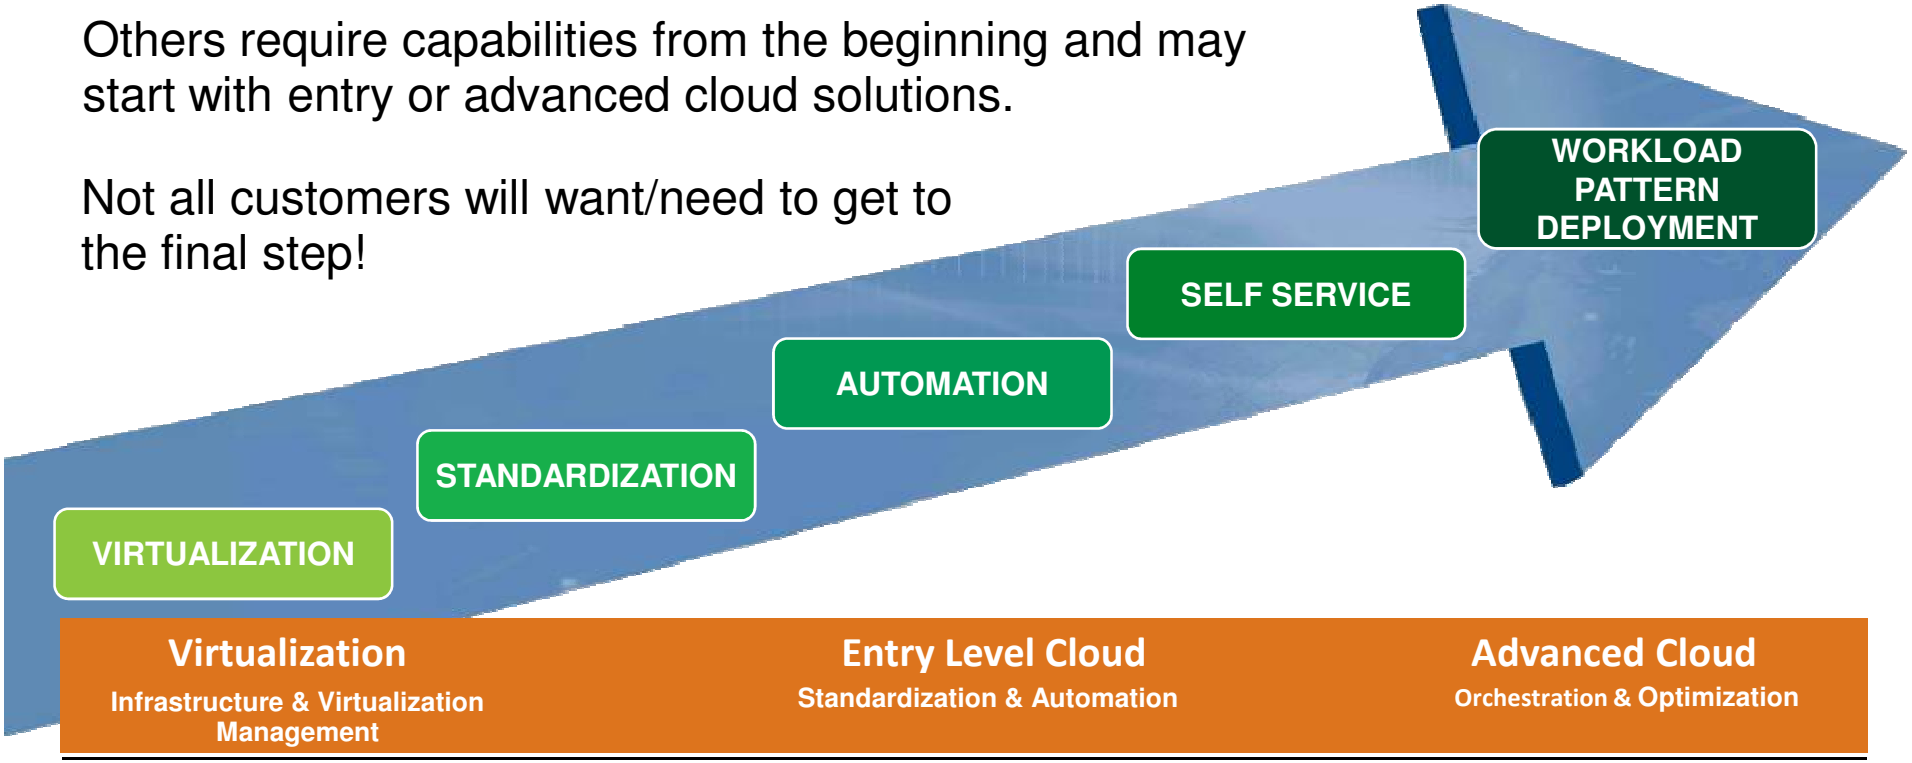

Cloud Computing Journey

The steps in the cloud journey offer different levels of capability for the customer IT environment.

Customers can embark on their cloud journey at any step.

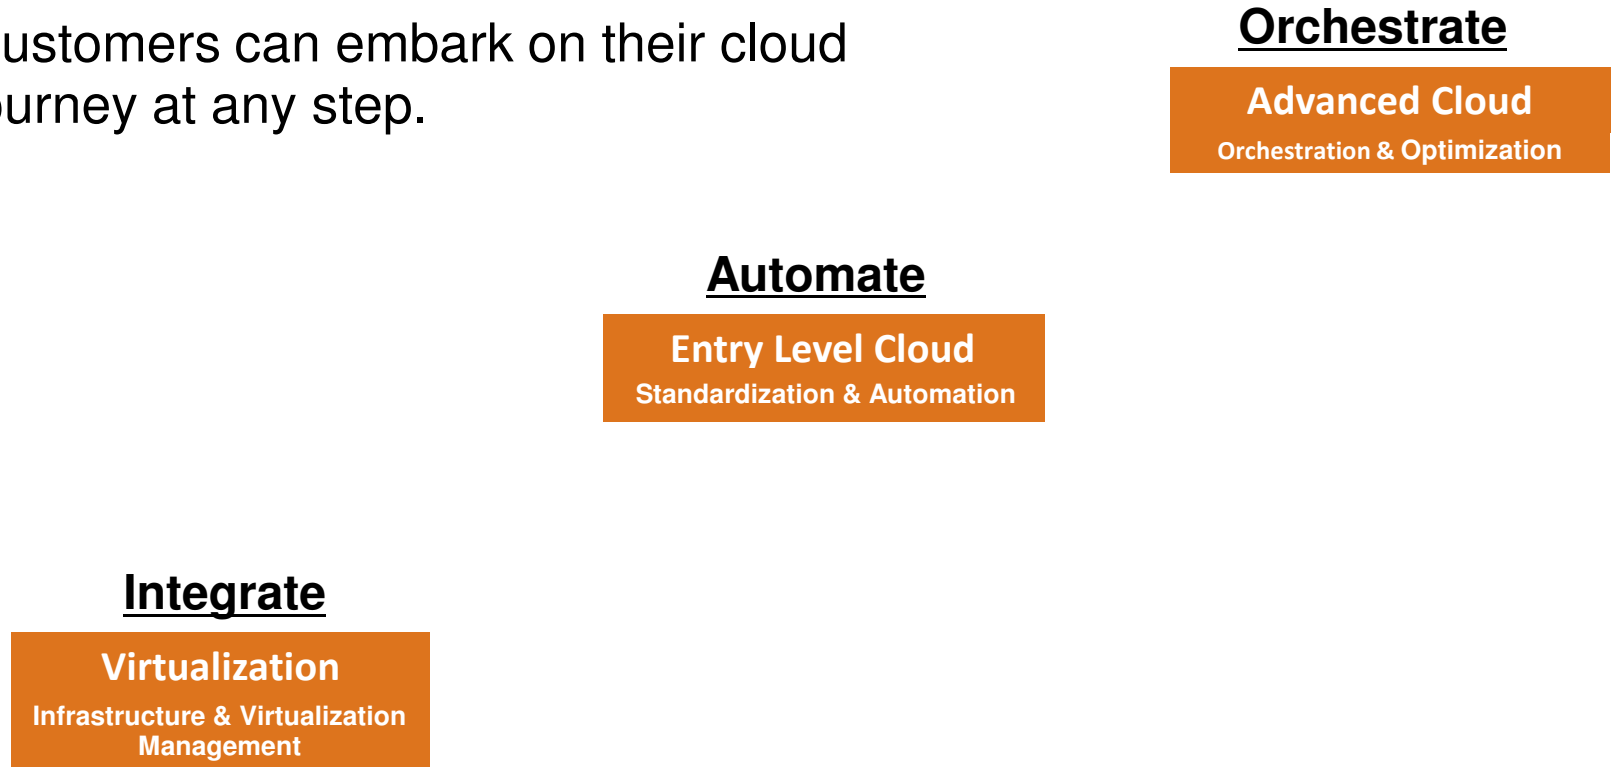

z Systems Cloud Blueprint

Orchestrate

Advanced Cloud
Orchestration & Optimization

Automate

Entry Level Cloud
Standardization & Automation

Integrate

Virtualization
Infrastructure & Virtualization Management

This is where z Systems drives differentiation!

- Infrastructure Scalability: Consolidate more workloads per core; elastic scaling using Capacity On Demand
- Virtualization Management: More virtual servers in a single footprint
- Security: Highest security rating for tenant isolation
- Reliability & Availability: Unparalleled in the industry

z Systems Cloud Blueprint

Orchestrate

Advanced Cloud
Orchestration & Optimization

Automate

Entry Level Cloud
Standardization & Automation

- Customers begin to standardize their environments for faster delivery of services.
- Automation is employed to provision and deprovision virtual guest environments using a shared pool of resources.
- Some customers may choose to allow end-user self service provisioning/deprovisioning.

Integrate

Virtualization
Infrastructure & Virtualization Management

This is where z Systems drives differentiation!

- Infrastructure Scalability: Consolidate more workloads per core; elastic scaling using Capacity On Demand
- Virtualization Management: More virtual servers in a single footprint
- Security: Highest security rating for tenant isolation
- Reliability & Availability: Unparalleled in the industry

z Systems Cloud Blueprint

Orchestrate

Advanced Cloud
Orchestration & Optimization

Finally, some customers will want to evolve and optimize their cloud environment to orchestrate application deployment based on reusable workload patterns in order deliver dynamic cloud services.

Automate

Entry Level Cloud
Standardization & Automation

- Customers begin to standardize their environments for faster delivery of services.
- Automation is employed to provision and deprovision virtual guest environments using a shared pool of resources.
- Some customers may choose to allow end-user self service provisioning/deprovisioning.

Integrate

Virtualization
Infrastructure & Virtualization Management

This is where z Systems drives differentiation!

- Infrastructure Scalability: Consolidate more workloads per core; elastic scaling using Capacity On Demand
- Virtualization Management: More virtual servers in a single footprint
- Security: Highest security rating for tenant isolation
- Reliability & Availability: Unparalleled in the industry

Virtualization and Cloud Portfolio for Linux on System z

Virtualization Infrastructure & Virtualization Management	Entry Level Cloud Standarization & Automation	Advanced Cloud Orchestration & Optimization

Servers: z13, zEC12, zBC12

- Massively scalable
- Characterized by great economics / efficiencies
- Highly secure / available

z/VM 6.3

- Support more virtual servers than any other platform in a single footprint
- Integrated OpenStack support

IBM Wave for z/VM

- A graphical interface tool that simplifies the management and administration of z/VM and Linux environments

Differentiation

Cloud Manager with OpenStack

- A simple, entry level cloud management stack
- Based on OpenStack
- Formerly known as SmartCloud Entry

Standardization

Cloud Orchestrator

- Based on OpenStack
- Builds on functionality of **Cloud Manager with OpenStack** and adds runbook automation and middleware pattern support for workload deployment
- Formerly known as SmartCloud Orchestrator

Service Lifecycle Management

Virtualization and Cloud Portfolio for Linux on System z

Virtualization Infrastructure & Virtualization Management	Entry Level Cloud Standarization & Automation	Advanced Cloud Orchestration & Optimization
	zPro • Provided by Velocity Software • Add-on feature to Velocity's zVPS product that provides performance management • Provides golden image creation, cloning, and operational controls	**Moab** • Provided by Adaptive Computing • Provides a policy based cloud management based on xCAT REST APIs
Differentiation	*Standardization*	*Service Lifecycle Management*

IBM Cloud Manager with OpenStack

- Formerly offered as IBM SmartCloud Entry V3

- Benefits:
  - Full access to OpenStack APIs, backed with IBM support
  - Cloud management solution that is easy to use
  - Self service portal for workload provisioning and virtualized image management
  - Heterogeneous support for IBM PowerVM®, z/VM, IBM PowerKVM and x86, and more.
  - Deploy, resize and capture
  - Linux server backup and restore
  - Manage z/VM from z/VM or manage z/VM from other platforms

- Requires z/VM 6.3 with appropriate service
  - Based on OpenStack

IBM Cloud Manager with OpenStack 4.1

- IBM Cloud Manager with OpenStack V4.1
  - Available: June 13, 2014
  - Based on Icehouse level of OpenStack
  - Manages “to z/VM” (i.e. needs server of z Systems)

- IBM Cloud Manager with OpenStack V4.1
  - Available: September 16, 2014
  - Based on Icehouse level of OpenStack
  - Manage z from z support added via “CMO appliance that ships with z/VM”

IBM Cloud Manager with OpenStack 4.2

- IBM Cloud Manager with OpenStack V4.2
  - Available: December 12, 2014
  - Based on Juno level of OpenStack
  - Manage to z only (appliance not updated)

- IBM Cloud Manager with OpenStack V4.2
  - Announced: February 24, 2015
  - Available: March 13, 2015
  - Based on Juno level of OpenStack
  - Manage from z to anywhere via “updated CMO appliance”

IBM Cloud Manager with OpenStack 4.3

- IBM Cloud Manager with OpenStack V4.3
  - Announced: May 11, 2015
  - Availability: June 19, 2015
  - Based on Kilo level of OpenStack
  - Manage to z only (appliance not updated)

IBM Cloud Orchestrator V2.4

- SmartCloud Orchestrator V2.3 renamed to IBM Cloud Orchestrator V2.4
  - Based on OpenStack enablement in z/VM
  - Icehouse level of OpenStack
- Announced October 7, 2014: US Announcement Letter 214-348
- Available October 10, 2014
- Provides:
  - Self Service
  - Rich provisioning
  - Integration to Business Processes
  - Pattern management
  - “Managed To”

z/VM & OpenStack Levels

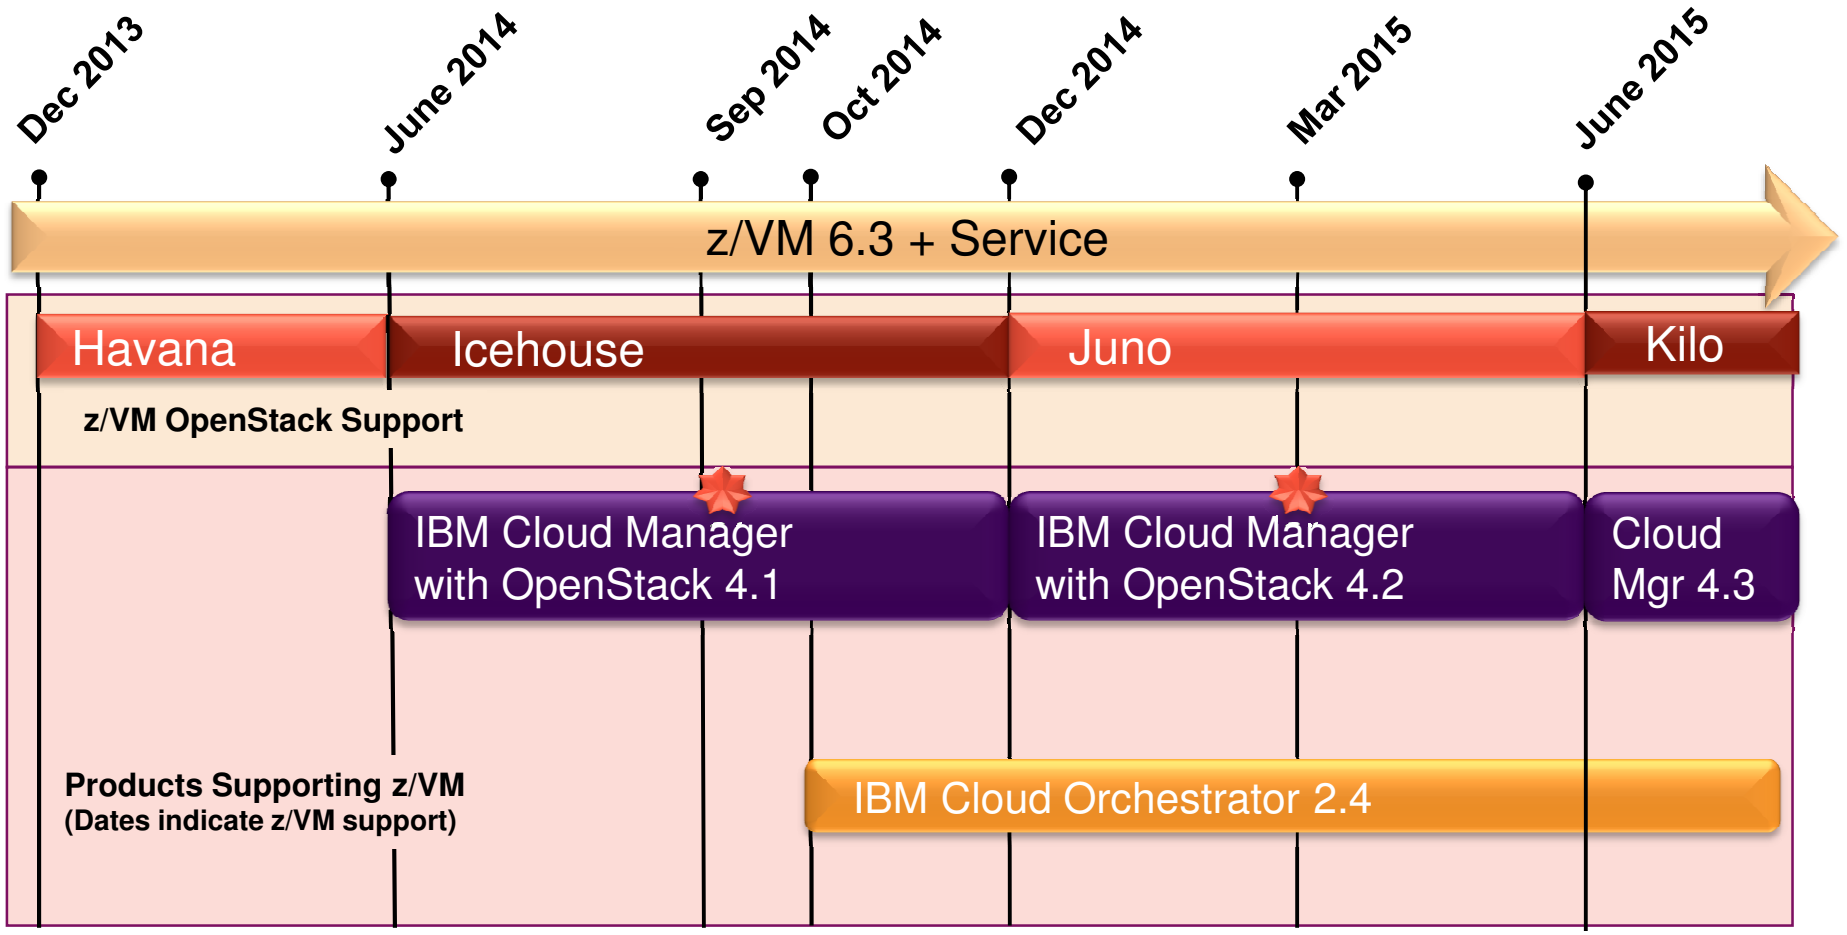

- z/VM Cloud Manager Appliance

http://www.vm.ibm.com/sysman/openstk.html
http://www.vm.ibm.com/sysman/osmntlvl.html

IBM Infrastructure Suite for z/VM and Linux 1.1.0

- Announced and Available
  - Announced September 2, 2014
  - Available September 5, 2014
  - Announcement Letter ENUS214-350

- Includes following products:
  - IBM Tivoli® OMEGAMON® XE on z/VM and Linux V4.3
  - IBM Tivoli Storage Manager, part of IBM Spectrum Protect, Extended Edition V7.1
  - IBM Operations Manager for z/VM V1.5
  - IBM Backup and Restore Manager for z/VM V1.3
  - IBM Wave for z/VM V1.1

Product Offerings (organized by Discipline)

Disciplines of Systems Management

Automation	• Scheduling and running common events and processes
Monitoring	• Monitor and acting on events and conditions of the system
Business Continuity	• Disaster Recovery and Backups/Restores
Data Management	• Tape and storage management
Accounting	• Charge back and billing; Inventory
Security	• Enforcing permissions, tracking, and auditing.
Performance	• Performance Management and Capacity Planning
Provisioning	• Creating and changing the resources and configurations of virtual machines and perhaps guest systems

Automation

• Scheduling and running common events and processes

- **Operations Manager for z/VM**
  - Priced Product
  - Facilitates automated operations
  - Monitor, view, and interact with consoles without logging on to service machines or Linux guests
  - Take actions based on service machine console messages and other system events
  - Schedule events for immediate execution or on a regular schedule
  - Interaction with OMEGAMON XE
- **WAKEUP** Utility
  - Part of z/VM
  - Rudimentary function for basic scheduling
- **IBM Wave for z/VM**
- **ISV Solutions** exist

Monitoring

• Monitor and acting on events and conditions of the system

- **Operations Manager for z/VM**
  - Priced Product
  - Facilitates automated operations
  - Monitor, view, and interact with consoles without logging on to service machines or Linux guests
  - Monitoring and Management of virtual machines, spool files, and error messages
  - Take actions based on service machine console messages and other system events
  - Schedule events for immediate execution or on a regular schedule
  - Interaction with OMEGAMON XE

- **Programmable Operator (PROP)**
  - Part of z/VM
  - Rudimentary function for monitoring

- **ISV solutions** exist

Business Continuity

• Disaster Recovery and Backups/Restores

- **Backup and Restore Manager for z/VM**
  - Priced Product
  - Backup and restore file level data for CMS minidisks and Shared File System
  - Backup and restore images of Linux guests and/or z/VM volumes
    - Use Tivoli Storage Manager for file level backup and restore of Linux data
  - Compression
  - Exit available for encryption

- **GDPS®** Offering
  - Priced Offering

- **Live Guest Relocation**
  - Part of z/VM Single System Image priced feature
  - Address planned outages

Data Management

• Tape and storage management

- **Backup and Restore Manager for z/VM**
  - Priced Product
  - Backup and restore file level data for CMS minidisks and Shared File System
  - Backup and restore images of Linux guests and/or z/VM volumes
    - Use Tivoli Storage Manager for file level backup and restore of Linux data
- **Tape Manager for z/VM**
  - Priced Product
  - Manage tapes: retention, access control, data security erase
  - Manage devices: share with other z/VM and non-z/VM systems
  - Manage mount requests for ATL, VTS, and manual mount devices
    - Supports IBM, Oracle STK libraries, and EMC libraries
    - TS7700 needs firmware update is available as code level 8.21.0.165 (EC: M13120 / PN: 2727271 & 2727272 (DVD1&2.))
- **Archive Manager for z/VM**
  - Priced Product
  - Users and administrators manage disk space more efficiently and effectively
  - Archive infrequently used or large files to tape or other disk
- **Other Options**
  - Rudimentary with DDR or from z/OS
- **ISV solutions** exist

Accounting

• Charge back and billing; Inventory

- CP System Service *ACCOUNT – interface to accounting information (processor, I/O, memory, … resources)
- **RETRIEVE** – VM utility that captures accounting records
- **ACCOUNT** – VM utility to produce reports from accounting records. Very simple.
- **ISV Solutions** exist for processing records.
- Do It Yourself
- **IBM Tivoli Usage and Accounting Manager** (ITUAM)
- Overlap with performance management
  - Some customers are using performance data for charge back

Security

• Enforcing permissions, tracking, and auditing.

- CP System Service *RPI – allows interaction with and ESM (External Security Manager)

- Base and extensions for
  - Virtual machine authentication
  - Virtual machine authorization
  - Auditing of violations

- **RACF for VM** is a priced feature that is pre-installed on the base system.

- **zSecure™ Manager for RACF z/VM**
  - Automate complex, time consuming z/VM security management tasks
  - Quickly identify and prevent problems in RACF
  - Create comprehensive audit trails

- **ISV Products** - External Security Managers (ESMs)

Performance

• Performance Management and Capacity Planning

- **Performance Toolkit for z/VM**
  - Priced feature of z/VM that is pre-installed
  - Basic realtime monitoring and report generation
- **OMEGAMON® XE on z/VM and Linux**
  - Priced Product
  - Performance monitoring of z/VM and Linux guests
  - Part of the OMEGAMON and IBM Tivoli Monitoring infrastructure, including Tivoli Enterprise Portal
  - Uses IBM Performance Toolkit for VM as its data source
  - Optionally uses Linux agent in virtual machines
- **IBM Wave for z/VM**
  - Very small subset based on Performance Toolkit data
- **IBM Tivoli Decision Support (TDS)**
  - Capacity Planning
- **ISV Solutions**
- **VM Resource Manager**
  - Part of z/VM product
  - Like Workload Manager, but less function and less effective

Provisioning

- Creating and changing the resources and configurations of virtual machines and perhaps guest systems

- **xCAT**
  - Integrated in base of z/VM 6.3
- **IBM Cloud Orchestrator**
- **IBM Cloud Manager with OpenStack**
- **IBM Wave for z/VM**
- **ISV Solutions**

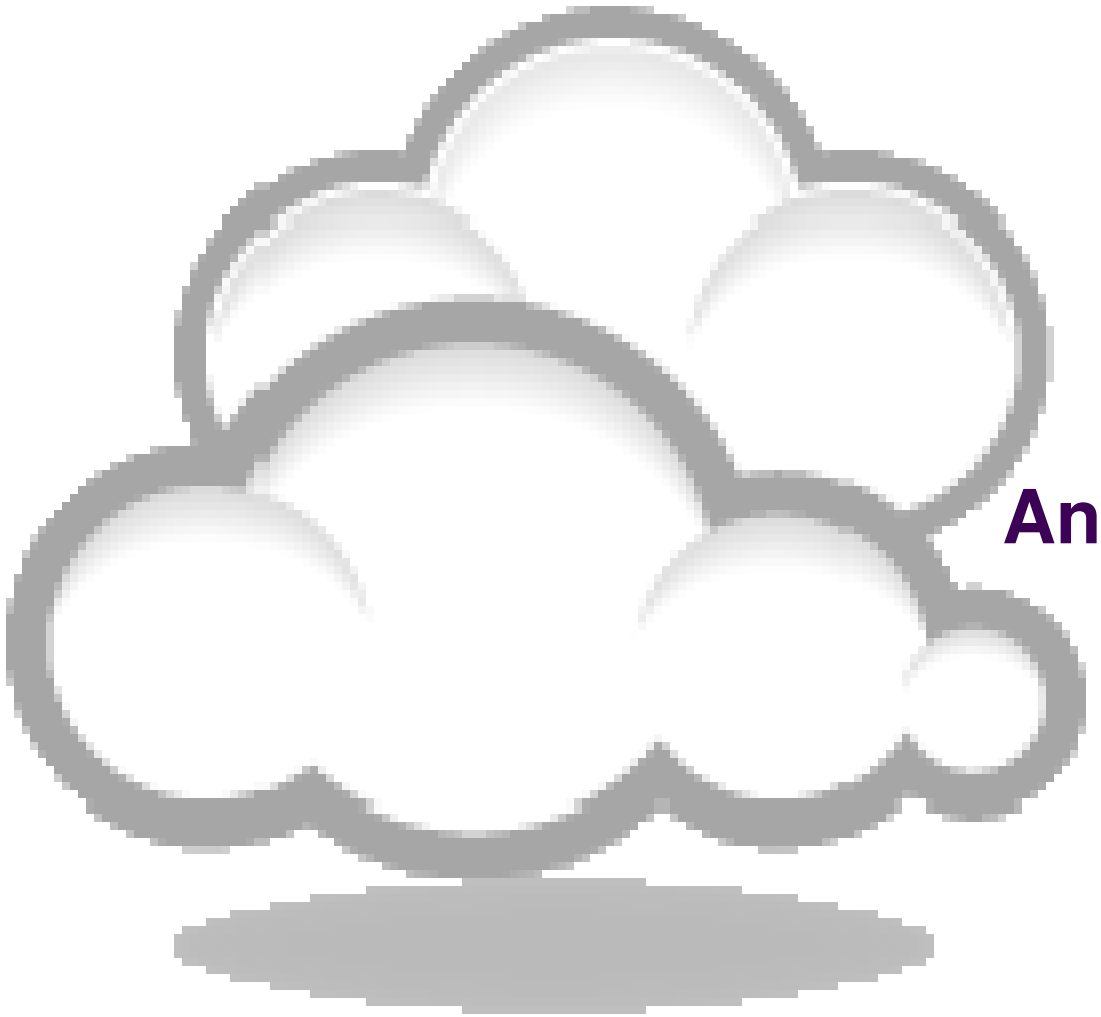

And in conclusion …

Putting together the pieces

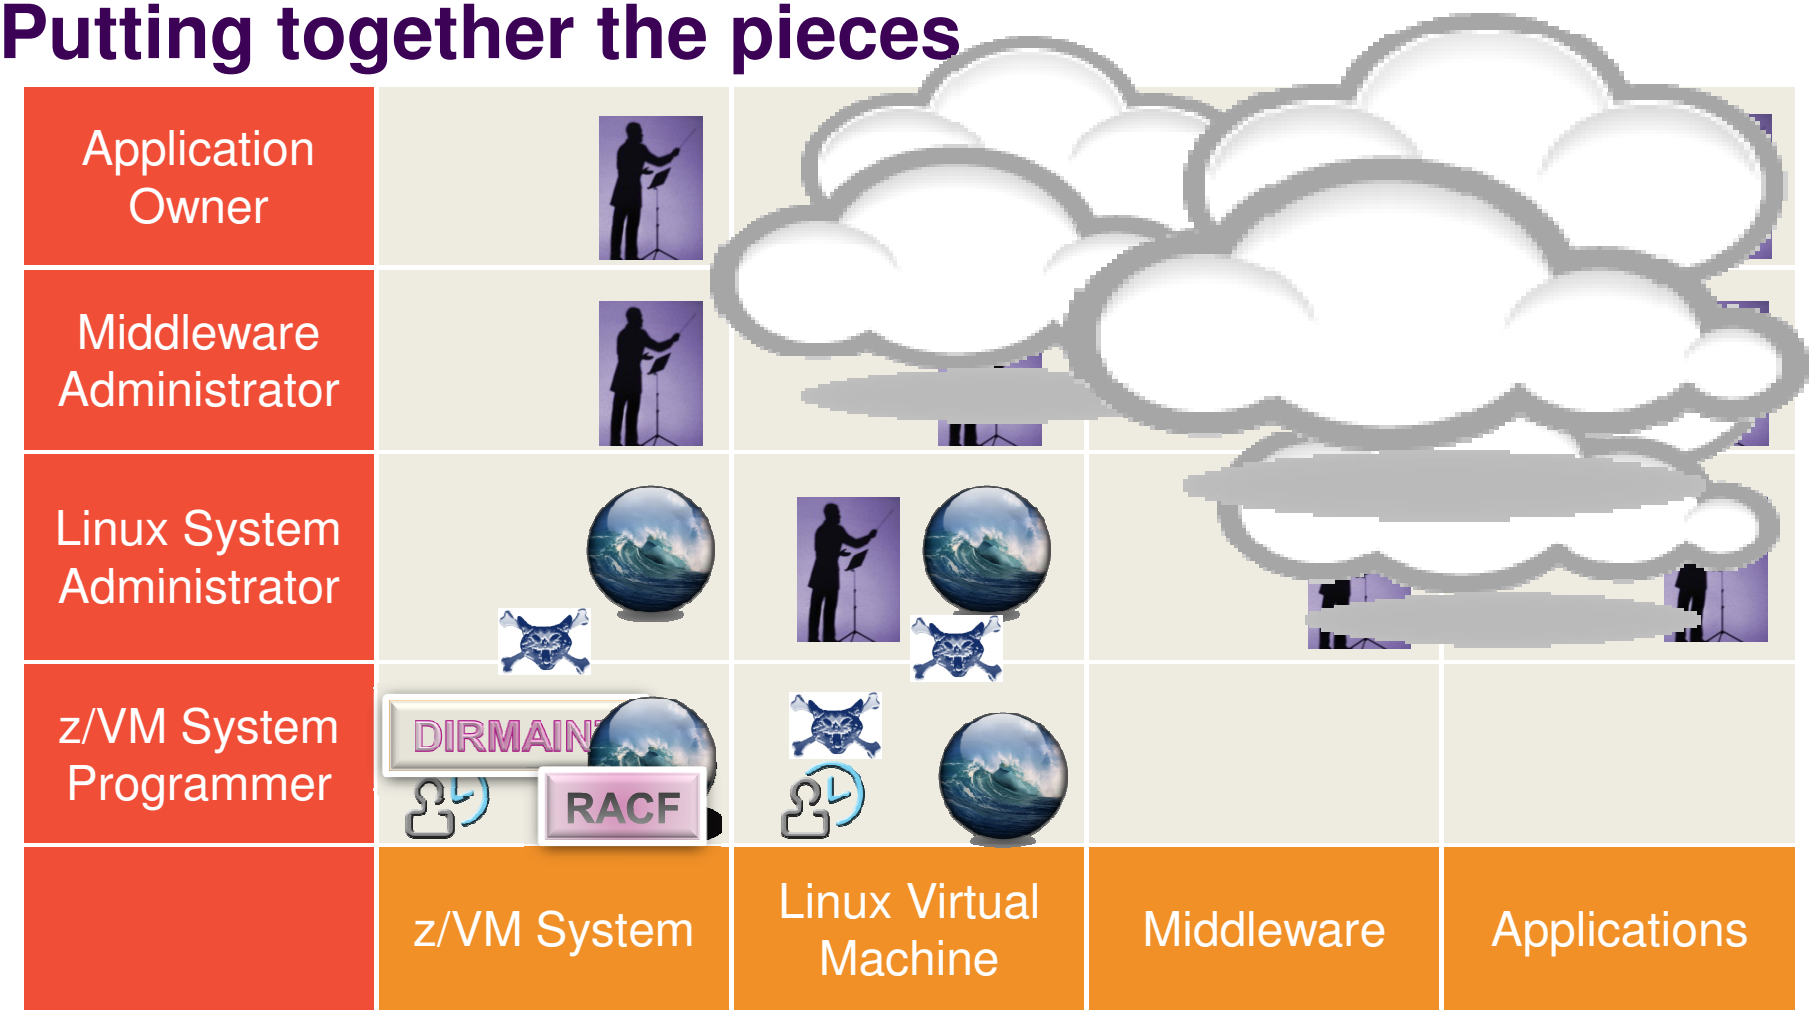

Questions to Ask when Selecting Systems Management Software

- What Systems Management Solutions are already in place in my enterprise?
- What do I want to manage?
- Who will be using the management solution? What are their skills?
- Was it designed and created by people who have actually managed systems?
- Which Systems Management Discipline does it address?
  - Which ones are missing and how do I fill those gaps?
- Does the solution work well with other products?
  - Any side effects of manual management?
- What key features or capabilities of z/VM are critical to our success? Does the solution support those key features?

Summary

- Systems Management is multi-facetted
  - What Systems and what Management?
  - Self Integrate vs. Product Solutions
- Must address all the disciplines to be enterprise-ready
- OpenStack Enablement is Strategic (z/VM 6.3)
  - IBM Cloud Manager with OpenStack
  - IBM Cloud Orchestrator
- Expect…
  - IBM to continue to enhance z/VM to support the ecosystem
  - changes outside of the standard cadence of z/VM releases
  - the need for multiple products to address various needs
  - new solutions from IBM and ISVs for challenges that arise
- http://www.vm.ibm.com/sysman/

References

- z/VM Systems Management main page
  - http://www.vm.ibm.com/sysman/
- z/VM OpenStack Maintenance
  - http://www.vm.ibm.com/sysman/osmntlvl.html

Backup Material
Additional OpenStack Information

OpenStack Strategy

http://openstack.org/

OpenStack is a global collaboration of developers and cloud computing technologists that seek to produce a ***ubiquitous Infrastructure as a Service (IaaS) open source cloud computing platform*** for public and private clouds. OpenStack was founded by Rackspace Hosting and NASA jointly in July 2010. 160 companies and close to 3,000 developers.

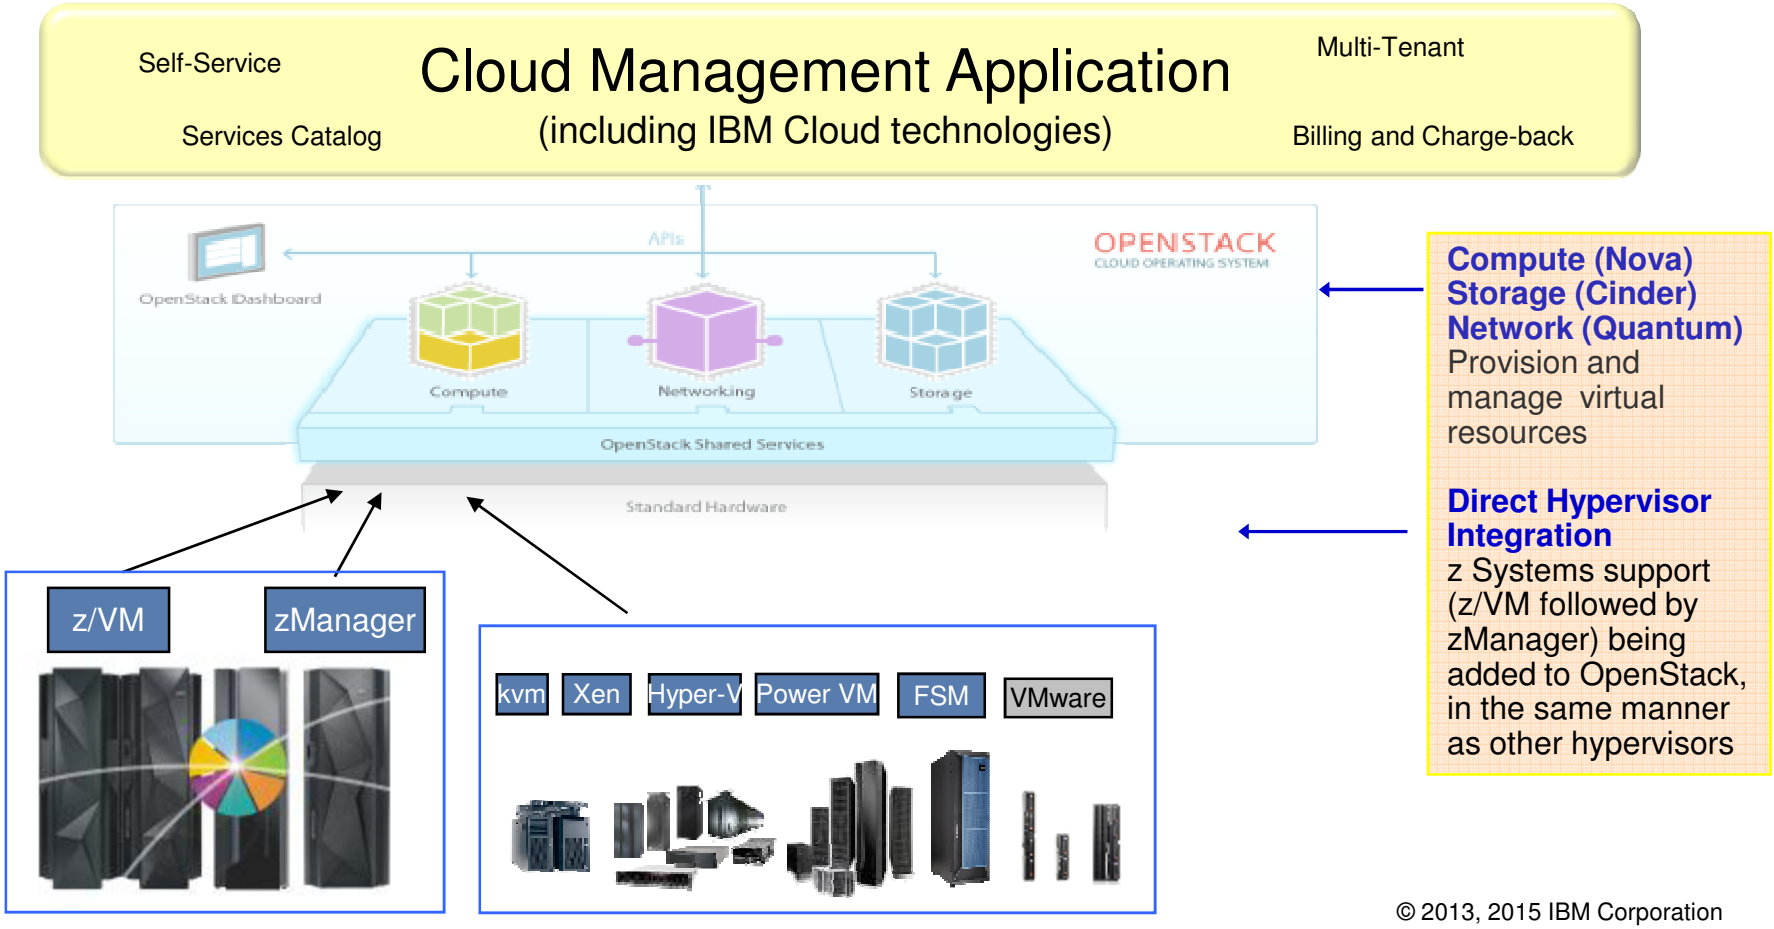

Compute (Nova)
Storage (Cinder)
Network (Quantum)
Provision and manage virtual resources

Direct Hypervisor Integration
z Systems support (z/VM followed by zManager) being added to OpenStack, in the same manner as other hypervisors

OpenStack Project Commitment – Top Ten by Commits

Essex	
Company	**Number**
Rackspace	1848
Red Hat	532
HP	398
Nebula	314
Canonical	97
Nicira	93
Sina	74
Citrix	68
Delta	60
eNovance	57

Folsom	
Company	**Number**
Red Hat	1660
Rackspace	1467
HP	1173
Nebula	414
IBM	210
Sina	204
Cloudscaling	172
Nimbis Services	156
VMware	147
Cononical	128

Grizzly	
Company	**Number**
Red Hat	1858
HP	1456
Rackspace	1160
IBM	972
VMware	362
Nebula	295
eNovance	279
Mirantis	156
Intel	147
OpenStack Foundation	128

Havana	
Company	**Number**
Red Hat	1858
HP	1456
Rackspace	1160
IBM	972
Mirantis	362
OpenStack Foundation	295
SUSE	279
eNovance	156
VMware	147
NEC	128

Top Ten OpenStack Project Contributors

Folsom September 2012		Grizzly April 2013		Havana September 2013		Icehouse April 2014	
Company	**Commits**	**Company**	**Commits**	**Company**	**Commits**	**Company**	**Commits**
Red Hat	1660	Red Hat	1858	Red Hat	1858	Red Hat	2975
Rackspace	1467	HP	1456	HP	1456	IBM	1782
HP	1173	Rackspace	1160	Rackspace	1160	HP	1581
Nebula	414	IBM	972	IBM	972	Rackspace	1351
IBM	210	VMware	362	Mirantis	362	Mirantis	1151
Sina	204	Nebula	295	OpenStack Foundation	295	SUSE	902
Cloudscaling	172	eNovance	279	SUSE	279	OpenStack Foundation	711
Nimbis Services	156	Mirantis	156	eNovance	156	eNovance	527
VMware	147	Intel	147	VMware	147	Intel	427
Canonical	128	OpenStack Foundation	128	NEC	128	NEC	417

Top Ten OpenStack Project Contributors

Grizzly April 2013		Havana September 2013		Icehouse April 2014		Juno September 2014	
Company	**Commits**	**Company**	**Commits**	**Company**	**Commits**	**Company**	**Commits**
Red Hat	1858	Red Hat	1858	Red Hat	3475	HP	3223
HP	1456	HP	1456	HP	2174	Red Hat	2859
Rackspace	1160	Rackspace	1160	IBM	1801	IBM	1407
IBM	972	IBM	972	Rackspace	1730	Mirantis	1371
VMware	362	Mirantis	362	Mirantis	1478	Rackspace	1280
Nebula	295	OpenStack Foundation	295	SUSE	985	*independent	898
eNovance	279	SUSE	279	OpenStack Foundation	716	SUSE	762
Mirantis	156	eNovance	156	eNovance	704	eNovance	414
Intel	147	VMware	147	*independent	686	B1 Systems	410
OpenStack Foundation	128	NEC	128	VMware	554	VMware	405

Deployments consist of projects interfacing over public APIs, with each project composed of multiple services interfacing via private APIs over RPC

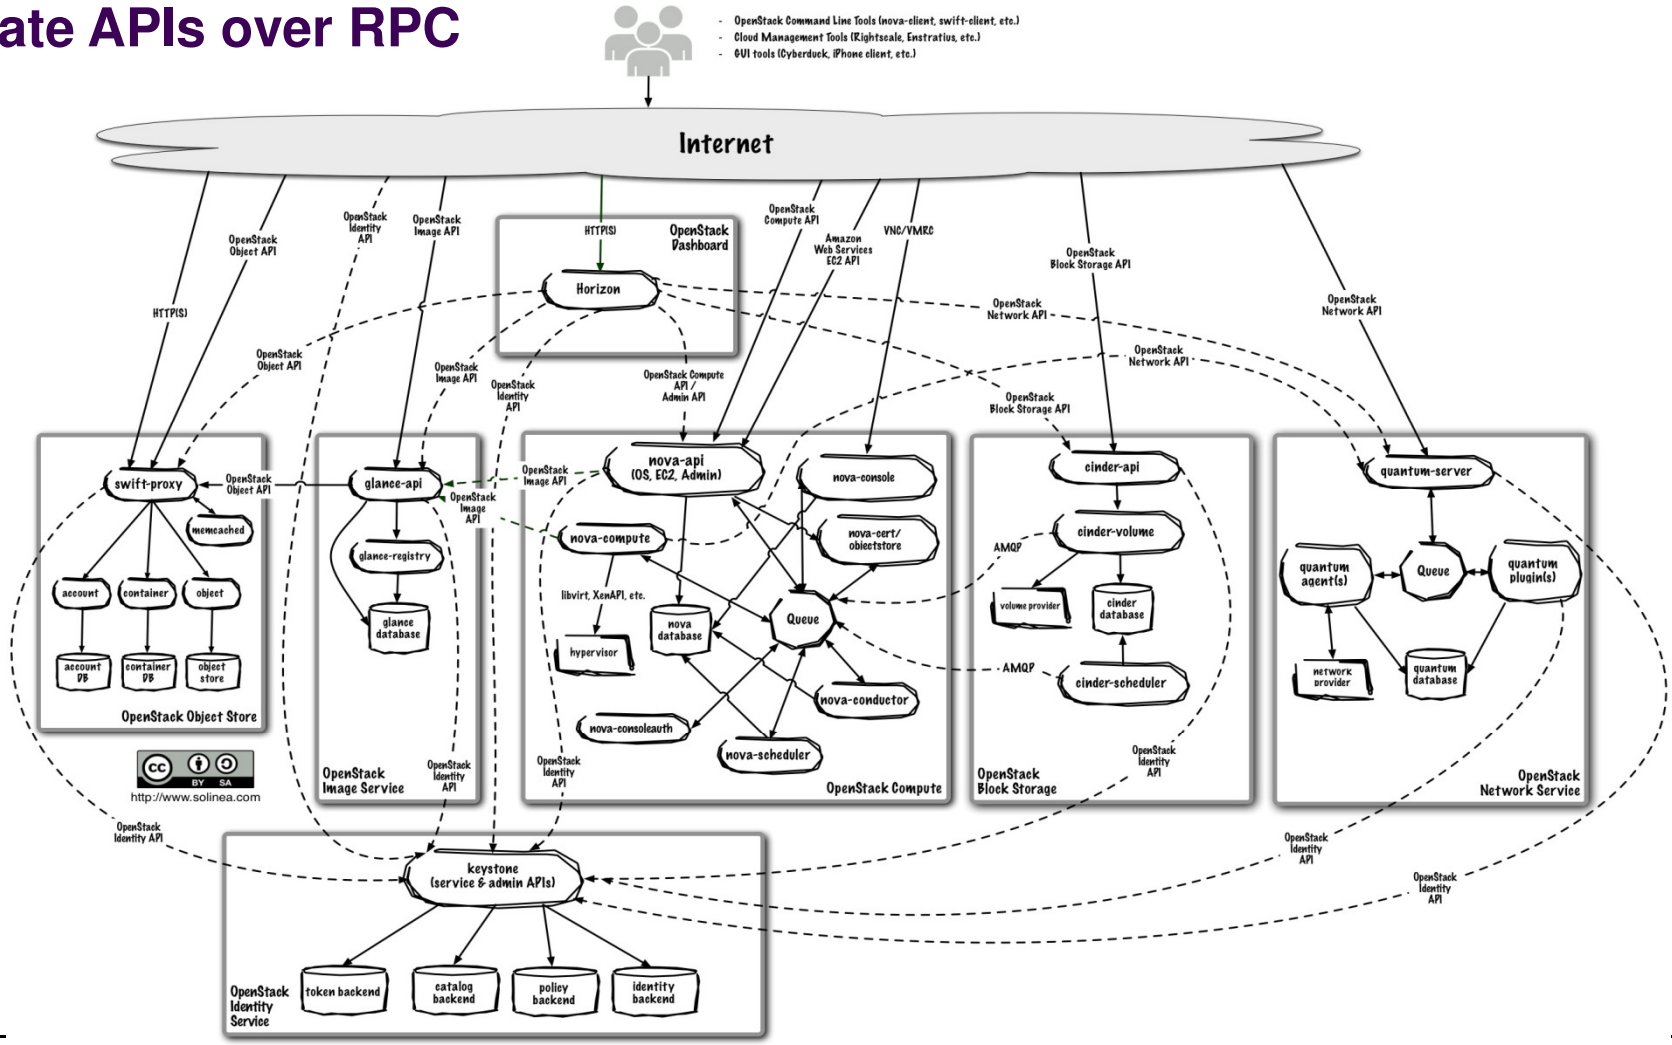

Compute (Nova) is a horizontally scalable offering on-demand compute resources by provisioning and managing

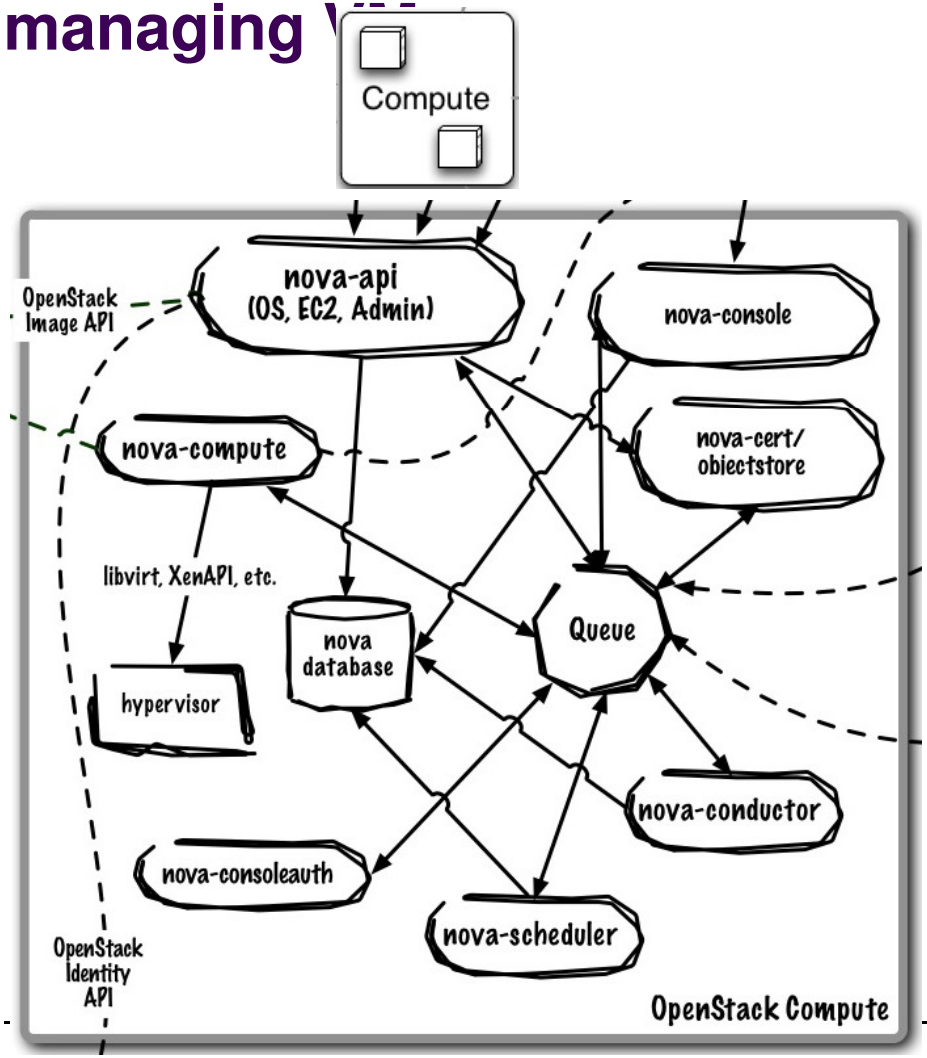

Core Use Case:

- Provision and manage virtualized compute resources (CPU, memory, disk, network)

Key Capabilities:

- REST-based APIs with rate limiting and authentication
- Manage Local Area Networks (LAN)
- Live migration of guests
- VM management (Instance)
  - Run, reboot, suspend, resize, terminate instances
- Floating IP addresses
- Security Groups
- RBAC with Projects & Quotas
- Manage to KVM, Xen (XenServer, Xen Cloud Platform), LXC, VMware vSphere 4.1+, Hyper-V, Bare Metal, PowerVM (limited)

Image Source: http://www.solinea.com/2013/04/17/openstack-summit-intro-to-openstack-architecture-grizzly-edition/

Database and Queue are central to the Nova control plane

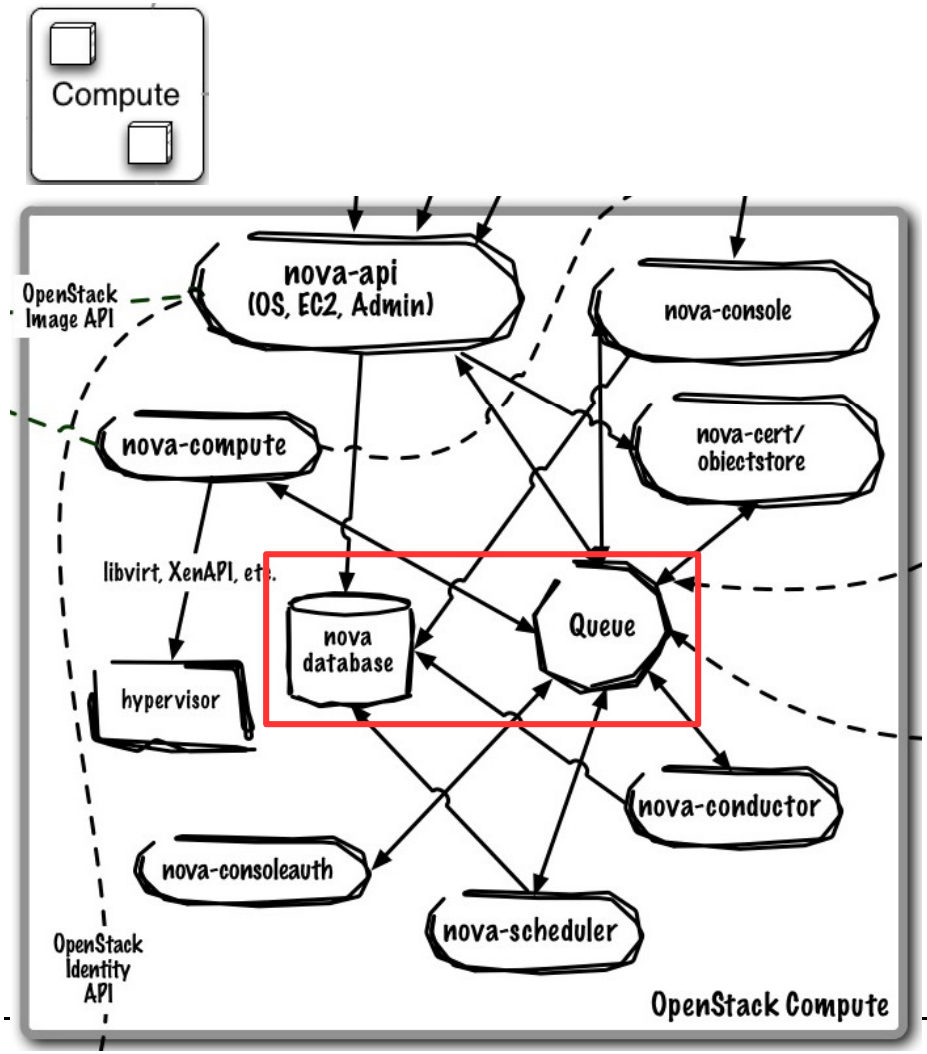

Core Use Case:

- Queue provides RPC messaging between services
- Database provides data persistence

Runs As: Controller Service

Deployment Considerations:

- Use DB and Queue clustering/HA methods
- ZeroMQ implementation available to decentralize queue

Key Capabilities:

- Community uses RabbitMQ as default queue, MySQL DB (IBM uses Apache Qpid and DB2)
- Single “cell” (1 Queue, 1 Database) typically scales from 500 – 1000 physical machines
  - Cells can be rolled up to support larger deployments
- Communications route through queue
  - API requests are validated and placed on queue
  - Workers listen to queues based on role or role + hostname
  - Responses are dispatched back through queue

nova-compute manages individual hypervisors and compute nodes

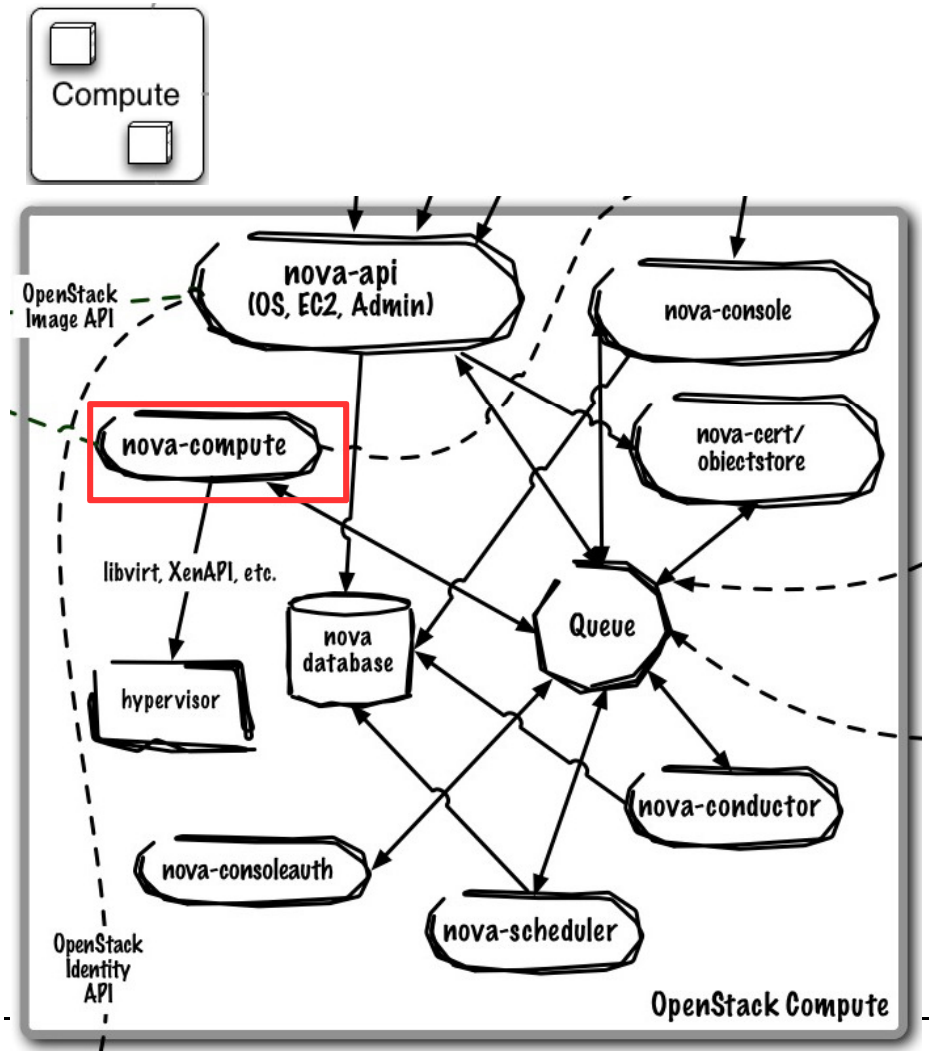

Core Use Case:

- Manage all interactions with single hypervisor control point

Runs As: Distributed Service

Deployment Considerations:

- Many nova-compute instances exist in the environment to ensure compute provisioning is always available
- Single nova-compute is not HA, manage single hypervisor to minimize failure domain
- No direct database acces is required

Key Capabilities:

- Create and manage virtual machines on hypervisor
- Attach networks and volumes to physical host (iSCSI, FC), expose to guest virtual machines
- Implementation point for security groups defining firewall rules for guest network traffic
- Uses plug-in model to manage to different hypervisors

nova-scheduler allocates virtual resources to compute nodes

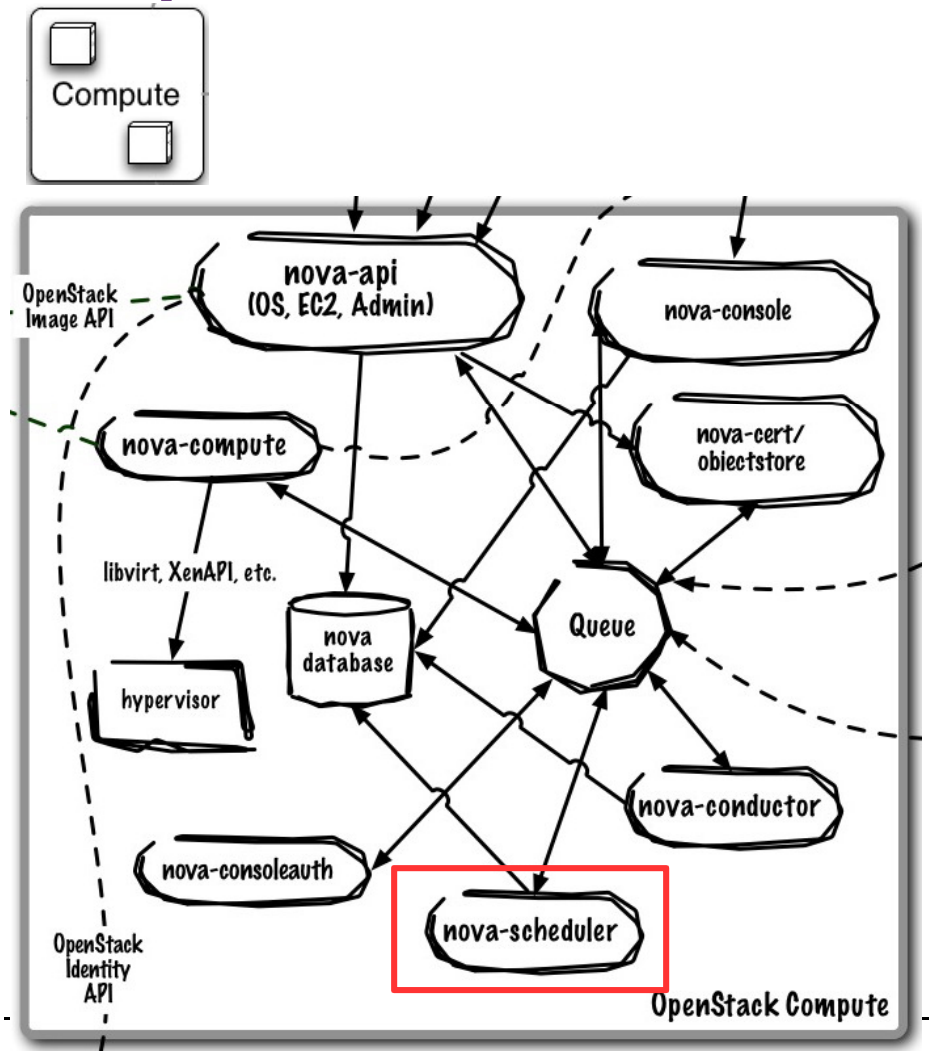

Core Use Case:

- Selects compute node to run virtual machine on

Runs As: Controller Service

Deployment Considerations:

- Default scheduler is horizontally scalable
- For other schedulers (e.g. Platform EGO), follow their specific best practice

Key Capabilities:

- Default scheduled is allocation-based using a series of filters to reduce set of applicable hosts and uses costing functions to provide weight
- Platform EGO adds utilization-based scheduling to default allocation based

nova-api supports multiple API implementations and is t try point into the cloud

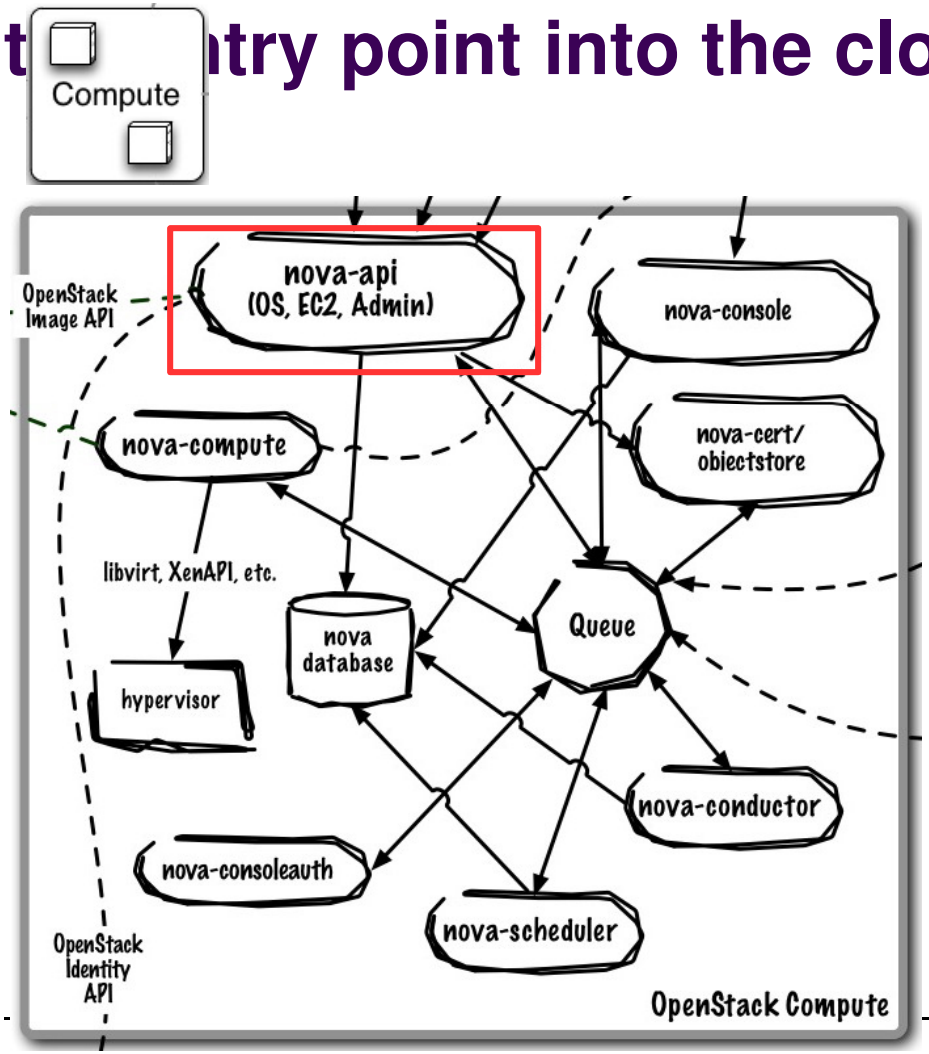

Core Use Case:

- Accept, validate, authenticate, and distribute incoming REST API requests

Runs As: Controller Service

Deployment Considerations:

- Horizontally scalable, start many instances
- Front with load-balancer to present as single endpoint

Key Capabilities:

- APIs supported
  - OpenStack Compute API
  - EC2 API (subset)
- Robust extensions mechanism to add new capabilities

nova-conductor manages database interactions on behalf of compute nodes

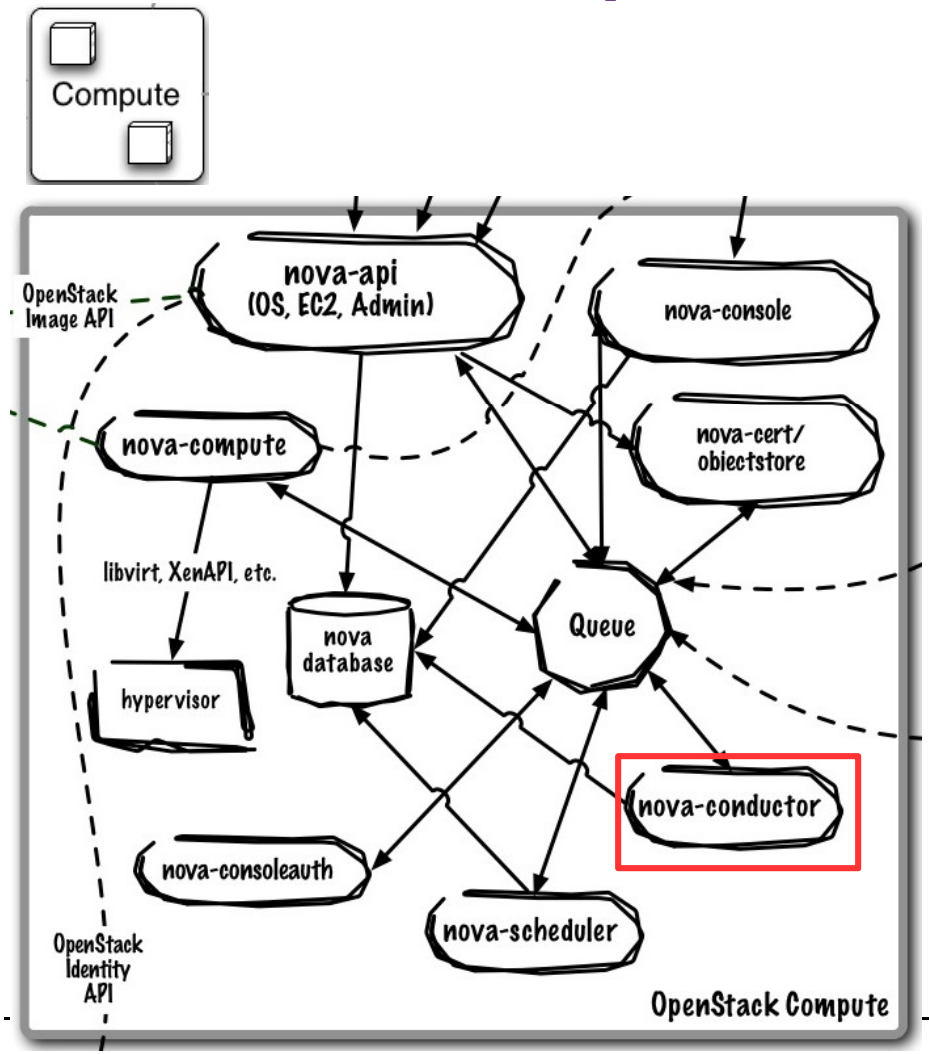

Core Use Case:

- Handles all database requests for nova-compute service

Runs As: Controller Service

Deployment Considerations:

- Horizontally scalable, start many instances

Key Capabilities:

- Talks directly to database on behalf of compute nodes

Network (Quantum) is a pluggable, scalable and API-driven system for managing networks and IP a[illegible]s

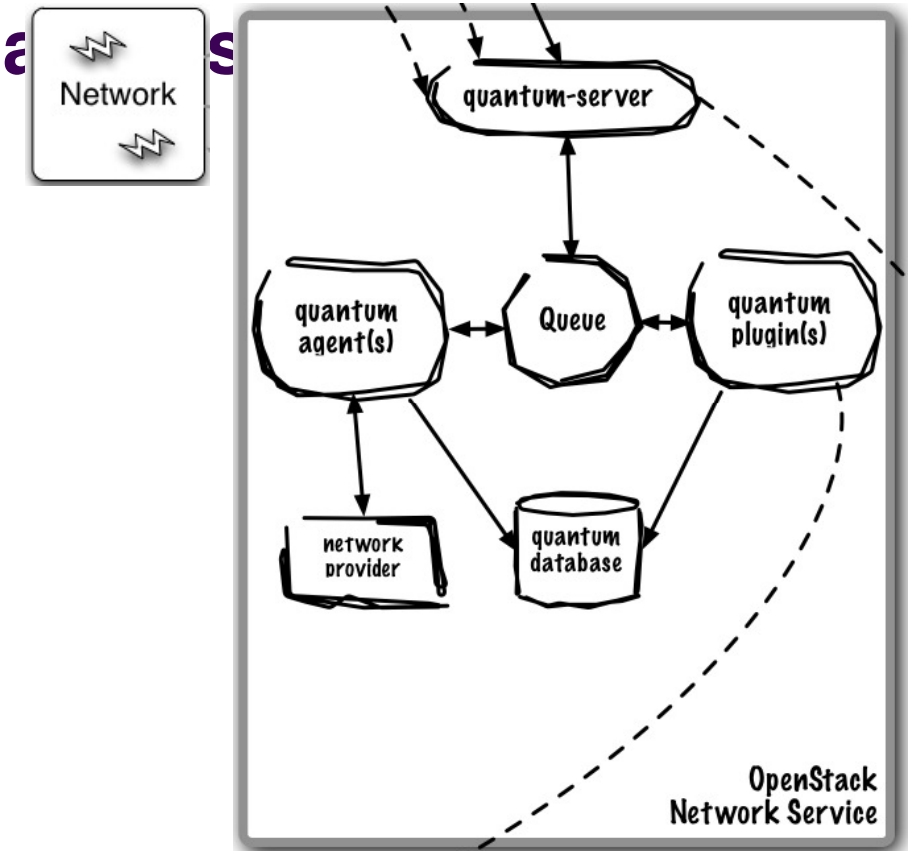

Core Use Cases:

- Provision and manage virtualized network resources (networks, ports, attachments)

Key Capabilities:

- Flexible networking models to suit the needs of different applications or user groups
- Create/delete tenant-specific L2 networks
- Attach / Detach host to network
- L3 support (dedicated static and DHCP, Floating IPs, DHCP, Routing)
- L4-7 Support (Load Balancers)
- Extension framework enabling deploy and management of additional network services: intrusion detection systems (IDS), load balancing, firewalls and virtual private networks (VPN)
- Support for
  - OpenFlow (Big Switch, Floodlight, NEC controllers)
  - Numerous SDN and network virtualization providers (e.g Niciria, Midokura, Plum Grid, Brocade, Mellanox)
  - OpenVswitch
  - Cisco Nexus

Image Source: http://www.solinea.com/2013/04/17/openstack-summit-intro-to-openstack-architecture-grizzly-edition/

Database and Queue are central to the Quantum control plane

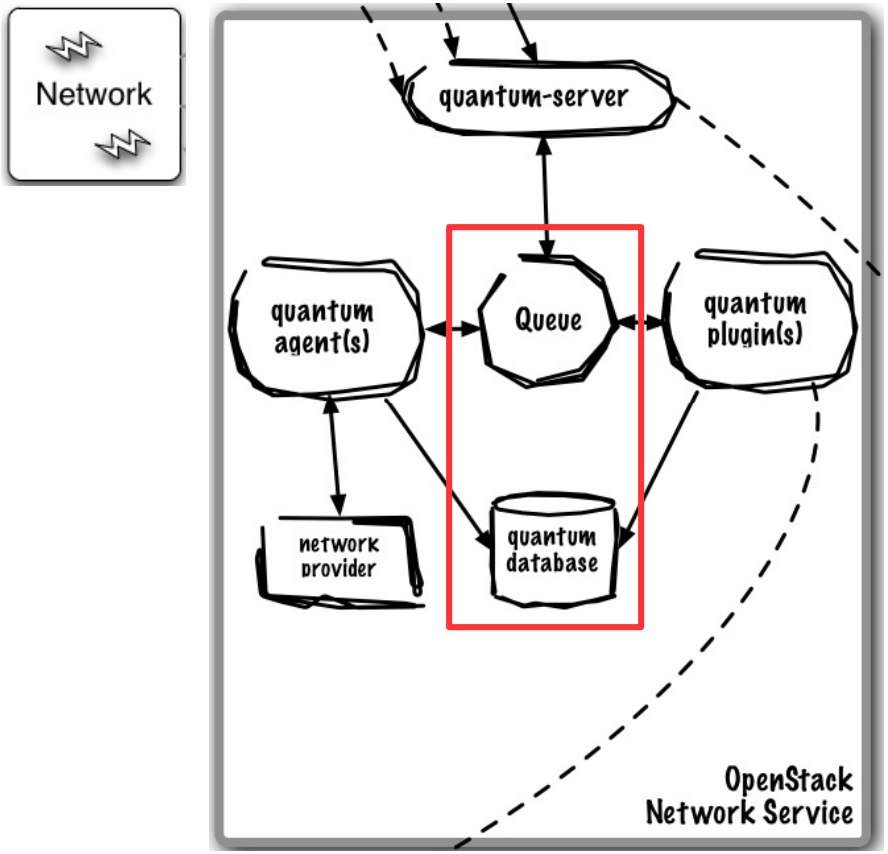

Core Use Case:

- Queue provides RPC messaging between services
- Database provides data persistence

Runs As: Controller Service

Deployment Considerations:

- Use DB and Queue clustering/HA methods
- ZeroMQ implementation available to decentralize queue
- Can use same Queue as Nova

Key Capabilities:

- Community uses RabbitMQ as default queue, MySQL DB (IBM uses Apache Qpid and DB2)

quantum-server implements the OpenStack Network API

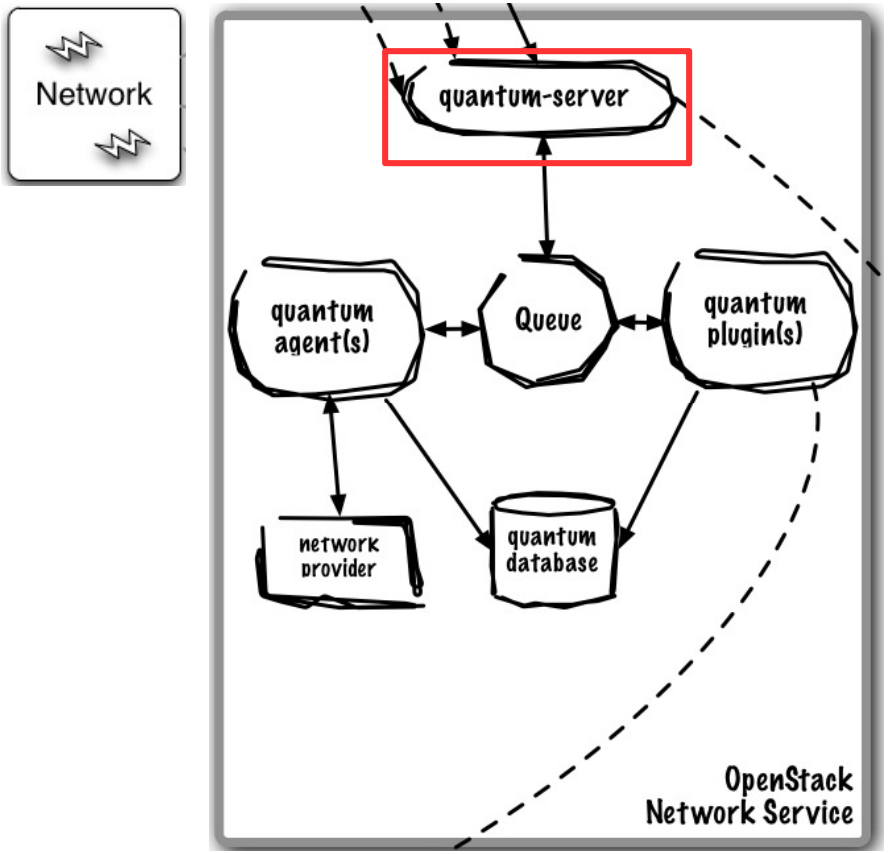

Core Use Case:

- Accept, validate, authenticate, and distribute incoming REST API requests

Runs As: Controller Service

Deployment Considerations:

- Use active/passive or active/active for HA using Linux HA methods (e.g. corosync)

Key Capabilities:

- Requires access to a database for persistent storage
- Passes user requests to the configured OpenStack Networking plug-in for additional processing
- Relies on the OpenStack Identity Project (Keystone) for authentication and authorization of all API request.

Quantum uses an agent model to add additional functionality to a deployment

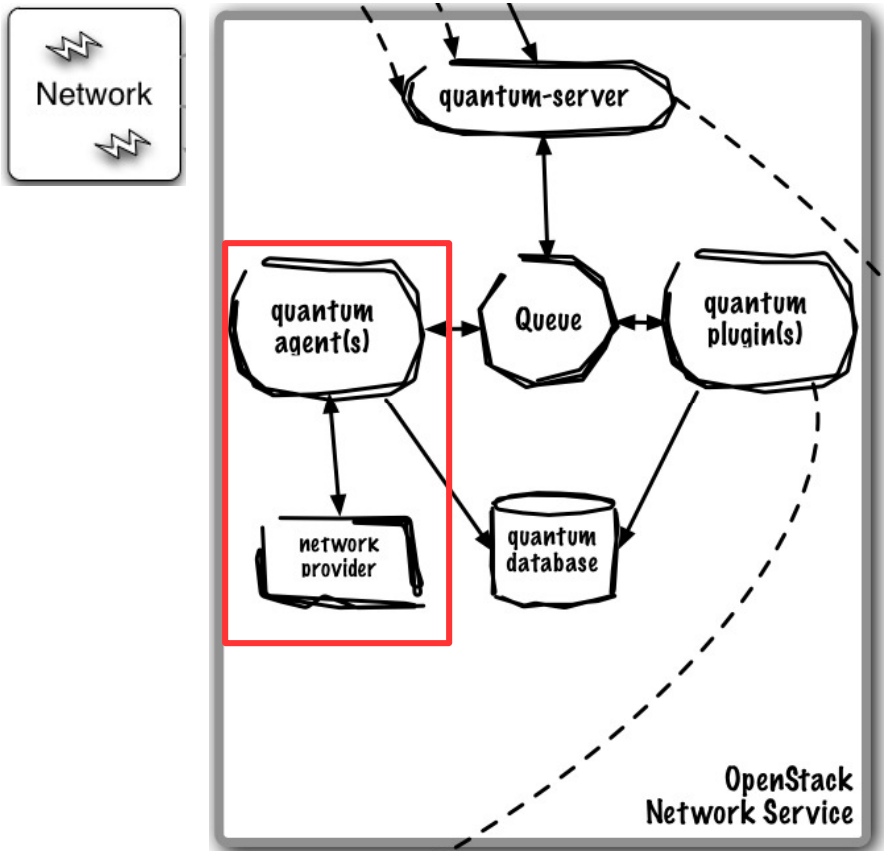

Core Use Case:

- plugin-agent: runs alongside nova-compute to manage physical host network connectivity
- dhcp-agent: provides DHCP to tenant networks
- l3-agent: provides L3/NAT forwarding for external network access

Runs As: Distributed Service (plugin-agent) or Controller Service (dhcp-agent, l3-agent)

HA:

- plugin-agent: same as nova-compute, single instance is not HA, minimize failure domain
- dhcp-agent, l3-agent: running many ensure ensures availability to provision new, can use active/passive or active/active for HA of provisoined node.

Key Capabilities:

- plugin-agent: runs alongside nova-compute to manage physical host network connectivity
- dhcp-agent: provides DHCP to tenant networks
- l3-agent: provides L3/NAT forwarding for external network access

Quantum plugins are vendor or technology-specific plugins that map virtual network topology onto infrastructure

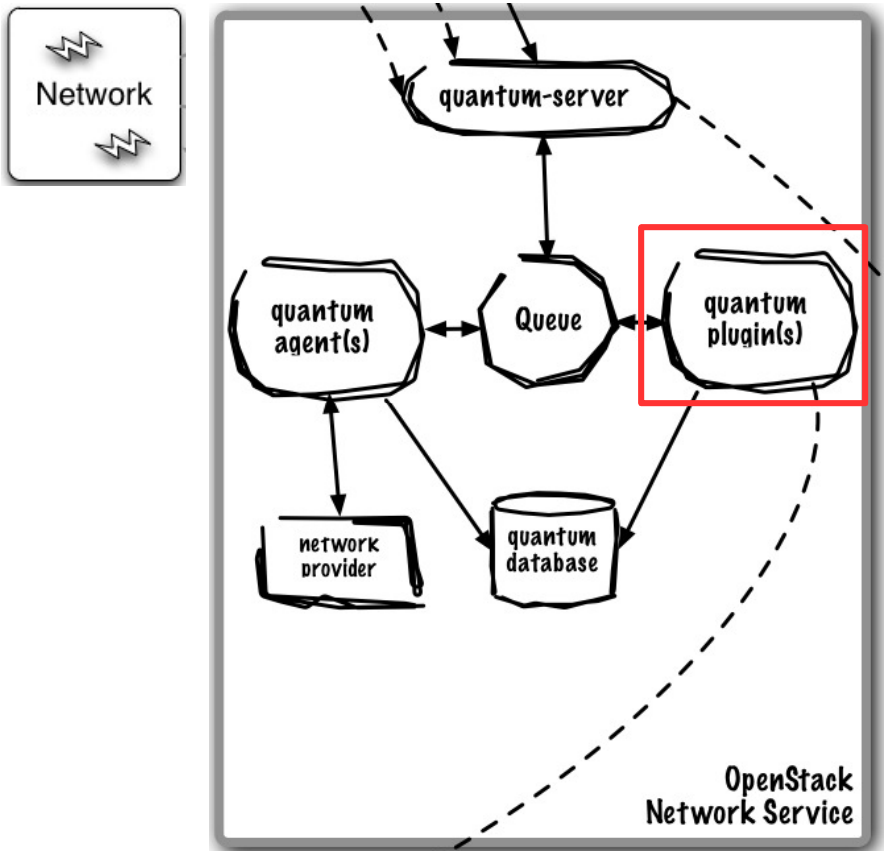

Core Use Case:

- Map virtual network topology onto infrastructure

Runs As: Controller Service

HA:

- Dependent on implementation

Key Capabilities:

- Uses plug-in model to support vendor-specific or technology-specific implementation that translates virtual networks to physical network

Storage (Cinder) exposes block devices to be connected to compute instances for expanded storage, better performance and enterprise storage platform integration

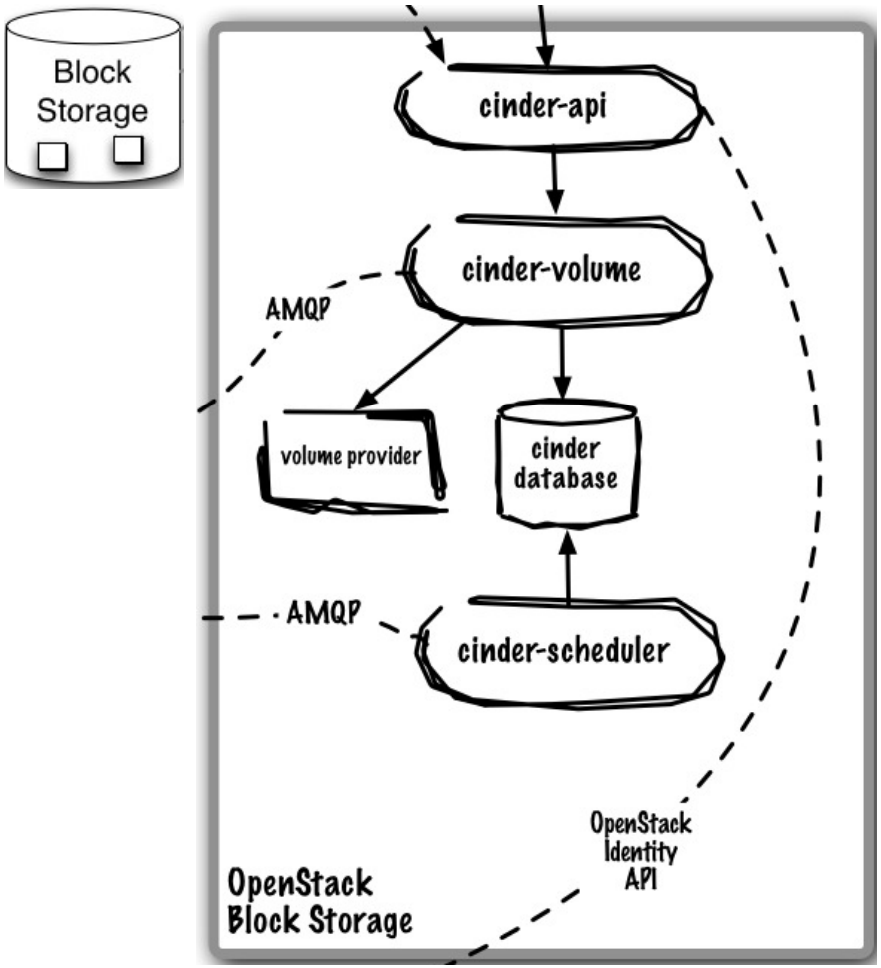

Core Use Cases:

- Provision and manage lifecycle of volumes and their exposure for attachment

Key Capabilities:

- Persistent block level storage devices for use with OpenStack compute instance
- Manage the creation, attaching and detaching of the block devices to servers
- Support for booting virtual machines from Cinder-backed storage
- Snapshot and restore functionality
- Supports following
  - LVM-backed volumes (iSCSI)
  - XIV (iSCSI)
  - SVC (iSCSI and Fiber Channel)
  - NetApp (iSCSI and NFS)
  - EMC (iSCSI)
  - HP/Lefthand (iSCSI)
  - RADOS block devices (e.g. Ceph distributed file system)
    (full list at Cinder Support Matrix)

Image Source: http://www.solinea.com/2013/04/17/openstack-summit-intro-to-openstack-architecture-grizzly-edition/

Database and the Queue are the core of Cinder's control plane

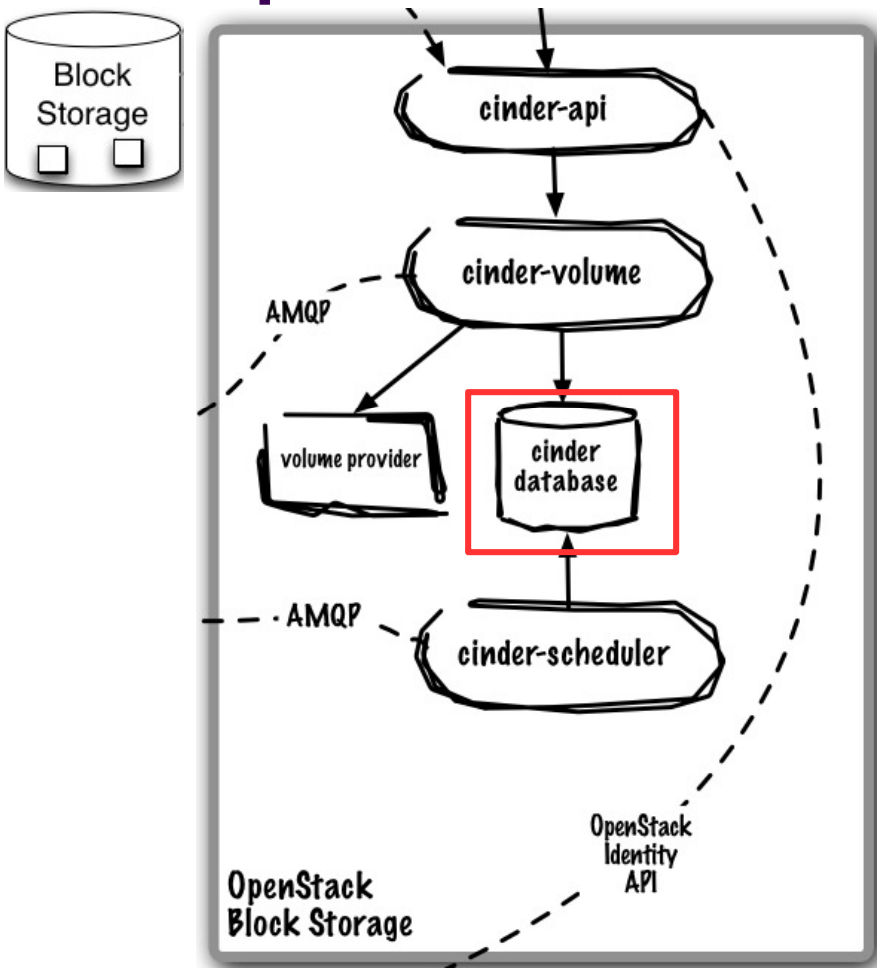

Core Use Case:

- Queue provides RPC messaging between services
- Database provides data persistence

Runs As: Controller Service

Deployment Considerations:

- Use DB and Queue clustering/HA methods
- ZeroMQ implementation available to decentralize queue
- Can use same queue/database as Nova

Key Capabilities:

- Community uses RabbitMQ as default queue, MySQL DB (IBM uses Apache Qpid and DB2)

cinder-api is the entry point to OpenStack Volume Service

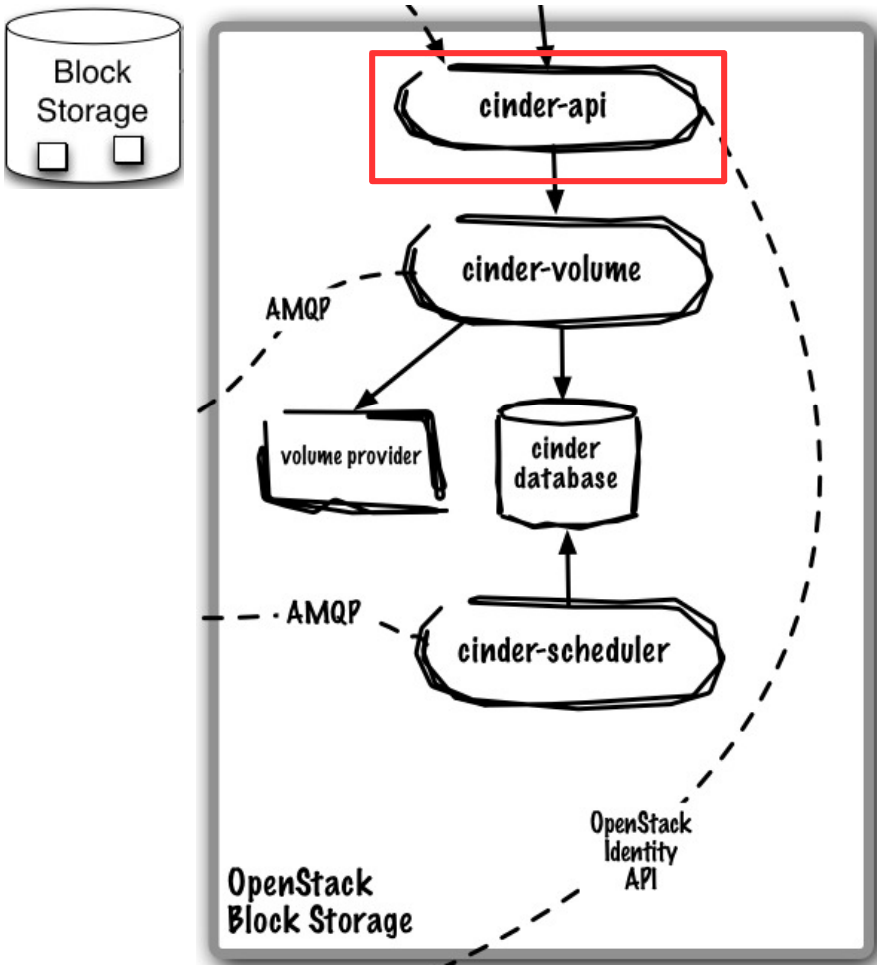

Core Use Case:

- Accept, validate, authenticate, and distribute incoming REST API requests

Runs As: Controller Service

Deployment Considerations:

- Horizontally scalable, start many instances
- Front with load-balancer to present as single endpoint

Key Capabilities:

- APIs supported
  - OpenStack Volume API
- Robust extensions mechanism to add new capabilities

cinder-volume manages individual block-based volume providers

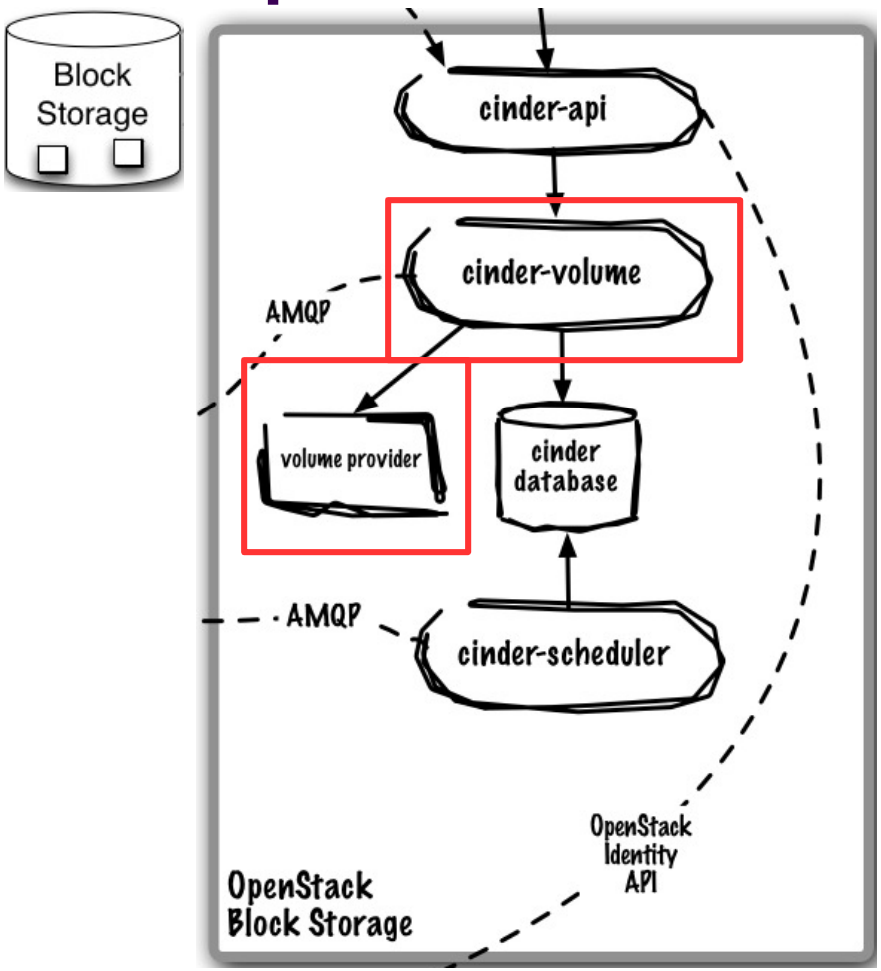

Core Use Case:

- Manages interactions with single block volume provider

Runs As: Distributed Service

Deployment Considerations:

- Many cinder-volume instances exist in the environment to ensure volume provisioning is always available
- Single cinder-volume is not HA, manage single provider to minimize failure domain

Key Capabilities:

- Create and manage volumes on storage backend
- Expose volumes to physical host (e.g. iSCSI, FC)
- Uses plug-in model to support differing storage systems

cinder-scheduler selects cinder-volume instance to place volume on

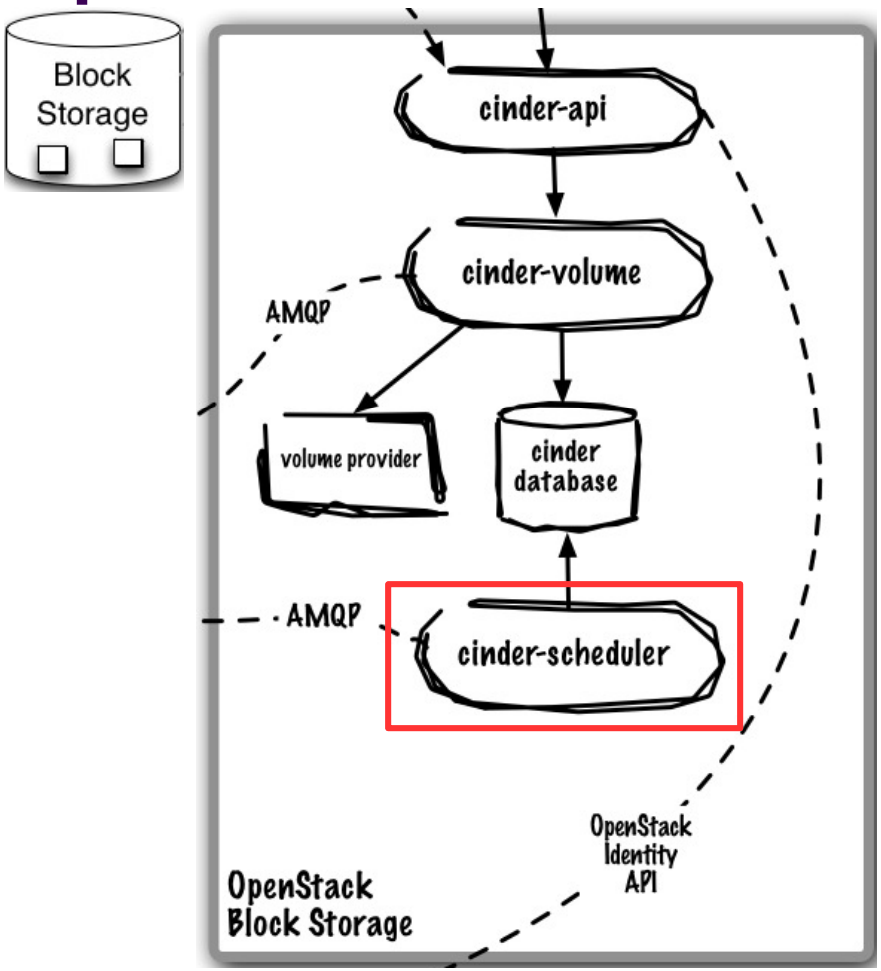

Core Use Case:

- Selects cinder-volume service to place volume on

Runs As: Controller Service

Deployment Considerations:

- Default scheduler is horizontally scalable

Key Capabilities:

- Default scheduled is allocation-based using a series of filters to reduce set of applicable hosts and uses costing functions to provide weight

Identity Service (Keystone) offers project-wide identity, token, service catalog, and policy services designed for integration with existing systems

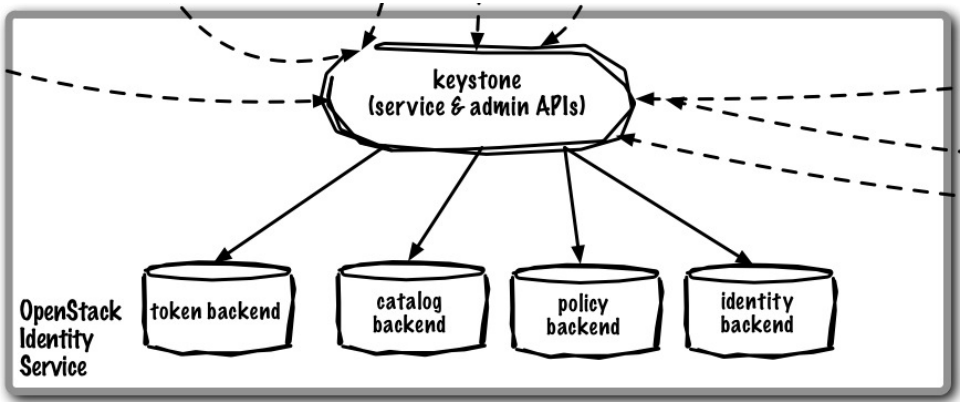

Core Use Cases:

- Installation-wide authentication and authorization to OpenStack services

Key Capabilities:

- Authenticate user / password requests against multiple backends (SQL, LDAP, etc) (Identity Service)
- Validate / manage tokens used after initial username/password verification (Token Service)
- Endpoint registry of available services (Service Catalog)
- Authorize API requests (Policy Service)
- Domain / Project / User model with RBAC for access to compute, storage, networking
- Policy service provides a rule-based authorization engine and the associated rule management interface.

Image Source: http://www.solinea.com/2013/04/17/openstack-summit-intro-to-openstack-architecture-grizzly-edition/

keystone service is the entry point for all AuthN and AuthZ in OpenStack

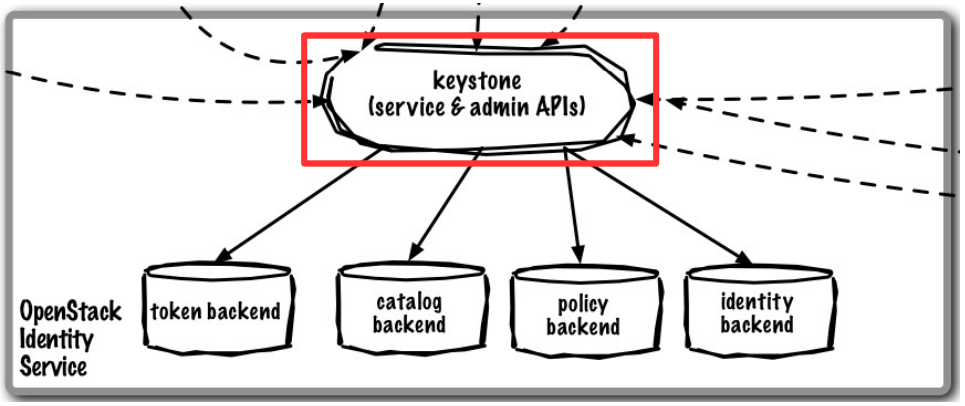

Core Use Case:

- Handle and service all Identity REST API requests

Runs As: Controller Service

Deployment Considerations:

- Horizontally scalable, start many instances
- Front with load-balancer to present as single endpoint

Key Capabilities:

- APIs supported
  - OpenStack Identity API
- Pluggable backends for each function: identity, token, catalog, and policy

Glance database persists all image related metadata

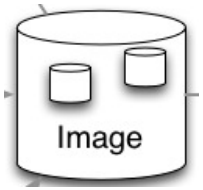

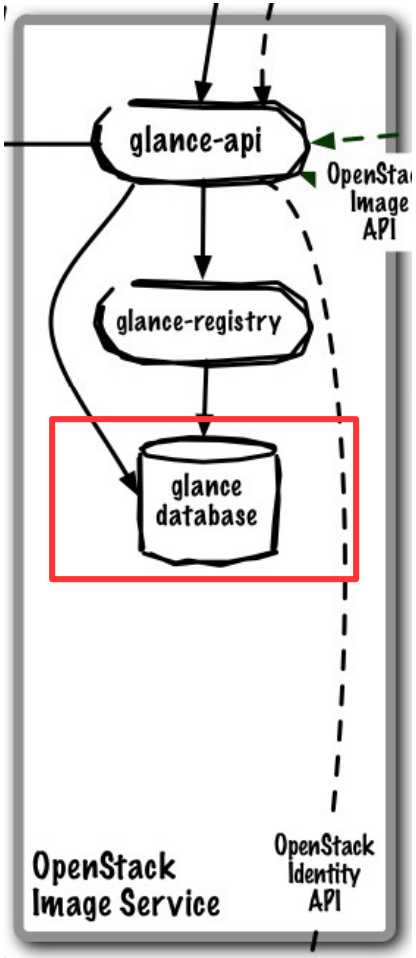

Core Use Case:

- Persist image-related metadata

Runs As: Controller Service

Deployment Considerations:

- Use DB and Queue clustering/HA methods
- Can use same queue/database as Nova

Key Capabilities:

- Persists image-related metadata

Image Service (Glance) provides registration, discovery, and delivery services for virtual disk a erve

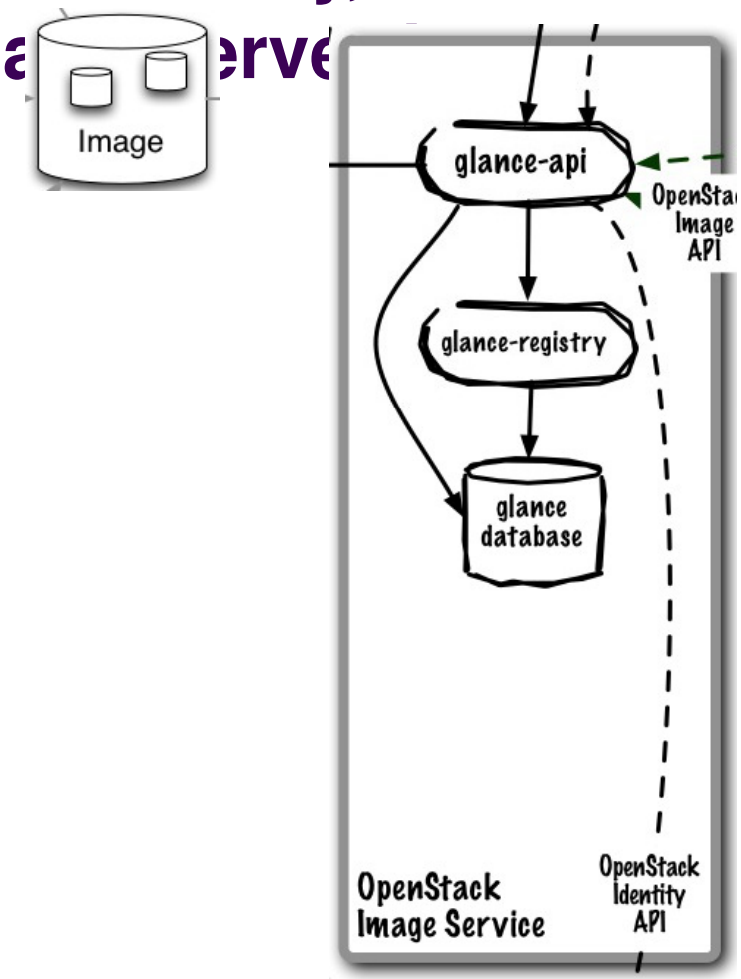

Core Use Cases:

- Administrator registers available guest images
- End-user discovers available guest images
- Deliver image to compute node on provisioning

Key Capabilities:

- Image Registry (storage optional and is delegated to a configurable store)
- Administrators can create base templates from which users can start new compute instances
- Users can choose from available images, or create their own from existing servers
- Snapshots can also be stored in the Image Service so that virtual machines can be backed up quickly
- Supported formats: Raw, Machine (a.k.a. Amazon AMI), VHD (Hyper-V), VDI (VirtualBox), qcow2 (Qemu/KVM), VMDK (VMWare), OVF (VMWare, others)

Image Source: http://www.solinea.com/2013/04/17/openstack-summit-intro-to-openstack-architecture-grizzly-edition/

glance-api routes incoming REST API Requests

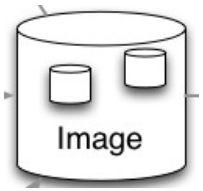

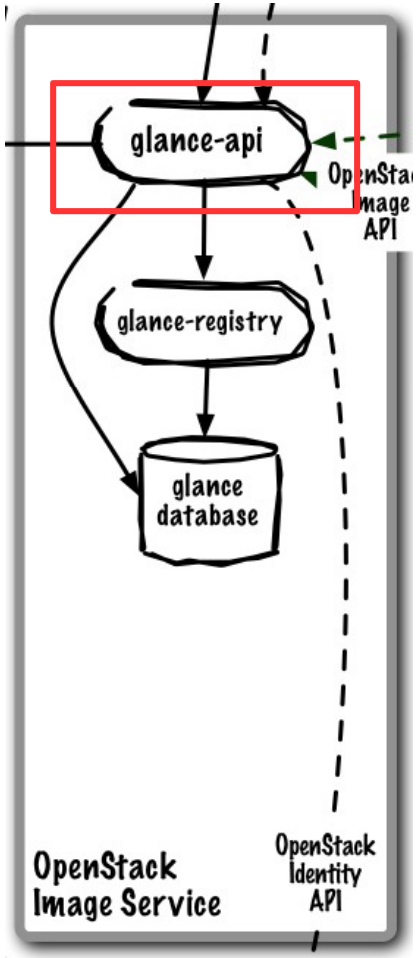

Core Use Case:

- Routes REST API requests to the appropriate handler

Runs As: Controller Service

Deployment Considerations:

- Horizontally scalable, start many instances
- Front with load-balancer to present as single endpoint

Key Capabilities:

- APIs supported
  - OpenStack Image API
- Routes requests from clients to registries of image metadata and to its backend stores
- Pluggable image store backends

glance-registry services Image Service API requests

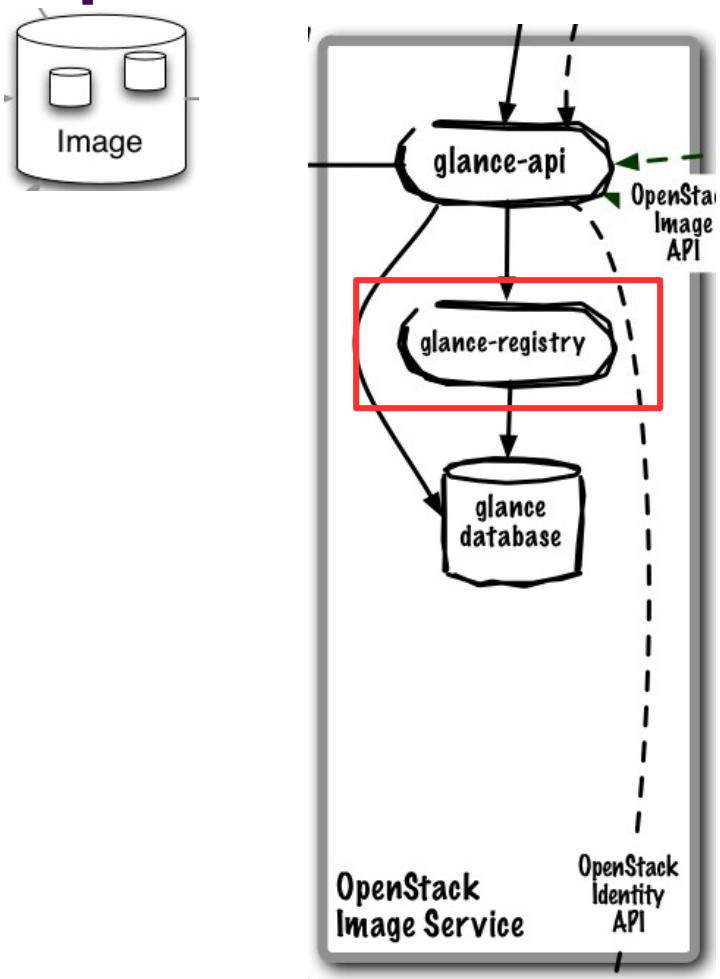

Core Use Case:

- Services Identity REST API requests

Runs As: Controller Service

Deployment Considerations:

- (to be determined)

Key Capabilities:

- APIs supported
  - OpenStack Image API

Horizon (Dashboard) enables administrators and users to access, provision, and manage resources through a self-service portal GUI

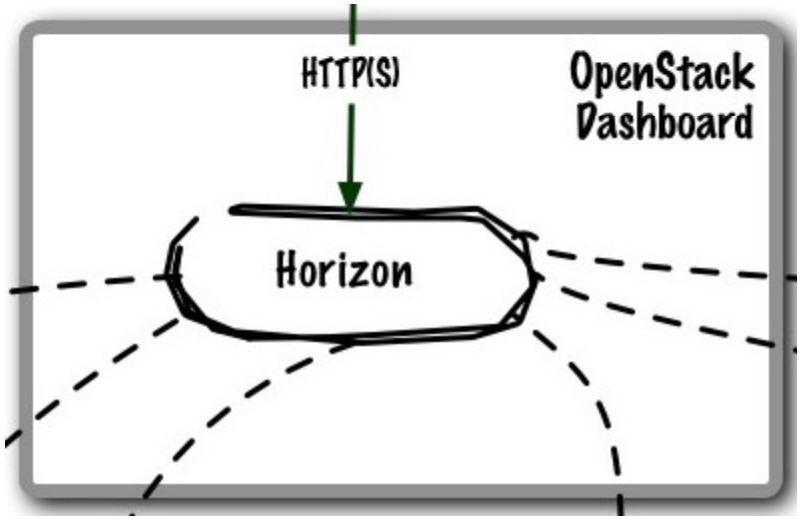

NOT SHIPPED BY IBM

Core Use Cases:

- Self-service portal for compute and object storage
- Cloud administration (users/projects, quotas, etc.)

Key Capabilities:

- Thin wrapper over APIs, no local state
- Registration pattern for applications to hook into
- Out-of-the-box support for all core OpenStack projects.
- Anyone can add a new component as a “first-class citizen”.
- Visual and interaction paradigms are maintained throughout.

Image Source: http://www.solinea.com/2013/04/17/openstack-summit-intro-to-openstack-architecture-grizzly-edition/

horizon is the self-service portal implementation

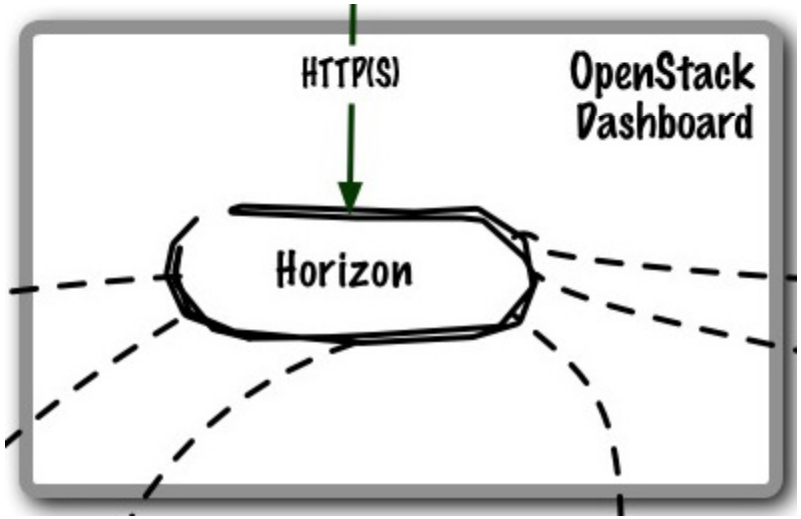

NOT SHIPPED BY IBM

Core Use Case:

- GUI access to OpenStack APIs

Runs As: Controller Service

Deployment Considerations:

- (to be determined)

Key Capabilities:

- Provision and manage virtual servers, volumes, and networks
- Create and manage tenants and users

Putting it all together....

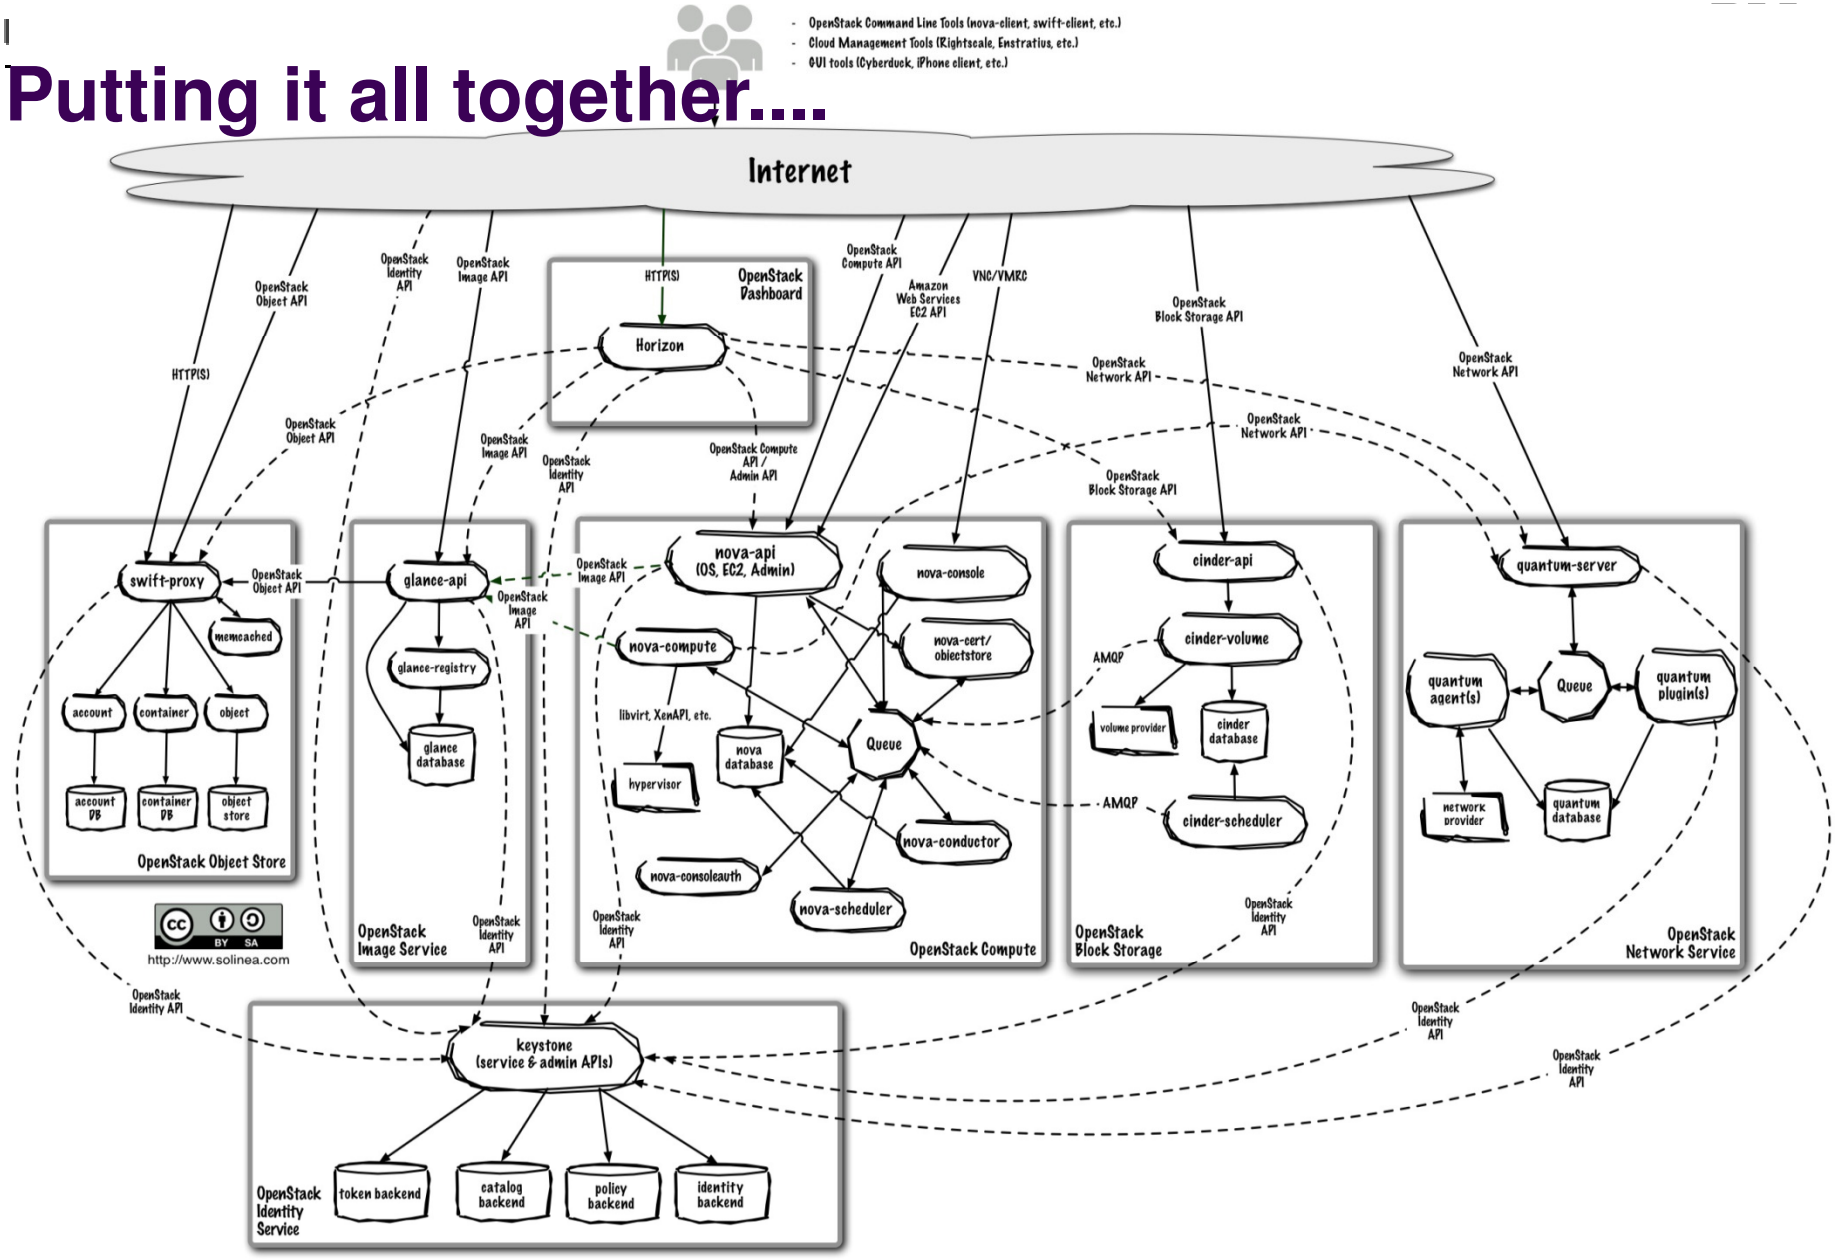

Backup Material
Additional xCAT Information

xCAT to be shipped and installed with z/VM

- xCAT Extreme Cloud Administration Toolkit
- Allows for Provisioning
- Function added to xCAT that will be ready for z/VM 6.3
  - **Capture** – ability to capture a virtual machine into an image that can be deployed at a later time.
  - **Deploy** – ability to deploy an image from previously captured image data.
  - **Export** – ability to export an image from a local repository for retention or transmission elsewhere.
  - **Import** – ability to import an image (possibly created elsewhere) into the local image repository.
  - **Delete Image** – ability to remove an image from a local image repository.
  - **List** – ability to discover images in a local image repository and information about the images. This is of primary use to xCAT.
- RESTful API's or interface for OpenStack communication to z/VM

xCAT Benefits to customers

- Provides out-of-the-box ease of use for many customer tasks
  - Server Deployment
  - Server Lifecycle management
  - Basic Server monitoring
- Infrastructure for Additional products to interface to z/VM using the OpenStack interfaces

xCAT Configuration into z/VM system

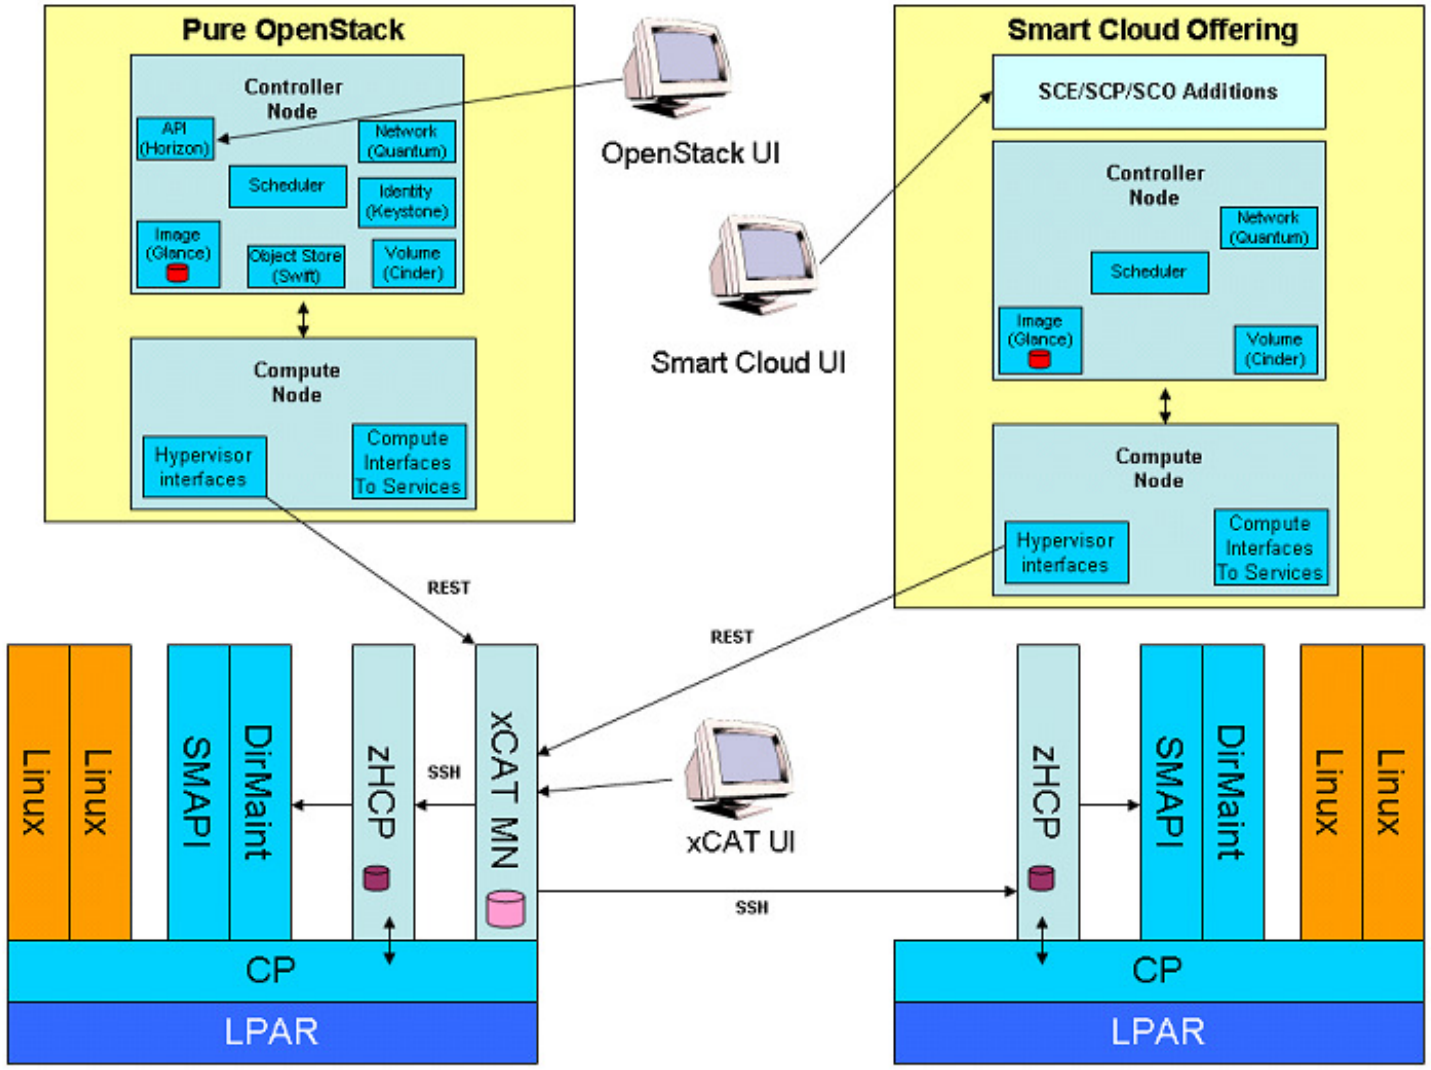

xCAT Node (or server) screen

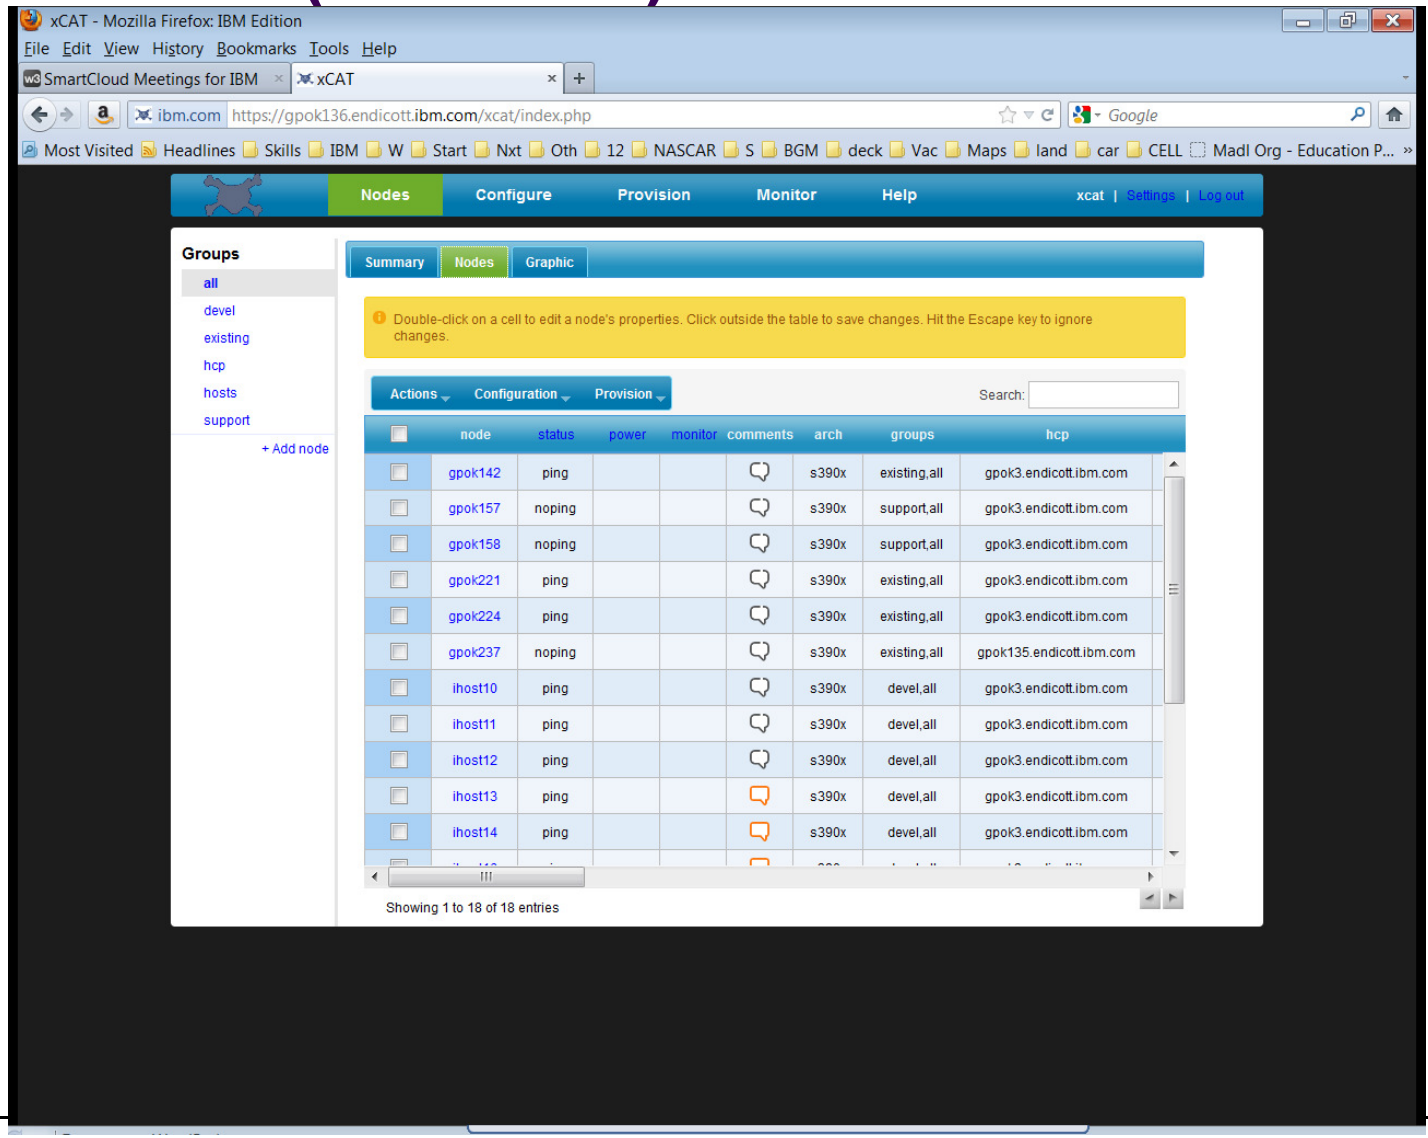

xCAT Node Details

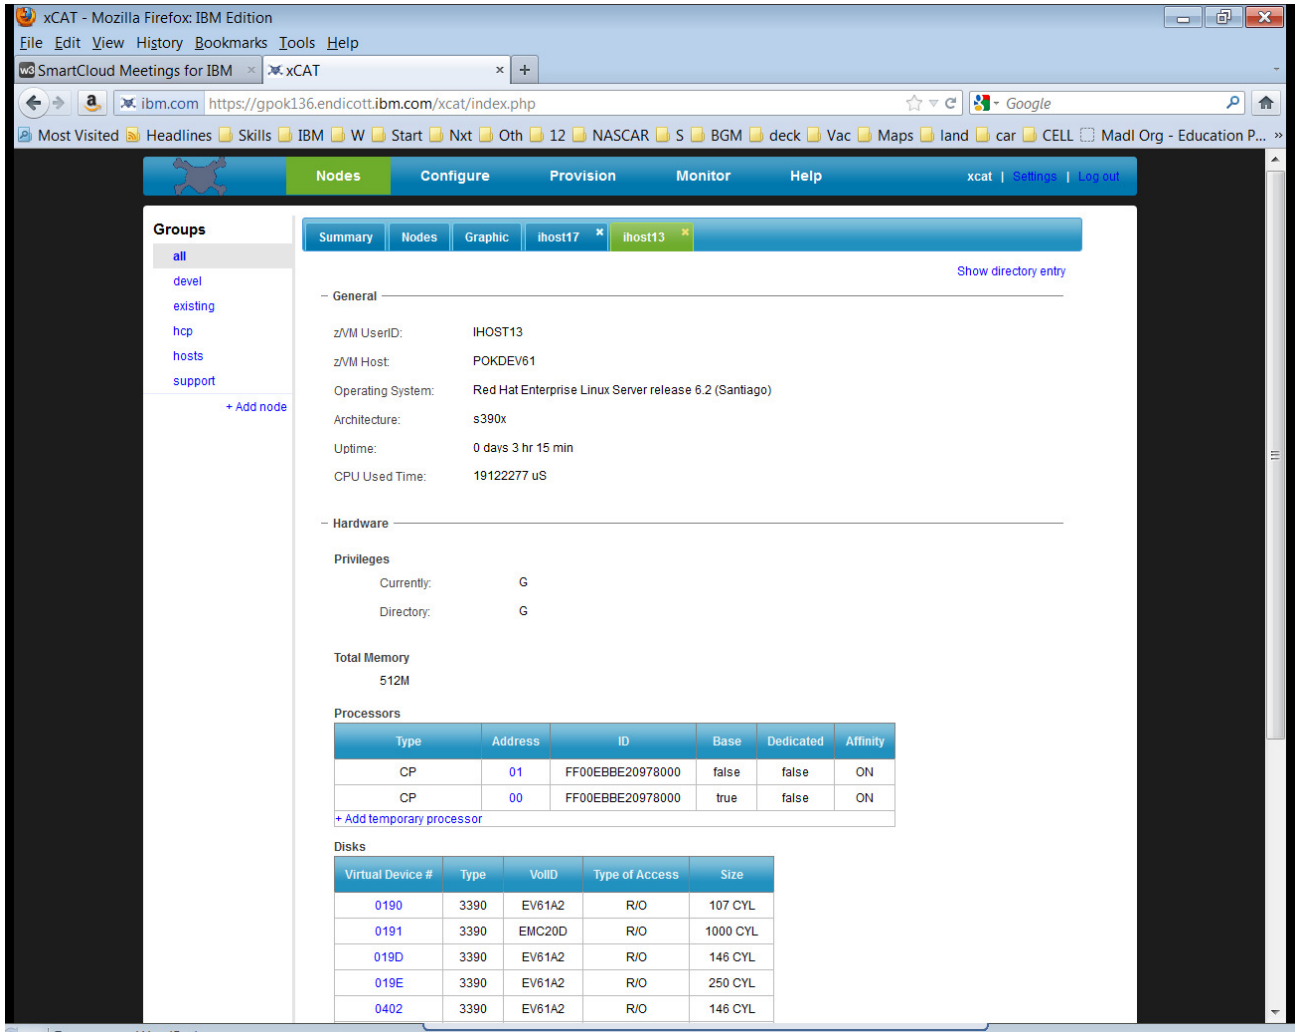

xCAT Node Details - cont.

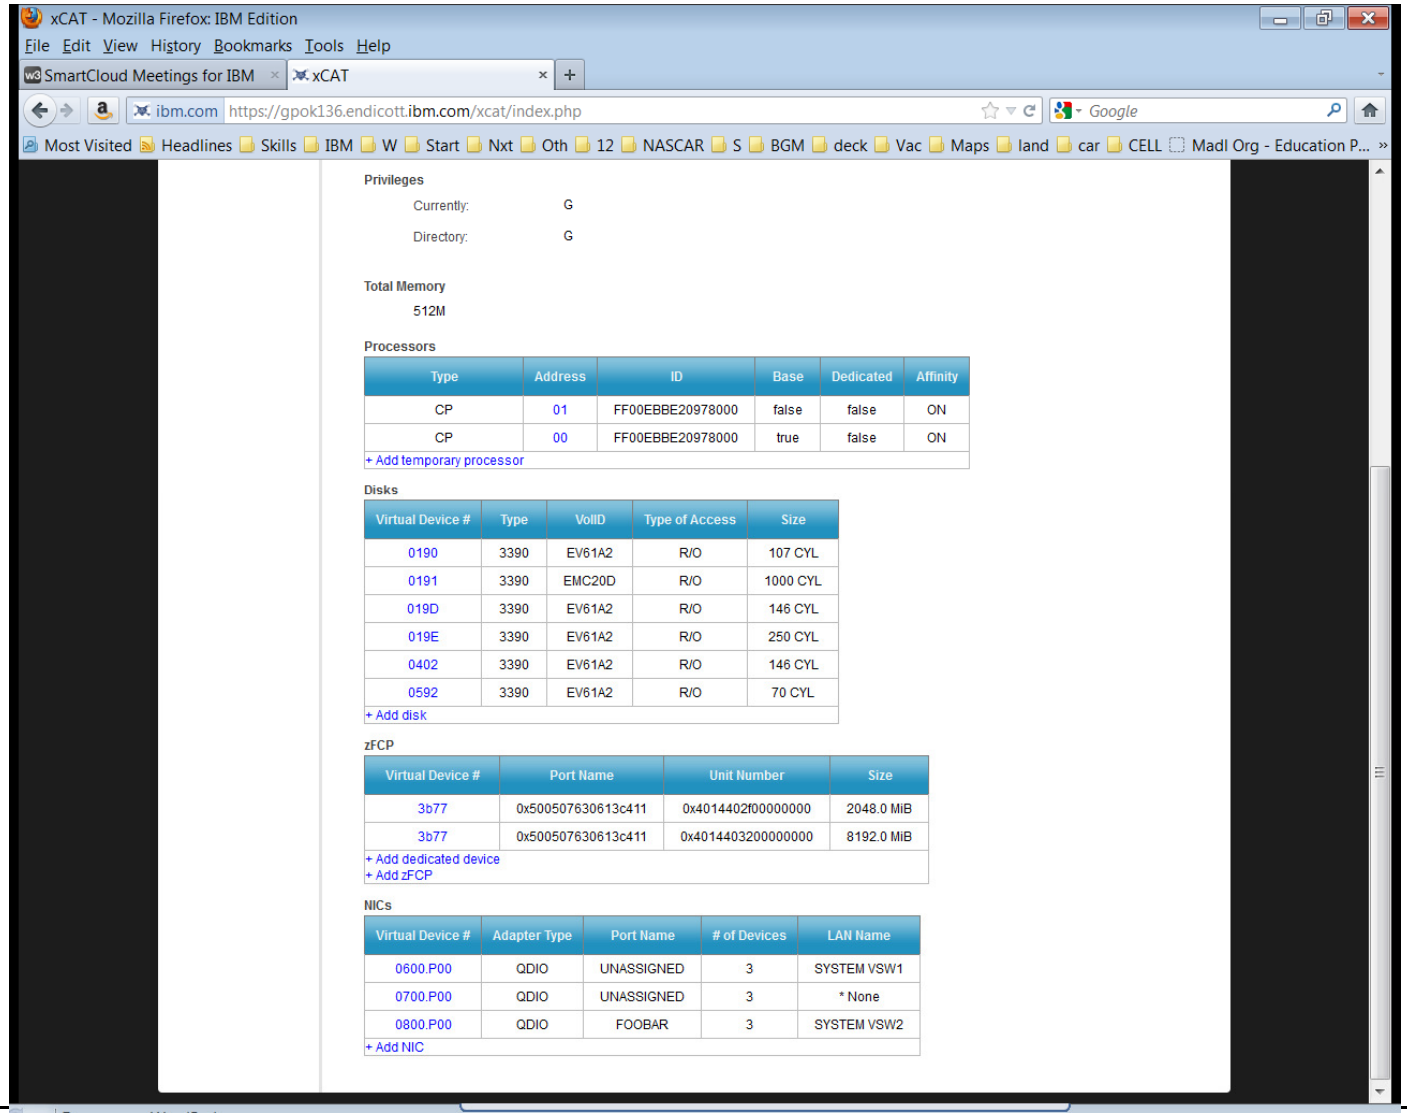

Privileges

Currently: G

Directory: G

Total Memory

512M

Processors

Type	Address	ID	Base	Dedicated	Affinity
CP	01	FF00EBBE20978000	false	false	ON
CP	00	FF00EBBE20978000	true	false	ON

+ Add temporary processor

Disks

Virtual Device #	Type	VolID	Type of Access	Size
0190	3390	EV61A2	R/O	107 CYL
0191	3390	EMC20D	R/O	1000 CYL
019D	3390	EV61A2	R/O	146 CYL
019E	3390	EV61A2	R/O	250 CYL
0402	3390	EV61A2	R/O	146 CYL
0592	3390	EV61A2	R/O	70 CYL

+ Add disk

zFCP

Virtual Device #	Port Name	Unit Number	Size
3b77	0x500507630613c411	0x4014402f00000000	2048.0 MiB
3b77	0x500507630613c411	0x4014403200000000	8192.0 MiB

+ Add dedicated device
+ Add zFCP

NICs

Virtual Device #	Adapter Type	Port Name	# of Devices	LAN Name
0600.P00	QDIO	UNASSIGNED	3	SYSTEM VSW1
0700.P00	QDIO	UNASSIGNED	3	* None
0800.P00	QDIO	FOOBAR	3	SYSTEM VSW2

+ Add NIC

xCAT Directory Entry (editable directly)

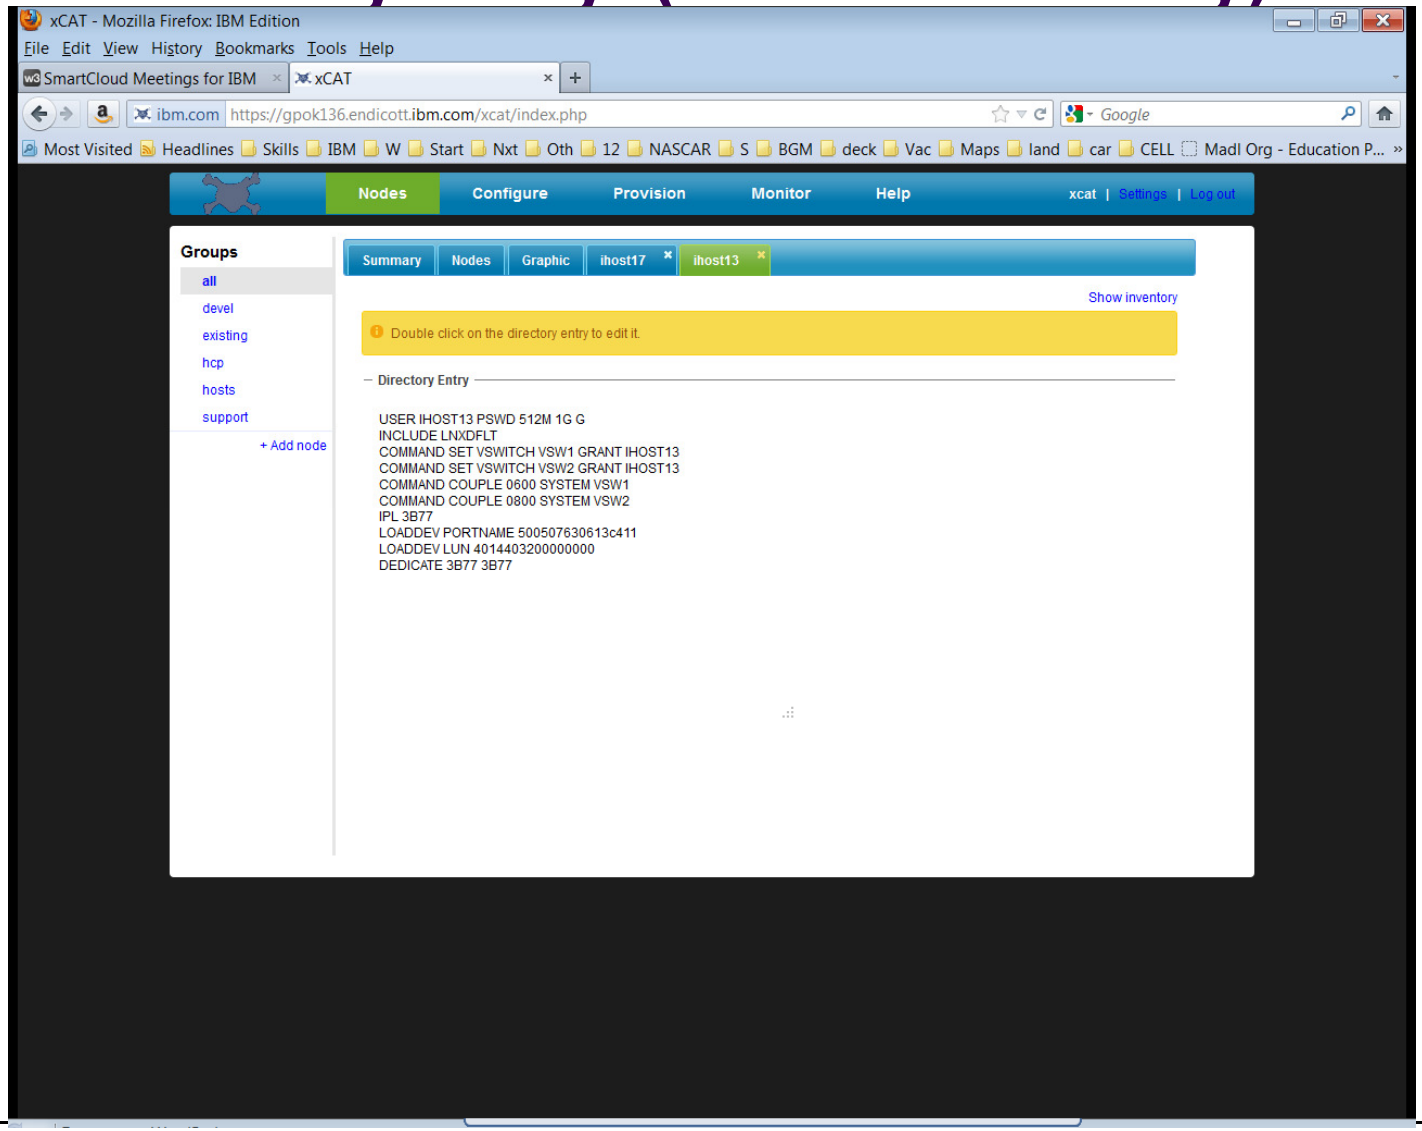

Backup Material
Additional IBM Wave Information

IBM Wave Intelligent Visualization

Quickly Understand the Status of System Resources

- **Get a current and accurate view of your managed environment**
  - Network Topology
    - Centralized view of the entire network topology per z/VM System, view Virtual LANS (VLANS)
    - Annotate network topology view to identify external resources - routers, switches, etc
  - Linux Servers
    - View performance gauges for all z/VM systems from one screen:
    - See resource consumption by guest or type
    - CPU, Virtual to Real, Paging, Spool
  - Storage
    - Visual representation of all storage resources

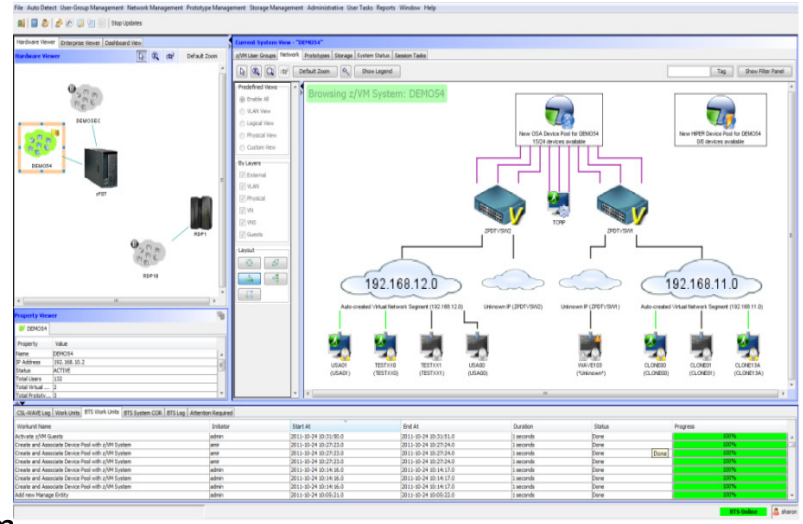

- **Visualize and control virtual resources**
  - Views can be graphical or easily switched to tabular mode
  - View relationships between resources easily and graphically
  - View the entire environment graphically and easily zoom in
- **Advanced filters, tagging, layout and layer based views for every display**

Intelligent Visualization

Mapping Attached Storage to Guests and Storage Controllers

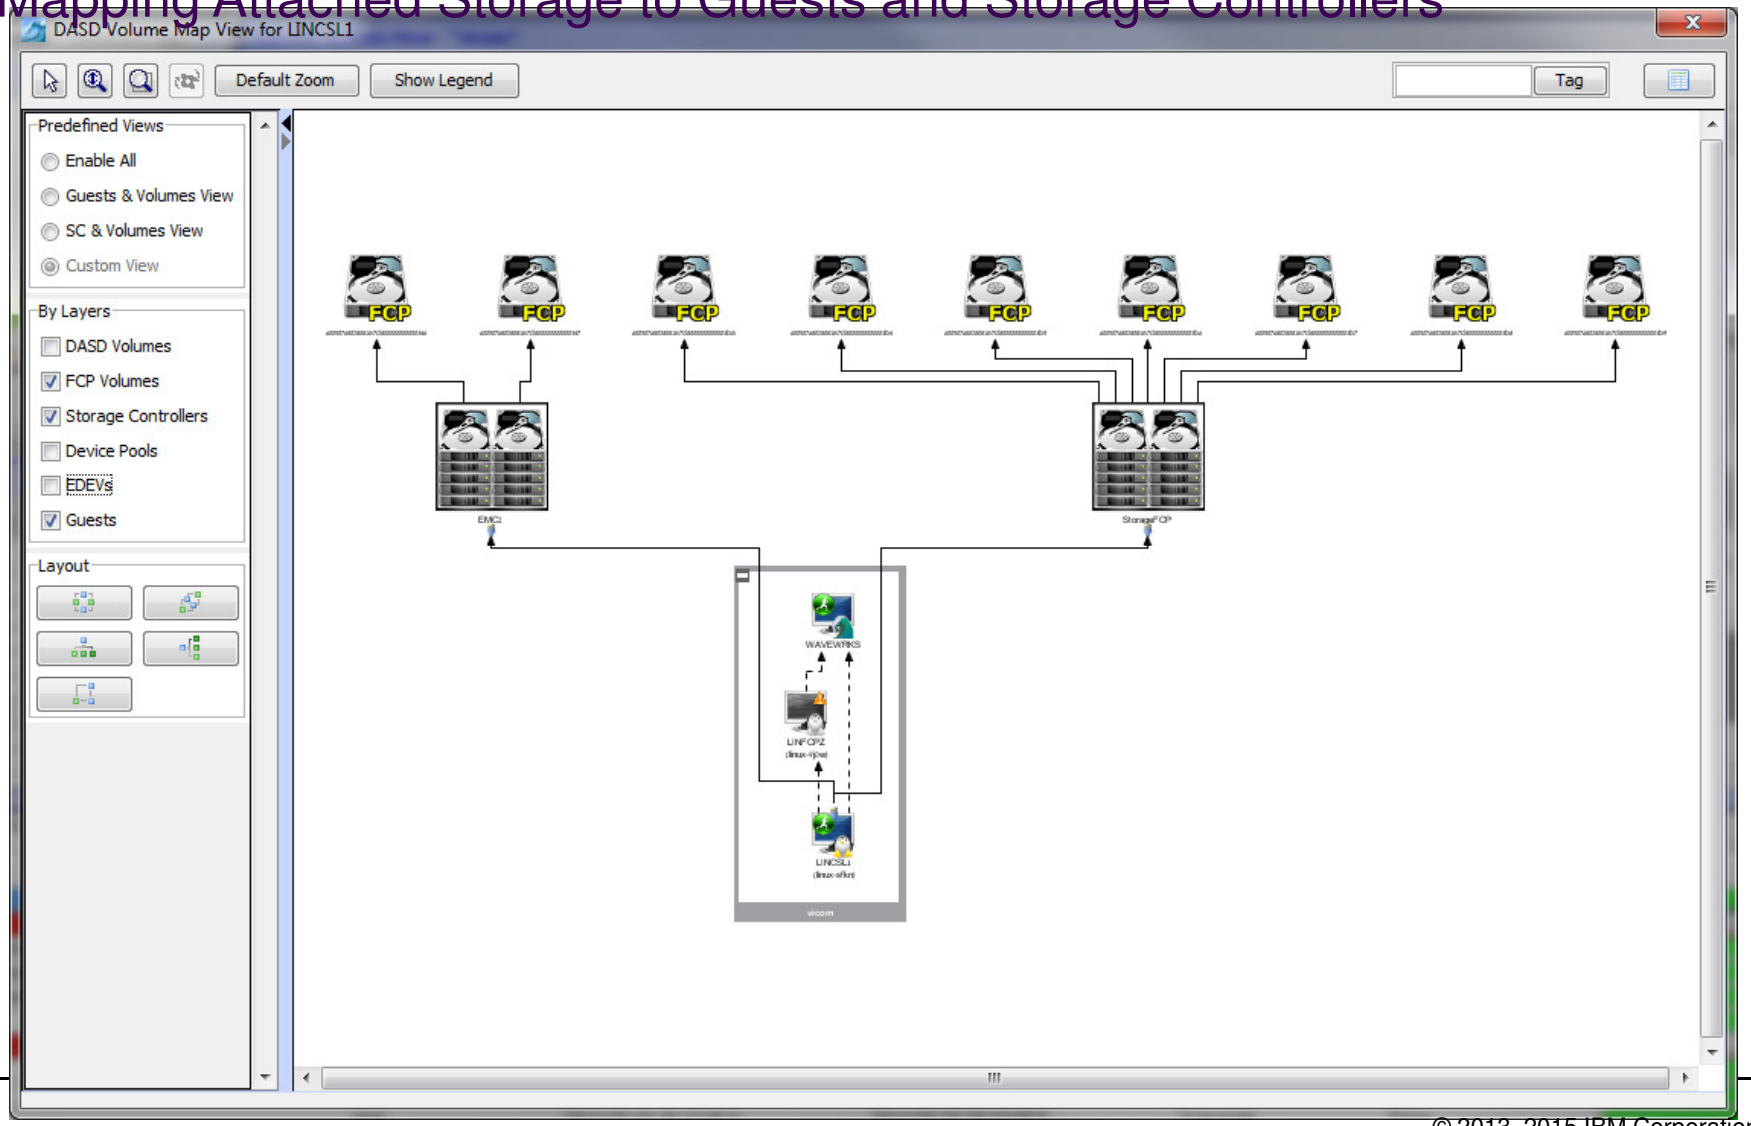

IBM Wave Simplified Monitoring

Automatic Detection and Monitoring of Resources

- **Agentless Resource Discovery**
  - Discover, manage and monitor z/VM resources and their relationships across multiple LPARs and CECs
  - Identify resource and relationship changes; reflect current environment in the user interface
- **Monitoring**
  - Allows the state of resources to be observed; icons show additional content for the resources
  - Use graphical and tabular displays with layered drill down to hone in on only the resources you need to view
  - Perform ongoing monitoring of changes that occur after initial auto-detection
- **Reporting**
  - Automatically generate charts like pie charts to report on utilization and more
  - All table-based views can be exported to a CSV file for import into other applications

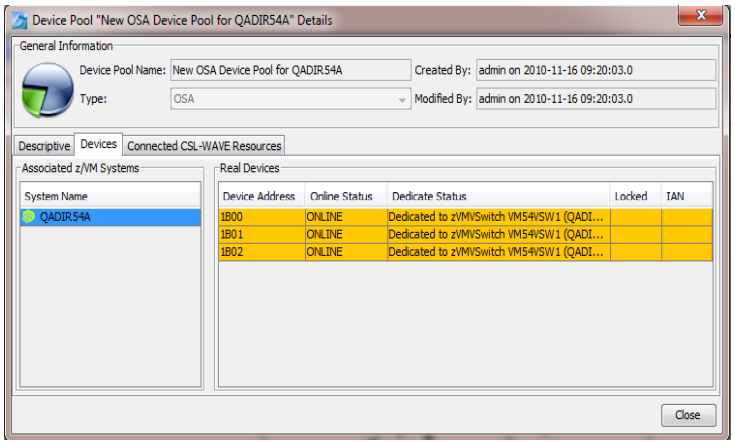

Performance Resource Monitoring

At a Glance Status of all z/VM instances

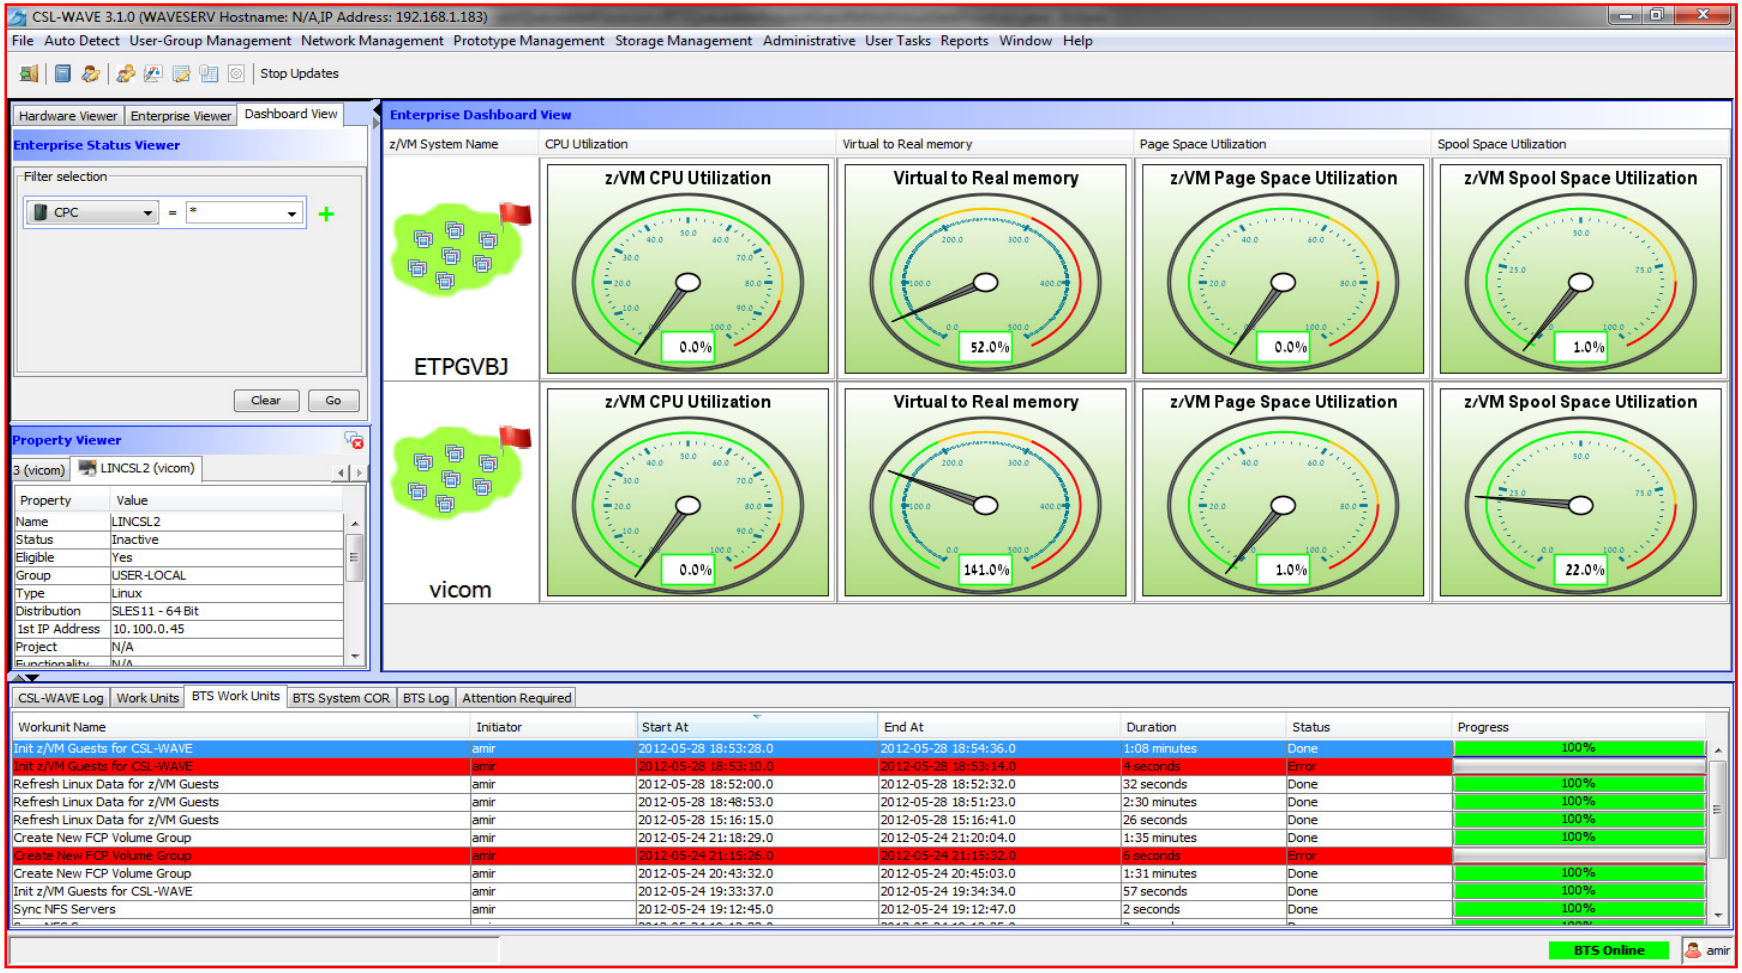

Performance Resource Monitoring

At a Glance Drill down to Virtual Memory Detail

Current Real and Virtual Memory Values

Type	Size	Ratio Display
Total Real Memory	7000M	7000
Total Virtual Memory	6020M	6020

Virtual to Real Ratio: 86%

Virtual Memory Utilization for All Guests - Last Updated: 08:45:24 EDT WEDNESDAY 06/19/13

Name	Type	Memory (MB)	VDisks (MB)	Total (MB)	% Used
CSLSSI2	CMS	1200	0	1200	19%
CSLSSI1	CMS	1200	0	1200	19%
NEWTEST1	Linux	128	500	628	10%
ISRAEL0	Linux	128	0	128	2%
TEMP0	Linux	128	0	128	2%
IONPWR2	Linux	128	0	128	2%
IONPWR1	Linux	128	0	128	2%
WAVEWRKC	CSL-WAVE	32	0	32	0%
WAVEWRKL	CSL-WAVE	32	0	32	0%
WAVEWRKS	CSL-WAVE	32	0	32	0%

Include IBM Service Machines

Virtual Memory Usage Distribution (By Guest)

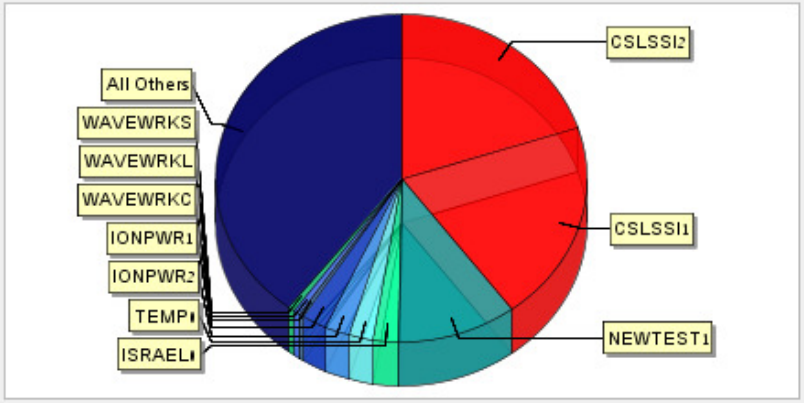

Virtual Memory Usage Distribution (By Type)

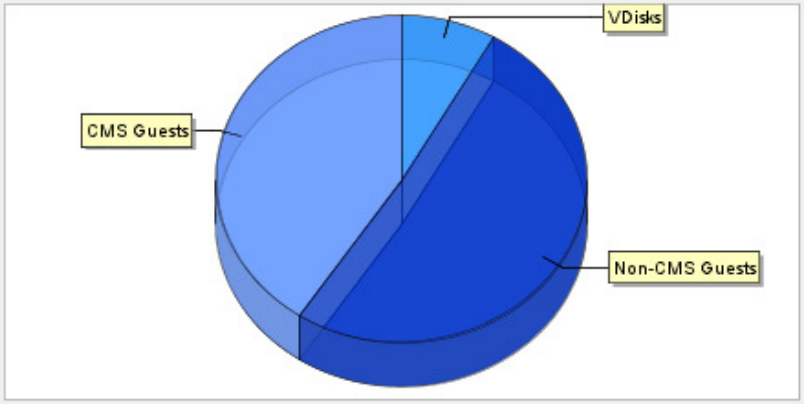

IBM Wave Unified Management

Managing the Entire Pool of Resources Intuitively

- **Simplification**
  - Simplify the process of performing a function across multiple z/VM or Linux systems
- **Manage Networks**
  - Centralized, layer based customizable view of the entire z/VM network topology
  - Define and control all network devices such as VSWITCHes and guest LANs
- **Manage Storage**
  - Manage devices and device pools
- **Provision Resources**
  - Clone resources and virtual servers, apply scripts for more customization
  - Install Linux on virtual machines
  - Adding attached storage to z/VM Guests using IBM Wave's Manage Storage Wizard
  - Activate/Deactivate servers in an ordered fashion
- **Improve Policy Management**
  - Use reminder notes attached to icons to provide advisory and policy notices

Simplify Systems Management Tasks

Provision resources quickly and easily

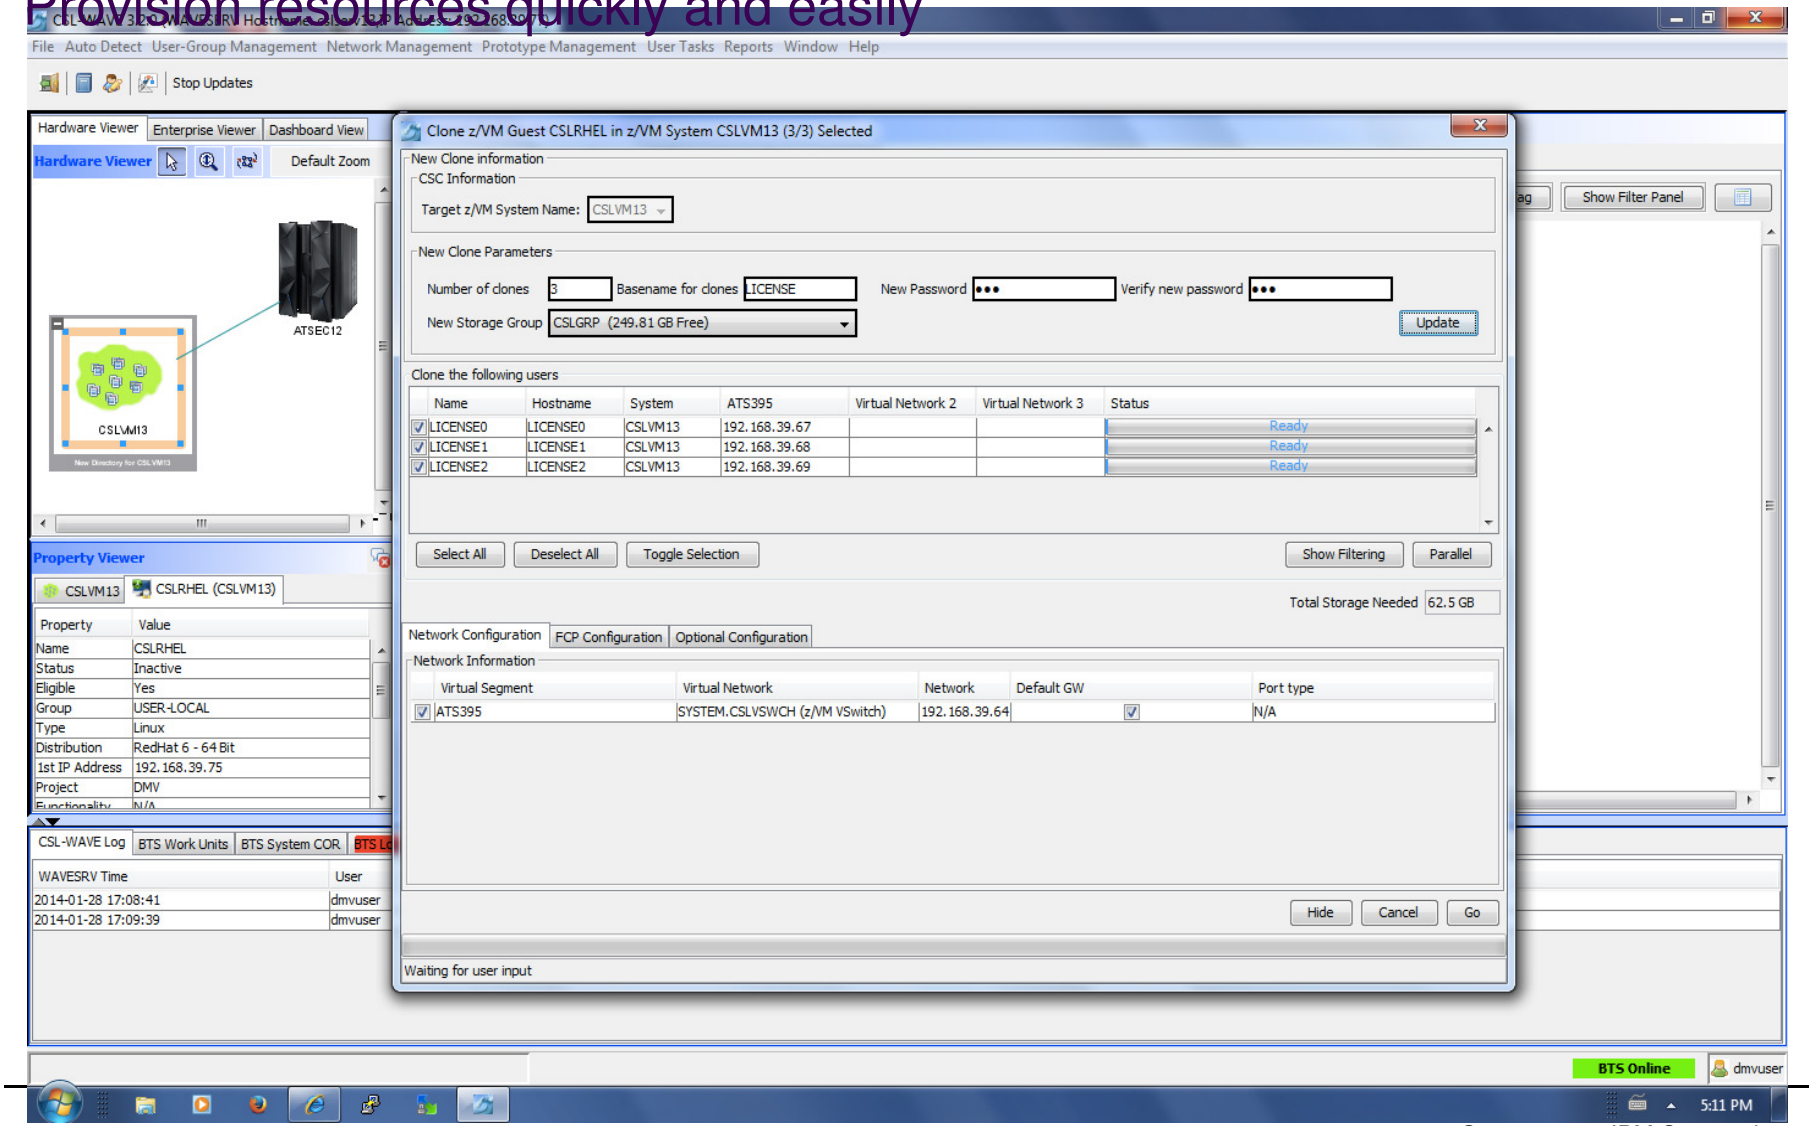

IBM Wave Systems Management Task Example

Add Disk Space To A Virtual Server

Without IBM Wave

1. Find requested disk space
2. Create disk definition
3. Activate definition
4. Connect storage to virtual server
5. Mount device
6. Create a File System

With IBM Wave

1. Open the “Add Storage” form
2. Fill the storage capacity requested
3. Press the “Go” button

Benefits:

- ✓Reduce reliance on scarce skills
- ✓Respond faster to IT customer needs
- ✓Reduce costs
- ✓Empower team to do more independently
- ✓Simplify management
- ✓Accurately depict current environment
- ✓Reduce manual procedure errors
- ✓Avoid problematic situations downstream

View Storage at a Glance

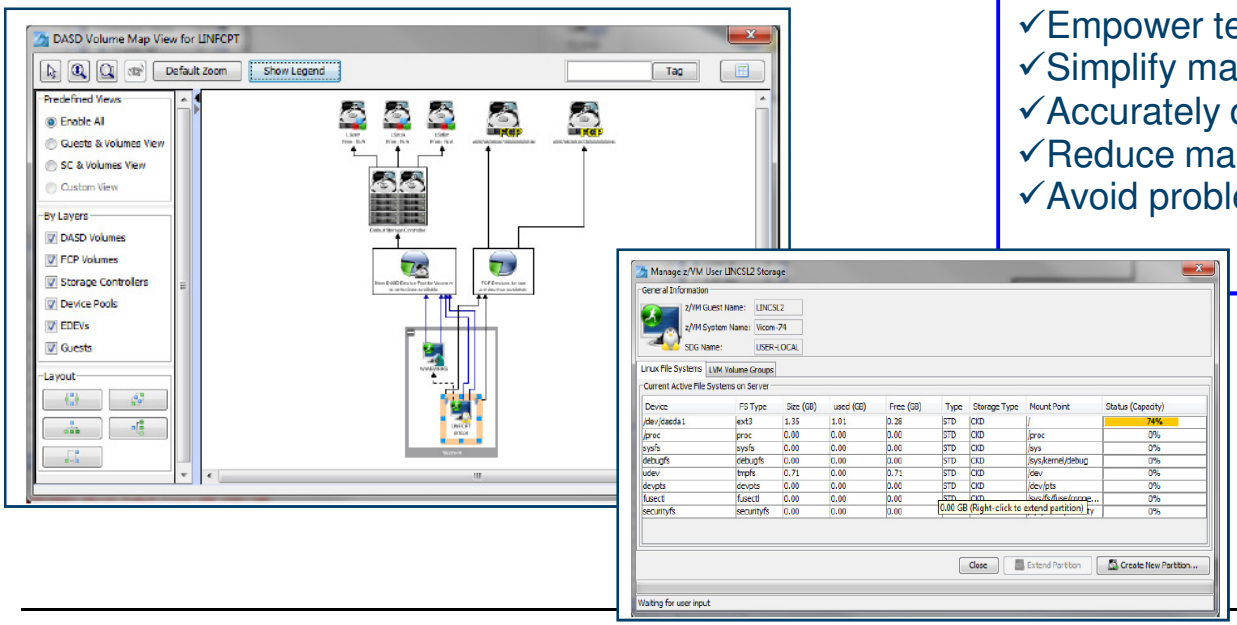

IBM Wave Systems Management Task Example:

Clone a Virtual Machine

Without IBM Wave

1. Determine if required resources exist
2. Create clone VM definition
3. Define clone VM resources
4. Create copies of private VM resources (server)
5. Create copies of private VM resources (disk)
6. Customize clone VM
7. Authorize clone VM access / VSwitch Access
8. Add clone to management groups
9. Activate clone
10. Configure the network
11. Run middleware configuration scripts
12. Monitor and report on cloning operation.

Clone a Linux Virtual Server

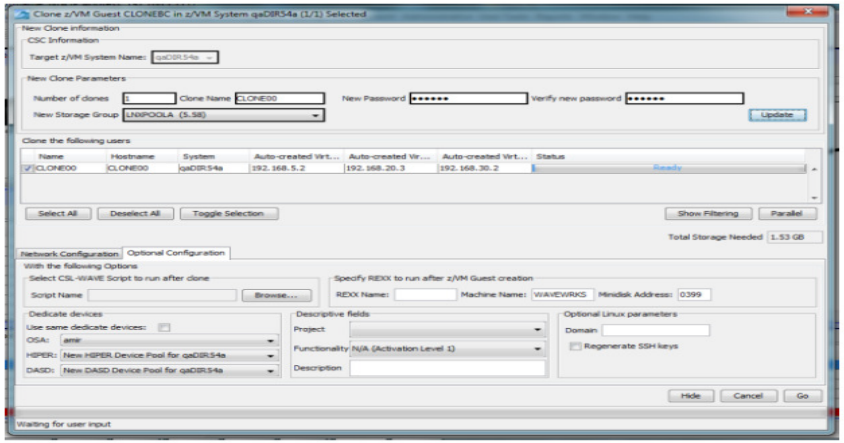

With IBM Wave

1. Open the “Clone” form
2. Fill in the needed information
3. Press the “Go” Button

Benefits:

- ✓Reduce time for a highly complex task
- ✓Reduce costs
- ✓Reduce reliance on scarce skills
- ✓Improve speed to clone
- ✓Simplify management
- ✓ Reduce errors associated with manual procedures
- ✓No need to monitor every step of the process

IBM Wave Systems Management Task Example

Live Guest Relocation

Without IBM Wave

- Using manual control program commands

With IBM Wave

- Using the GUI's Drag-and-Drop techniques
- Or Execute via menu selection

Task	Task Steps
Log into both z/VM instances	Login PBCVMA Login PBCVMB
Find out which instance has the running guest	q HTTP2 in PBCVMA q HTTP2 in PBCVMB
Verify the guest can be moved	vmrelo test HTTP2 to PBCVMB
Move the guest	vmrelo move HTTP2 to PBCVMB
Log out of both z/VM instances	Logoff PBCVMA Logoff PBCVMB

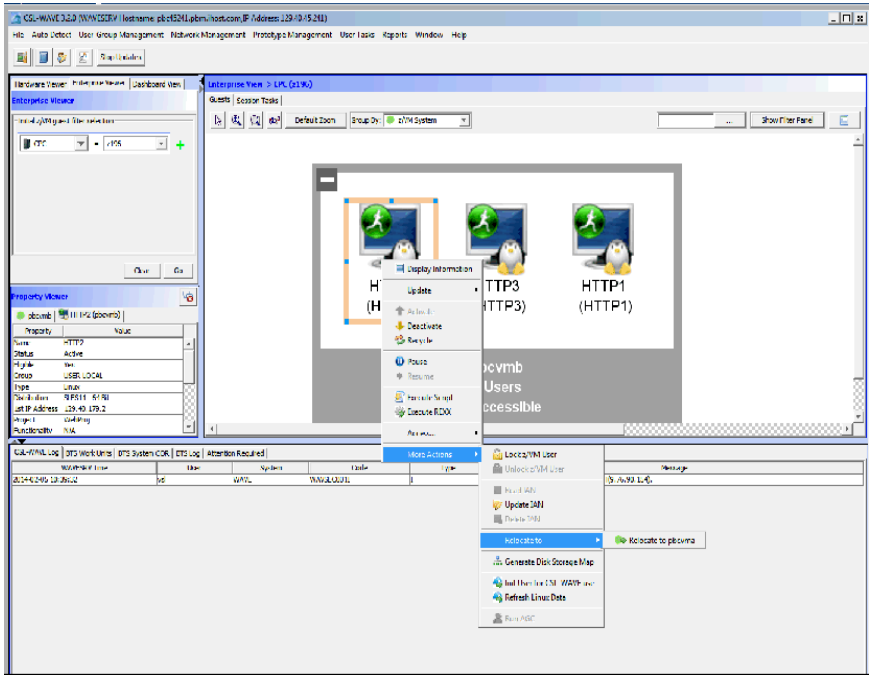

Benefits	IBM Wave for z/VM Capabilities
✓ Gain efficiencies in virtualization management ✓ Work with a current, accurate and complete view of your managed z/VM environment	▪ IBM Wave provides a high level view of performance, storage usage, networks at a glance with built-in reporting ▪ IBM Wave enables automation of management tasks and can incorporate scripts. ▪ By providing an up to date, accurate view of the IT environment through its “agent-less discovery” organizations can plan, change and optimize their virtualized resources accurately
✓ Simplify administrative, operations and systems functions ✓ Enable improved self service to reduce costs	▪ Tasks that would otherwise take hours and require significant z/VM knowledge such as Live Guest Relocation, Server Cloning and Storage provisioning can be performed quickly and easily ▪ Make common management tasks accessible to more user roles
✓ Respond quickly to changing business needs ✓ Reduce errors with appropriate delegation	▪ Easily delegate administrative capabilities to the appropriate users ▪ Enforce segregation policies at the individual administrator as well as the group level ▪ Set scope and permissions to match business requirements

Benefits	IBM Wave for z/VM Capabilities
✓ Improve service levels ✓ Easily respond to changing requirements. ✓ Reduce time spent on administrative efforts	▪ Offers easy, convenient access to performance and management information –at a glance ▪ Helps you quickly and easily administer and provision resources like servers, storage, user accounts. ▪ Tag resources with meaningful notes to help enforce installation defined rules.
✓ Easily manage virtualized environments ✓ Simplify and accelerate your journey to cloud	▪ Lets you provision new servers and easily clone Linux virtual servers and other resources ▪ Scripts allow customization of a golden master. ▪ Support early virtualization steps needed to get to a private cloud.
✓ Create audit trails of IBM Wave users' activities	▪ List tasks and status requested by the users with respect to their scope. ▪ Log each operation that changes the system including logon and logoff to provide an audit trail. The logs may be then routed to a centralized logging mechanism for further filtering or processing.
✓ Simplify your administration ✓ Extend the reach of your existing IT staff	▪ IBM Wave automates a sequence of VM commands, reducing steps needed to complete common administrative and management tasks—and improve consistency. ▪ IBM Wave helps your team manage additional servers even if you do not have a deep expert skills bench available.

IBM Wave Architecture

Client

- The Client can run on Microsoft® Windows®, running Java™ 1.7
- Graphic interpretation of the TVP through communication with WAVESRV using Point-and-Click and Drag-and-Drop operations

WAVESRV

- This server (virtual or physical) hosts the application database and Background Task Scheduler
- One BTS server can manage many Target Virtualization Platforms.

TVP

- The Target Virtualization Platform (TVP) represents the hypervisor which hosts the virtual guests that are managed.
- The BTS utilizes the TVP API to query and perform changes to the TVP and hosted virtual guests.

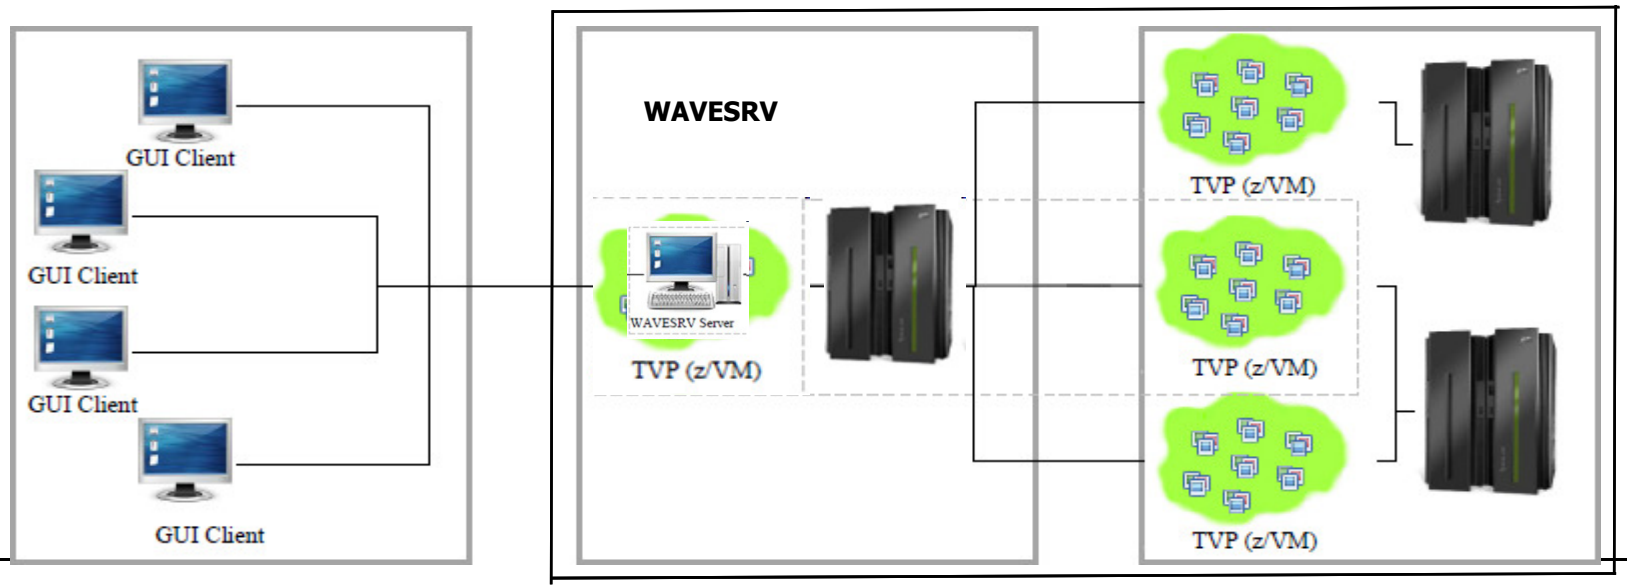

IBM Wave Requirements

Client

- Windows 7 Workstation
- Internet Explorer or Firefox
- Java Runtime 1.7 with Web Start Support
- PuTTY or equivalent telnet/SSH client

WAVESRV

- z/VM Guest or LPAR
- RHEL 6 or SLES 11
- MySQL V12.22 or higher
- Java SE Runtime 1.7
- Apache

TVP

- IBM z Systems10® or later
- z/VM V5.4, V6.2 or higher with Systems Management API configured
- IBM Directory Maintenance for z/VM (DirMaint™) or equivalent
- Performance Toolkit for VM™ (Perfkit, optional but suggested)