Emily K. Hugenbruch – Advisory Software Engineer

OpenStack and z/VM – What is it and how do I get it?

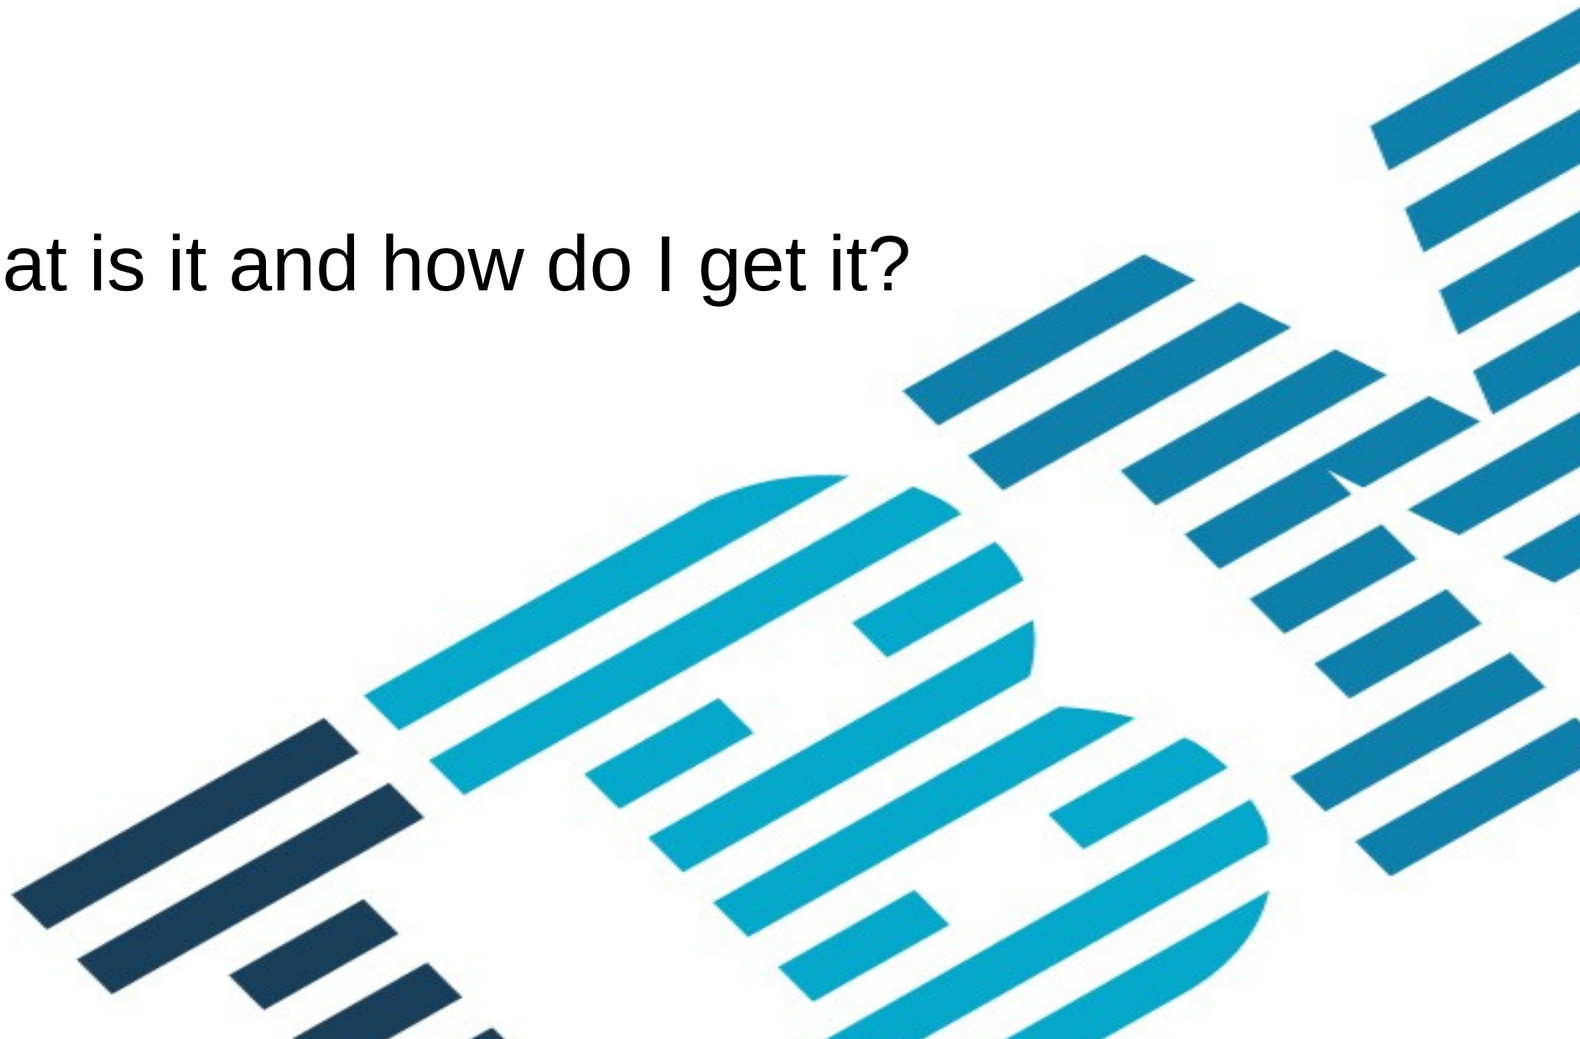

Trademarks

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

BladeCenter*
DB2*
DS6000*
DS8000*
ECKD
FICON*
GDPS*
HiperSockets
HyperSwap
IBM z13*
OMEGAMON*
Performance Toolkit for VM
Power*
PowerVM
PR/SM
RACF*
Storwize*
System Storage*
System x*
System z*
System z9*
System z10*
Tivoli*
zEnterprise*
z/OS*
zSecure
z/VM*
z Systems*

* 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 .
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.

Who is OpenStack for?

OpenStack targets these

End User Self-service IT requests via mobile & Web clients	**Cloud Manager** Manage cloud workloads Manage Cloud Users Manage Cloud environment
Platform Manager Manage physical resources (servers, storage and networking)	**Hypervisor Admin** Manage hypervisor resources (servers, storage and networking)

What is OpenStack?

Open source software for creating private and public clouds.

OpenStack software controls large pools of compute, storage, and networking resources throughout a datacenter, managed through a dashboard or via the OpenStack API.

What is OpenStack?

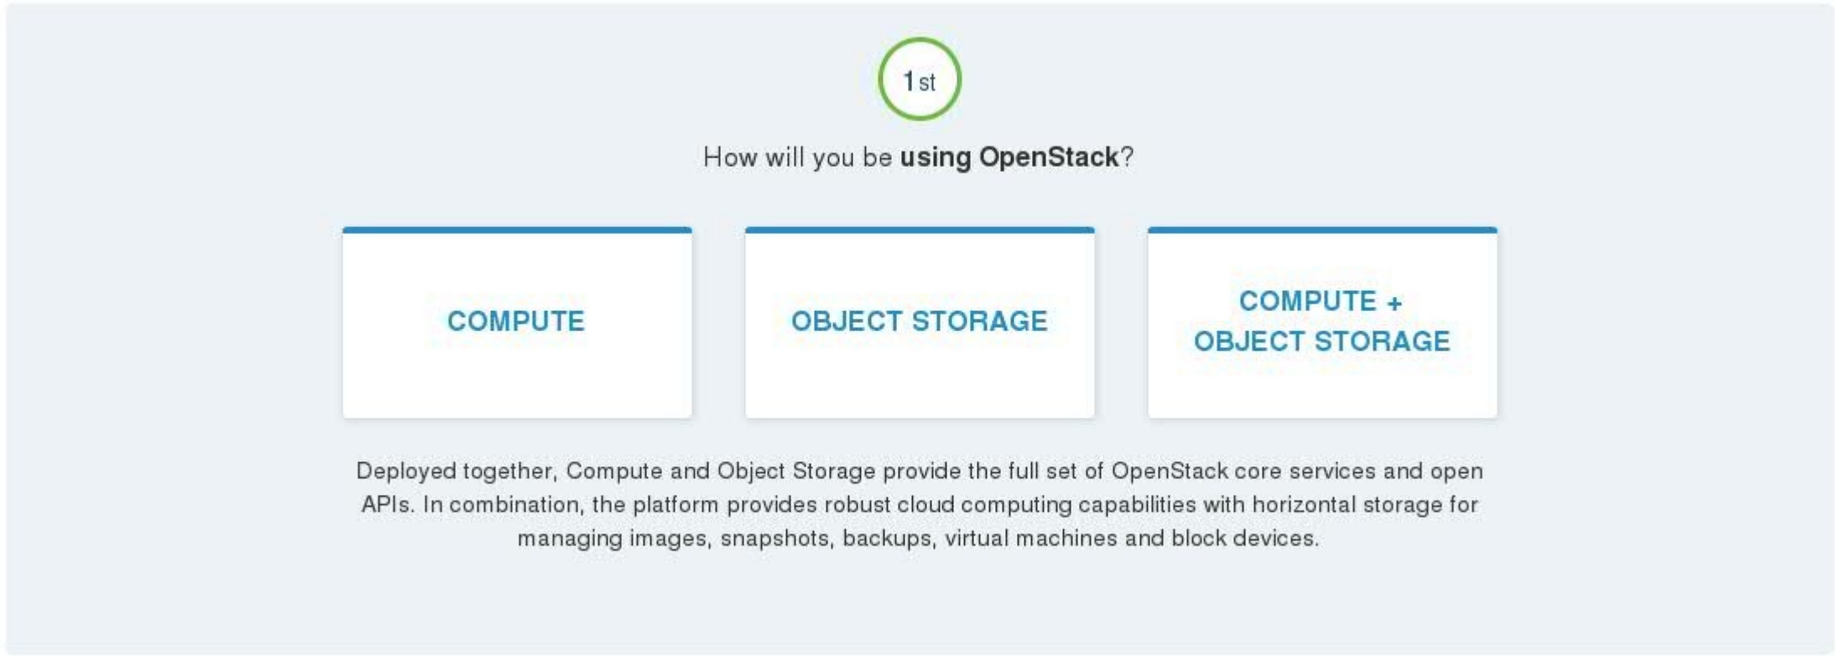

Source: https://www.openstack.org/software/

What is OpenStack?

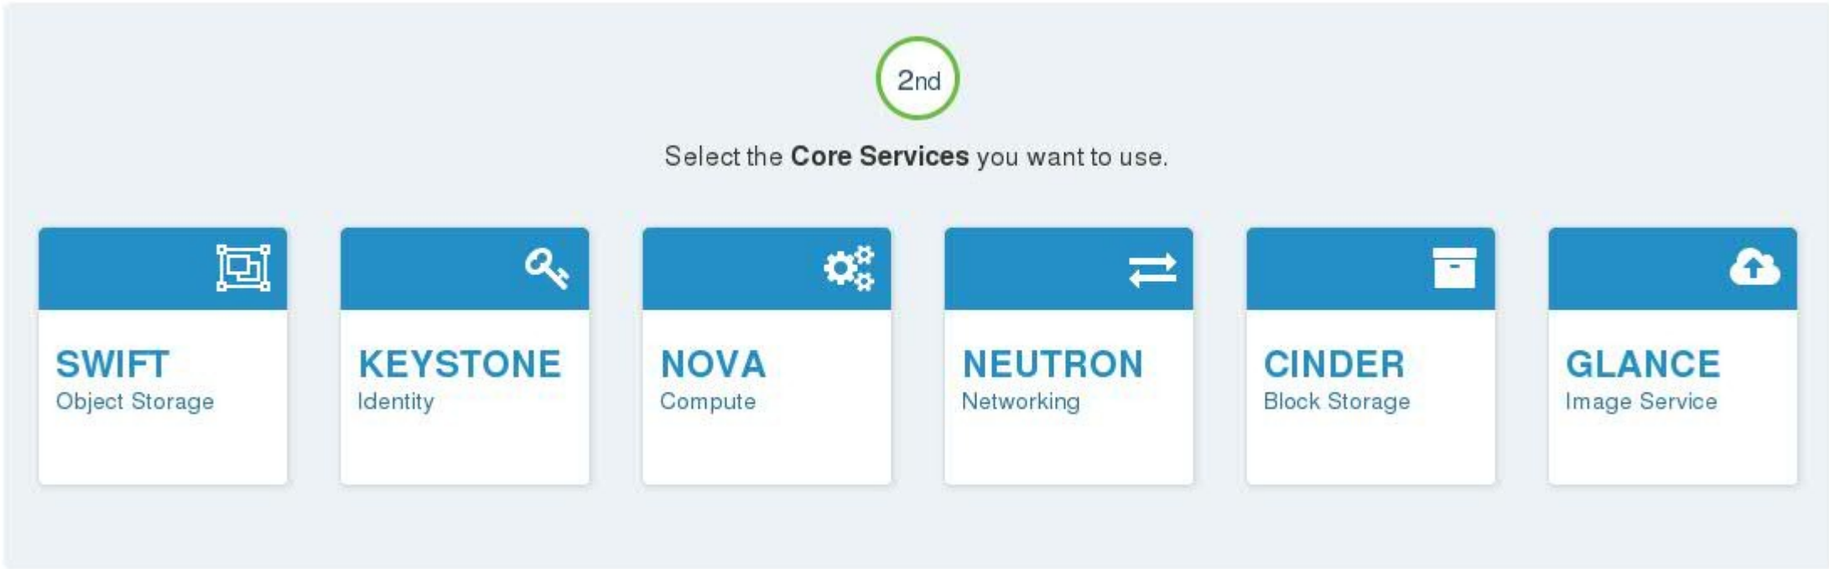

What is OpenStack?

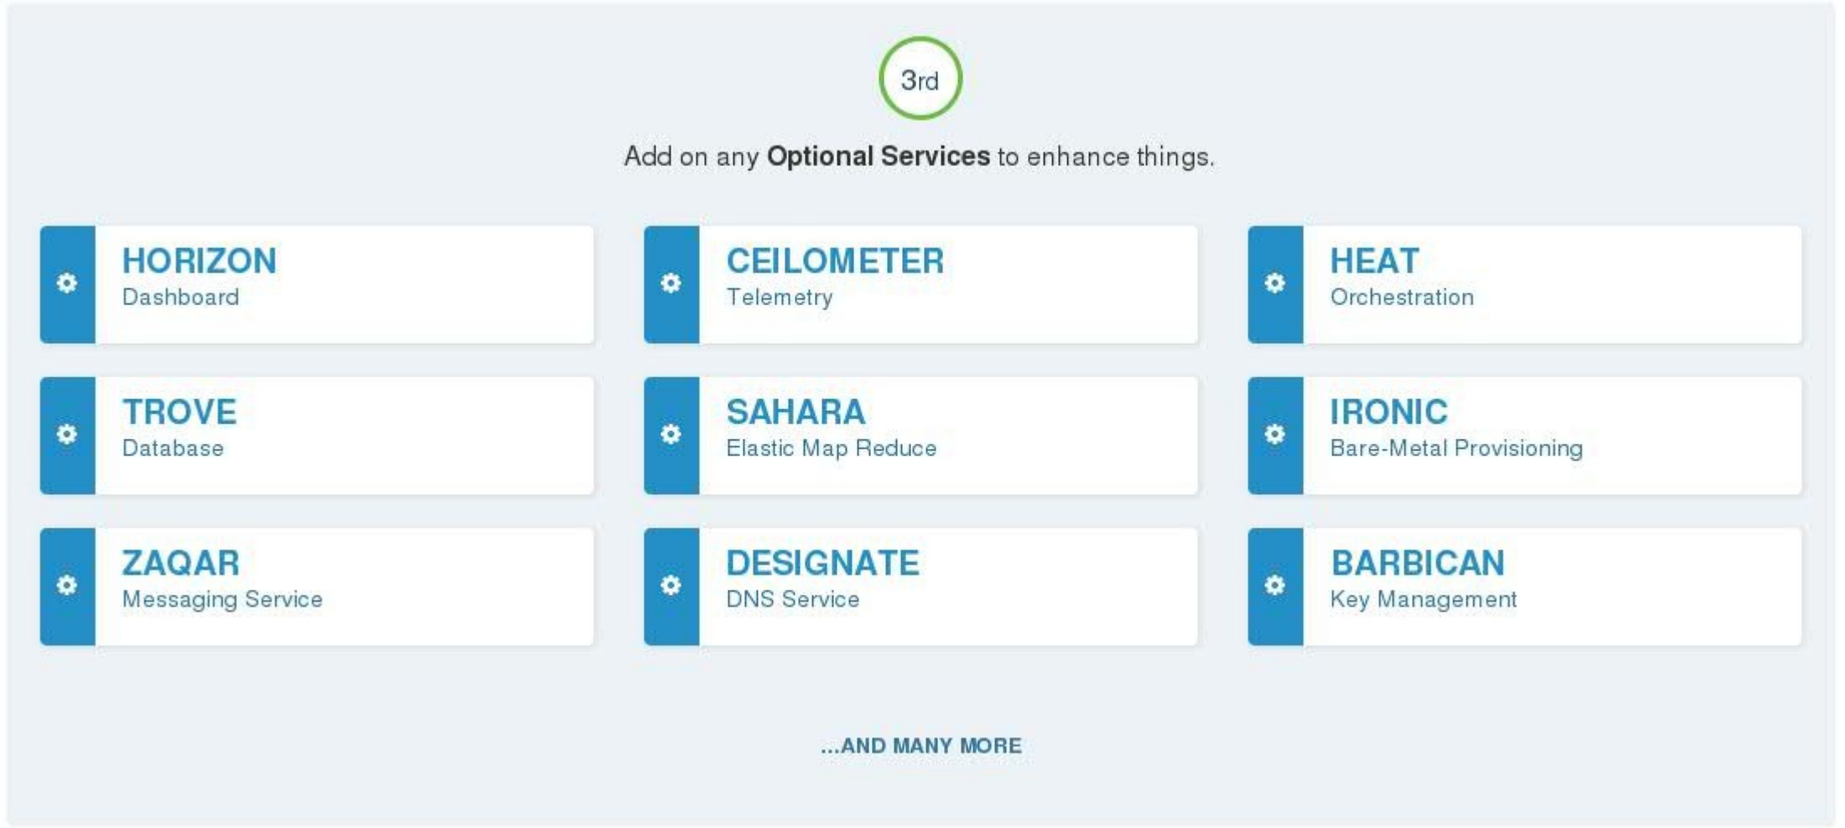

What is OpenStack?

Series	Status	GA Date
Newton	Under development	
Mitaka	Current stable release	Apr 7, 2016
Liberty	Security supported	Oct 15, 2015
Kilo	Security supported	Apr 30, 2015
Juno	EOL	Oct 16, 2014
Icehouse	EOL	Apr 17, 2014
Havana	EOL	Oct 17, 2013
Grizzly	EOL	Apr 4, 2013
Folsom	EOL	Sep 27, 2012
Essex	EOL	Apr 5, 2012
Diablo	EOL	Sep 22, 2011
Cactus	Deprecated	Apr 15, 2011

Source: http://releases.openstack.org/

What is OpenStack?

Most community interaction takes place over IRC (internet relay chat).

Etherpads (a collaborative notepad) and wikis are used for short or long term documentation.

The twice-yearly summit includes sessions for sales/managers, operators and the design summit for technical direction. Summit sessions are decided upon via a combination of voting and core members.

Each OpenStack project has a Project Team Lead (elected position) and several core members. These are people who have done enough code reviews and contributions to be considered experts.

OpenStack Roadmap

	Scalability		Resiliency		Manageability		Modularity		Interoperability	
	Mitaka	Newton	Mitaka	Newton	Mitaka	Newton	Mitaka	Newton	Mitaka	Newton
Ceilometer	X	X	X		X	X		X	X	X
Cinder			X	X	X	X	X			
Glance	X		X		X	X			X	X
Heat	X	X		X	X	X				
Horizon		X			X	X	X	X		X
Ironic	X		X		X	X				
Keystone					X	X		X		
Magnum			X		X	X				
Neutron	X	X	X		X	X			X	X
Nova		X	X	X		X			X	X
Oslo			X		X	X	X	X		X

Source: https://www.openstack.org/software/roadmap/

Current z/VM support

If the controller is on z/VM, then this is manage-from z. If the controller is outside z/VM, then this is manage-to z.

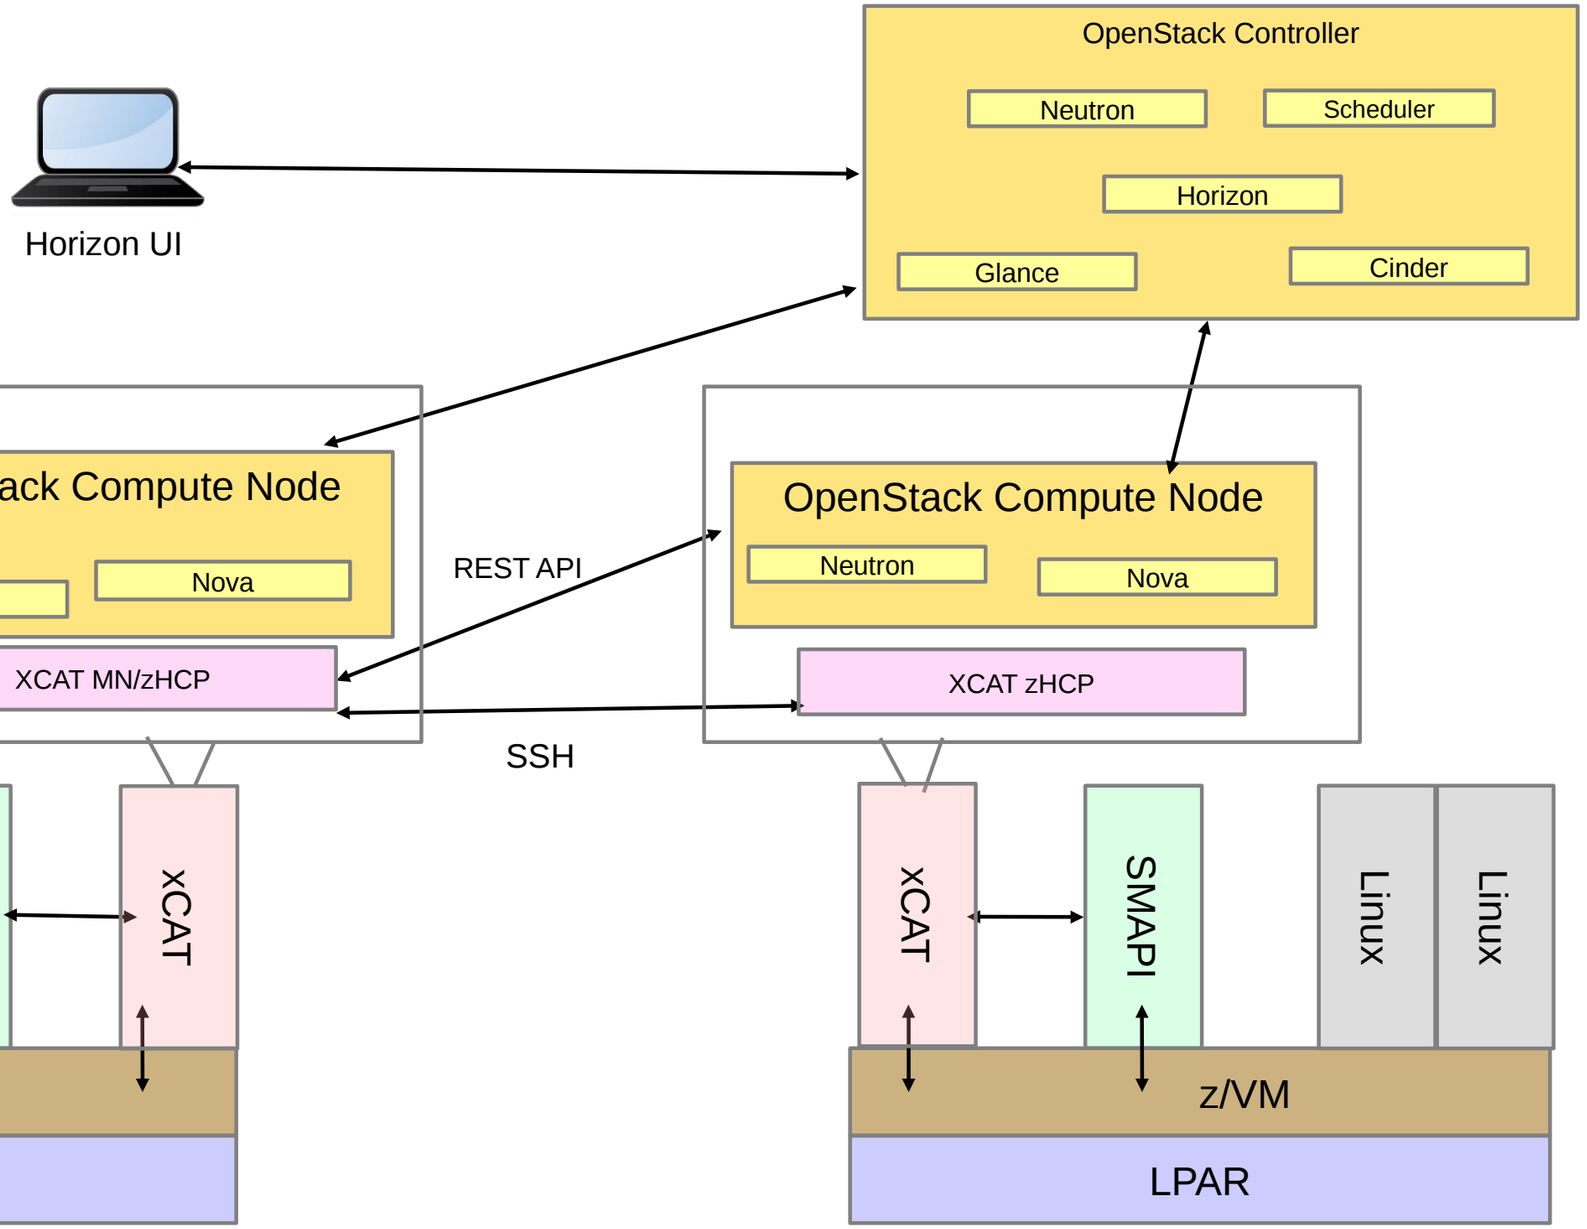

Supported features - nova

- Launch → Image_Definition_Create_DM
- Reboot → Linux reboot, or re-IPL
- Terminate
- Resize → Complete rebuild of the guest
- Pause →PA1
- Un-pause
- Live Migration

- Snapshot
- Fibre Channel
- Set Admin Pass
- Get Guest Info
- Get Host Info
- Glance Integration
- Config Drive
- Discovery of existing guests (driven via xCAT GUI)

Supported features - neutron

- Right now the z/VM agent only supports Layer 2
- VLAN Networking
- Flat Networking

Supported features - cinder

- Attach Volume
- Detach Volume
- Right now support is only for storage in the IBM Storwize family/SVC Fiber Channel Protocol

Latest deliverable

- OpenStack Liberty support, including
  - Ceilometer support
  - RHEL7 and SLES12 provisioning through OpenStack
  - Keystone v3

Latest deliverable

Different CMA modes

- 5 Options
  - Controller – OpenStack controller and compute node and xCAT MN and zHCP
  - Compute – OpenStack compute node and xCAT zHCP
  - Compute_mn – OpenStack compute node and xCAT MN and zHCP (if you have a non-CMA controller)
  - MN – xCAT MN and zHCP
  - ZHCP – xCAT zHCP
- This means a single service stream for xCAT and OpenStack updates
- The ZHCP userid is no longer needed, the whole appliance runs from XCAT

Things to know

- Liberty is the first release of our Cloud Manager Appliance that is not part of the IBM Cloud Manager with OpenStack product.
  - The ICM deployer is not included
  - Cannot manage from z to other platforms
  - Only the xCAT GUI and Horizon GUI are included, evaluate your self-service portal needs to see if an additional OpenStack product is right for you
- The Chef server is not included
  - The Chef client is still included
  - Chef cookbooks are included
    - For configuring the CMA to be managed by an external cross-platform OpenStack controller
    - For configuring the CMA to use an external Keystone server

How do I get OpenStack on z/VM?

- There are many ways to get OpenStack:
  - On-premises distribution: A customer downloads and installs an OpenStack distribution within their internal network. You could create your own using our community plugins, or choose a distribution like SUSE OpenStack Cloud 6
    - SUSE OpenStack Cloud 6 is available here https://www.suse.com/products/suse-openstack-cloud/
    - It uses open source OpenStack (community) Liberty drivers for z/VM and xCAT, which are available to anyone who wants to download them.
    - Only the xCAT MN and zHCP need to run on z/VM with the rest of the code running in an x86 Linux guest. SUSE OpenStack Cloud 6 includes the z/VM installation and configuration, so there are no manual steps!

How do I get OpenStack on z/VM?

- There are many ways to get OpenStack:
  - On-premises distribution: A customer downloads and installs an OpenStack distribution within their internal network. You could create your own using our community plugins, or choose a distribution like SUSE OpenStack Cloud 6
  - Hosted OpenStack Private Cloud: A vendor hosts an OpenStack-based private cloud: including the underlying hardware and the OpenStack software.
  - OpenStack-as-a-Service: A vendor hosts OpenStack management software (without any hardware) as a service. Customers sign up for the service and pair it with their internal servers, storage and networks to get a fully operational private cloud.
  - Appliance based OpenStack: z/VM includes an OpenStack appliance

Sources: https://en.wikipedia.org/wiki/OpenStack
https://www.suse.com/company/press/2015/suse-offers-beta-preview-of-suse-openstack-cloud-6.html

Before you install the z/VM appliance

- To set up the z/VM appliance, you will need some resources for the appliance and some resources for your cloud.
- For the appliance:
  - Disk space for your LVM, this must have enough space for any guest images you want to store. We recommend 50G of disk space.
  - If not already set up, the XCAT userid (where the appliance is installed) should also have a minidisk at address 101 and 102 and the MAINT630 must have minidisks at addresses 102, 103, 104, 105 each of size:
    - ECKD - 3338 cylinders
    - FBA/eDevice - 4806720 blocks
  - 1 externally available IP addresses, associated OSA card/gateway/mask (install will set up a layer 2 vSwitch from this)

Before you install the z/VM appliance

- To set up the z/VM appliance, you will need some resources for the appliance and some resources for your cloud.
- For your cloud:
  - Some disk space for you to install your guests (at least the equivalent of one 3390-9 for one guest)
  - Additional IP addresses on the same subnet as your appliance IP
  - If you want to use cinder to attach/detach additional disks to your guests, some storage connected via an IBM Storwize SAN
    - IP address of your SVC storage
    - Filename of the SAN private key file
    - Storwize SVC pool name
    - Storwize SVC io_group_id

Installing the z/VM appliance – Configure DirMaint and SMAPI

- On a fresh install, choose Yes on this screen

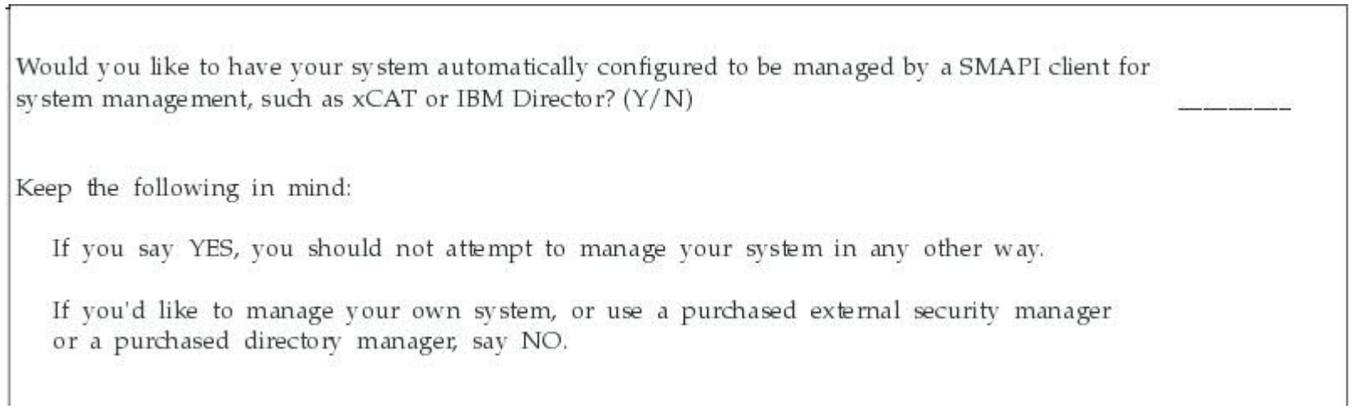

- Or, enable DirMaint and SMAPI on your own, following the instructions in the Directory Maintenance Facility Tailoring and Administration Guide and the Systems Management Application Programming
  - Be sure to follow the instructions in the DirMaint Appendix B to make DirMaint and SMAPI talk

Installing the z/VM appliance – Getting the latest service

- The appliance has two parts
  - One is installed via the latest in z/VM service for CMS (CMA120 FILE) so
    - Be sure that you're up to date on all your service.
    - Check this webpage for the latest service information and manuals http://www.vm.ibm.com/sysman/osmntlvl.html

Installing the z/VM appliance – FixCentral

- The appliance has two parts
  - The other is downloaded from FixCentral, so download this onto your laptop
  - http://www.ibm.com/support/fixcentral

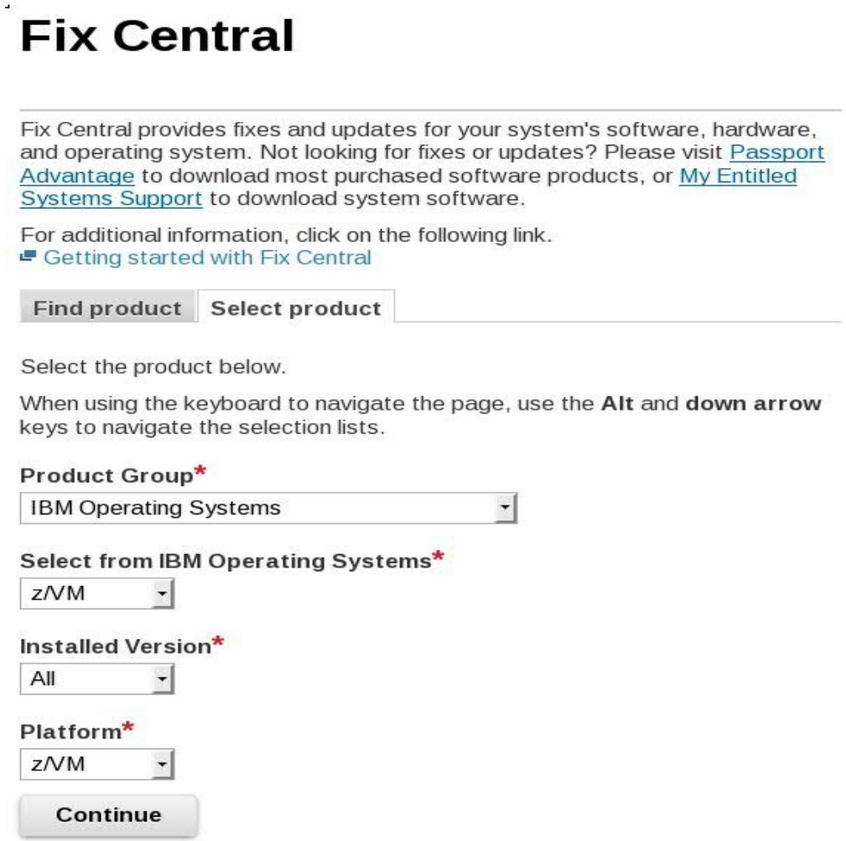

Installing the z/VM appliance - FixCentral

- From Fix Central
  - Depending on the type of DASD you're using choose:
    - CMA101.ECKDPACK and CMA102.ECKDPACK
    - CMA101.FBAPACK and CMA102.FBAPACK

Installing the z/VM appliance – Configuration files

- From MAINT630, configure your DMSSICNF (xCAT) and DMSSICMO (OpenStack) COPY files using these commands
  - LOCALMOD CMS DMSSICxx $COPY
  - SERVICE CMS BUILD
  - PUT2PROD

Installing the z/VM appliance – Configuration files

Table 2. DMSSICNF COPY File, With Hints for Updating

DMSSICNF COPY File	Hints for Updating
/***/	
/* XCAT server defaults */	
/***/	
XCAT_User = "XCAT" /* xCAT z/VM user ID */	
XCAT_Addr = "10.10.10.10" /* XCAT IP Address */	
XCAT_Host = "xcat" /* xCAT hostname */	
XCAT_Domain = ".yourcompany.com" /* xCAT domain name */	Your site, for example: ".ibm.com"
XCAT_vswitch = "XCATVSW1" /* xCAT Vswitch name */	
XCAT_OSAdev = "NONE" /* OSA address for xCAT */	
XCAT_zvmsysid = "zvmnode" /* xCAT z/VM system id */	System name where XCAT is running
XCAT_notify = "OPERATOR" /* Notify when xCAT started */	Notification, as desired
XCAT_gateway = "10.10.10.1" /* Network gateway IP addr. */	
XCAT_netmask = "255.255.255.0" /* Default network mask */	
XCAT_vlan = "NONE"	
XCAT_iso = "volid1 volid2 volid3 volid4 volid5 volid6 volid7 volid8	Volume labels of one or more
volid9 volidA"	disks to hold the xCAT image files
XCAT_MN_Addr = "x.xx.xx.xxx" /* xCAT mgmt node IP address */	IP address assigned to xCAT Management Node
XCAT_MN_vswitch = "XCATVSW2" /* xCAT MN Vswitch name */	
XCAT_MN_OSAdev = "NONE" /* OSA address for xCAT MN */	OSA address to attach to XCATVSW2
XCAT_MN_gateway = "NONE" /* Network gateway IP addr. */	Network gateway, if not *x.xx.xx*.1
XCAT_MN_Mask = "255.255.255.0" /* Netmask for xCAT MN */	Net mask for your network
XCAT_MN_vlan = "NONE"	
XCAT_MN_admin = "mnadmin" /* MN administrator userid */	Userid of xCAT maintenance ID for
	SSH access to XCAT
XCAT_MN_pw = "NOLOG" /* MN admin password */	Password for the XCAT_MN_admin user above.
	If set to the default 'NOLOG', the user
	will not be created.
	Notes:
	- XCAT_MN_pw is a requirement for using
	OpenStack with z/VM
	- If the password is set, it should
	then be changed after logging on to the
	XCAT_MN_admin user (via SSH).
/***/	- For a CMA, a valid password should
/* ZHCP server defaults */	be specified instead of "NOLOG".
/***/	
ZHCP_User = "ZHCP" /* zhcp z/VM user ID */	
ZHCP_Addr = "10.10.10.20" /* zhcp IP ADDRESS */	
ZHCP_Host = "zhcp" /* zhcp hostname */	Your site, for example: ".ibm.com"
ZHCP_Domain = ".yourcompany.com" /* zhcp domain name */	
ZHCP_vswitch = "XCATVSW1" /* zhcp Vswitch name */	
ZHCP_OSAdev = "NONE" /* OSA address for zhcp */	
ZHCP_gateway = "10.10.10.1" /* Network gateway IP addr. */	
ZHCP_netmask = "255.255.255.0" /* Default network mask */	
ZHCP_vlan = "NONE"	

Installing the z/VM appliance – Configuration files

Table 4. DMSSICMO COPY File Defining a Controller Node, With Hints for Updating

DMSSICMO COPY File	Hints for Updating
cmo_admin_password = "yourpassword"	Overall password used in OpenStack configurations(e.g. account, tenant, DB, etc.).
cmo_data_disk = "volid1 volid2 volid3 volid4"	Volume IDs used to configure controller servers.
openstack_default_network = "192.168.1.2-192.168.1.254/24"	IP address range and CIDR used to create the default network and its subnet. Change to your value.
openstack_system_role = "controller"	The role of the CMA. After initial configuration, do not change.
openstack_controller_address = "controller_ip_address"	Controller's IP address that this nova compute node will use.
openstack_zvm_diskpool = "ECKD:xcateckd"	DASD type and DASD pool name to be used when deploying virtual machines.
openstack_instance_name_template = "cmo%05x"	Instance name template. Compute role system should use same template as controller role system.
openstack_zvm_fcp_list = "1d2c"	FCPs used by instances. Contact your z/VM system administrator if you don't know which FCPs you should use.
openstack_zvm_scsi_pool = "NONE"	xCAT SCSI Pool name. Change to your value.
openstack_zvm_zhcp_fcp_list = "NONE"	FCPs used only by xCAT ZHCP node. Obtain the list from your z/VM system administrator.
openstack_san_ip = "NONE"	IP address of your SVC storage. Contact your SVC service manager if you don't know the address.
openstack_san_private_key = "id_rsa"	Filename of private key file. Contact your SVC service manager to get the file.
openstack_storwize_svc_volpool_name = "NONE"	VDISK pool used by cinder. Contact your SVC service manager to get the name of the pool.
openstack_storwize_svc_vol_iogrp = "NONE"	The io_group_id. Contact your SVC service manager to get the file.
openstack_zvm_image_default_password = "NONE"	Password for new instance. Change to your value.
openstack_xcat_mgt_ip = "192.168.2.1"	xCAT's management IP address. Used to communicate with new deployed instances.
openstack_xcat_mgt_mask = "255.255.255.0"	Network mask for xCAT management IP. Change to your value.
openstack_zvm_xcat_master = "xcat"	The xCAT node name for this CMA.
openstack_zvm_vmrelocate_force = "NONE"	The type of relocation to be performed, ask for your system administrator to get this type info.
openstack_zvm_xcat_service_addr = "192.168.2.1"	The xCAT management node IP address that is reachable by all compute nodes and ZHCP.
openstack_volume_enable_multipath = "TRUE"	Used to configure the nova configuration property zvm_multiple_fcp and the cinder configuration property storwize_svc_multipath_enabled.

Installing the z/VM appliance – Appliance installation

- Refer to the CMA120 FILE on the MAINT 400 disk for detailed instructions (this should have been installed/updated by the service you applied earlier).
- From MAINT630, create, link and format the new 102-105 minidisks
  - Access 102 as T
  - Access 103 as U
  - Access 104 as V
  - Access 105 as W

Installing the z/VM appliance – Appliance installation

- FTP the CMA10x.ECKDPACK or CMA10x.FBAPACK files you downloaded from Fix Central over to your MAINT630 101 and 102 disks using
  - BIN
  - QUOTE SITE FIXRECFM 1024
- Unpack this file to the MAINT630 103 and 104 disks using
  - COPYFILE CMA101 filetype T CMA101 filetype V (UNPACK OLDDATE
  - COPYFILE CMA102 filetype U CMA102 filetype W (UNPACK OLDDATE
- SIGNAL SHUTDOWN the XCAT userid and link the 101 and 102 disks
- Restore the image file to the XCAT 101 and 102 disks using DDRREST
  - ACCESS 193 T
  - DDRREST 101 CMA101 filetype V
  - DDRREST 102 CMA102 filetype W

Installing the z/VM appliance – Appliance installation

- DETACH the XCAT 101/102 disks
- Comment out the entry for ZHCP in DMSSISVR NAMES
  - This file lives on the MAINT.193 disk, comment out these lines
    * Node server for xcat
    * :server.ZHCP
    * :type.XCAT
    * :subtype.NODE
- Add OPTION LNKNOPAS to the XCAT id user directory
  - If you have the full DirMaint: “dirm for xcat setoptn add lnknopas”
- Add the XCAT id to your VSMWORK1 AUTHLIST file on VMSYS:VSMWORK1.
  - You can replicate the line for MAINT and change the id to XCAT

Installing the z/VM appliance – Appliance installation

- Start the appliance by
  - FORCE VSMGUARD
  - XAUTOLOG VSMGUARD
  - At this point you can also SIGNAL SHUTDOWN ZHCP WITHIN 60 to shut down that id forever. If you leave it running, don't worry, nothing bad will happen
- You should get a message when the appliance has finished IPLing, such as:
  - [time] * MSG FROM XCAT   : CMA VERSION IS: 1.2.0-20160315
    [time] * MSG FROM XCAT   : XCAT VERSION IS: 2.8.3.9
    [time] * MSG FROM XCAT   : CMA: NO MIGRATION STEPS NEEDED
    [time] * MSG FROM XCAT   : OPENCLOUD IS UP AND RUNNING
    [time] * MSG FROM XCAT   : CMA IS RUNNING IN CONTROLLER MODE
    [time] * MSG FROM XCAT   : XCAT SERVICES RUNNING: MN, ZHCP
    [time] * MSG FROM XCAT   : CMA SERVICES RUNNING: OPENSTACK CONTROLLER, COMPUTE

Installing the z/VM appliance – Appliance installation

- You'll want to log in to the Horizon dashboard as you finish your installation, use the admin id and password you set up in DMSSICMO COPY
- https://x.xx.xx.xxx/dashboard/admin/

The unauthorized message is okay the first time the screen comes up.

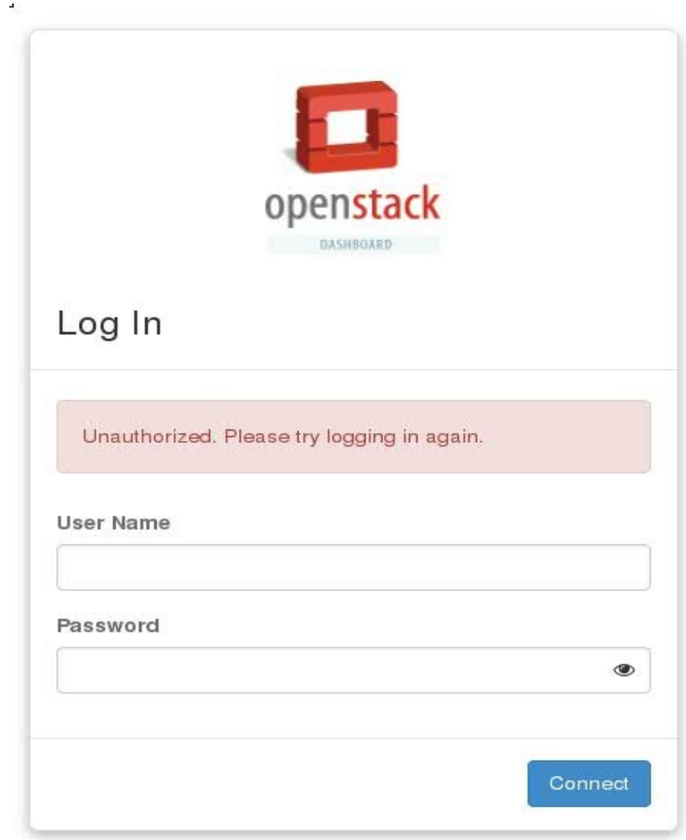

Installing the z/VM appliance – Verification

- Next, run the Installation Verification Program to ensure that your appliance is set up correctly. See Appendix A of Enabling z/VM for OpenStack for complete instructions.
- Run the prep script to get an IVP script that is tailored to your system
  - Log in to your Appliance by SSH/PuTTy as mnadmin
  - Run sudo perl /opt/xcat/share/xcat/tools/zvm/prep_zxcatIVP_LIBERTY.pl

Installing the z/VM appliance – Verification

- Go to https://xcat_mn_addr/xcat and log in as admin
- Go to Nodes->Nodes and select “xcat”

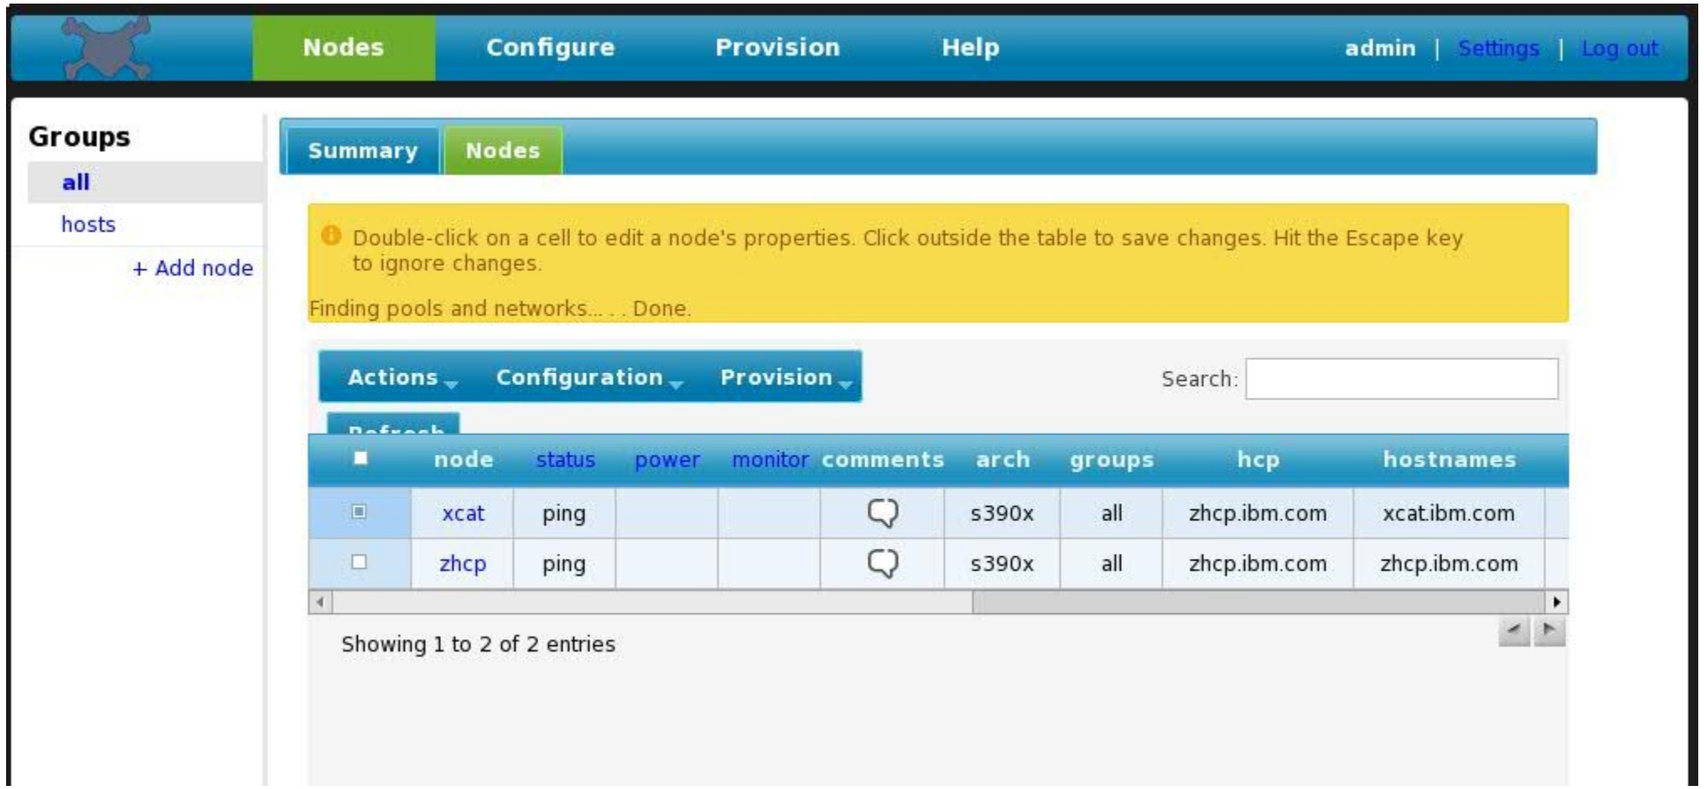

- The select Actions->Run script

Installing the z/VM appliance – Verification

- perl /home/mnadmin/zxcatIVPDriver_9.99.99.199.sh ← where this is the name of the IVP you created on the mnadmin id
- Click Run and note the messages

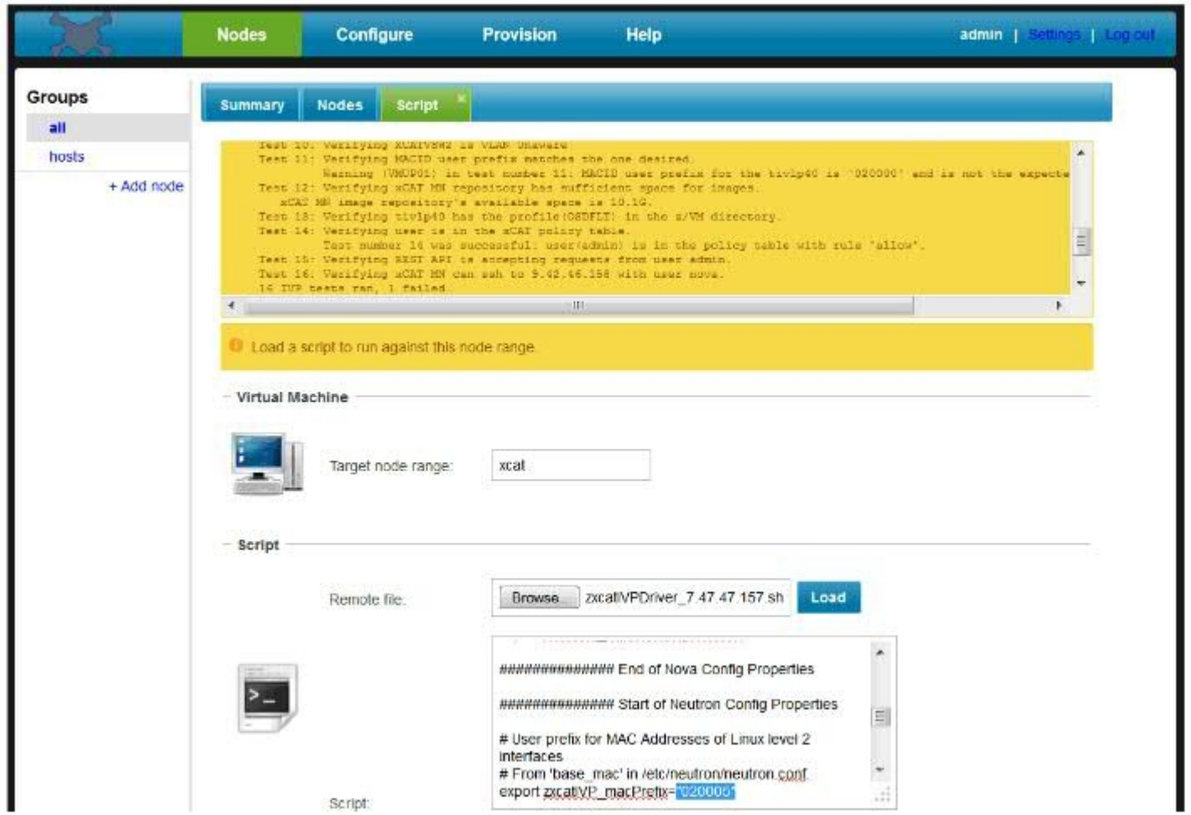

Running OpenStack commands

- If you want to issue OpenStack commands via the commandline, you can do this by SSHing into your appliance as mnadmin.
- Then issue “source openrc”, this sets up your authentication through OpenStack keystone so you can issue commands. IBM supplies this openrc file.

```
[mnadmin@xcat ~] $ source openrc
[mnadmin@xcat ~] $ nova show a1bec88f-45b9-4197-a915-39722dd6cc8d
+--------------------------------------+------------------------------------------------------------+
| Property                             | Value                                                      |
+--------------------------------------+------------------------------------------------------------+
|  network                             |                                                            |
| OS-DCF:diskConfig                    | MANUAL                                                     |
| OS-EXT-AZ:availability_zone          | nova                                                       |
| OS-EXT-SRV-ATTR:host                 | poktst63                                                   |
| OS-EXT-SRV-ATTR:hypervisor_hostname  | POKTST63                                                   |
| OS-EXT-SRV-ATTR:instance_name        | osp00012                                                   |
| OS-EXT-STS:power_state               | 1                                                          |
| OS-EXT-STS:task_state                | -                                                          |
| OS-EXT-STS:vm_state                  | active                                                     |
| OS-SRV-USG:launched_at               | 2016-01-07T10:21:18.000000                                 |
| OS-SRV-USG:terminated_at             | -                                                          |
| accessIPv4                           |                                                            |
| accessIPv6                           |                                                            |
| config_drive                         | True                                                       |
| created                              | 2016-01-07T10:21:18Z                                       |
| flavor                               | m1.tiny (1)                                                |
| hostId                               | cff390ba982119825cd70ffce3688ac8070539e0f308a9c3fd8405ae   |
| id                                   | a1bec88f-45b9-4197-a915-39722dd6cc8d                       |
| image                                | Image not found (305ecab9-a46f-4a0d-8d43-7cfc29b90eb5)     |
| key_name                             | -                                                          |
| metadata                             | {"dsmode": "local"}                                        |
| name                                 | osp00012-gpok189.endicott.ibm.com                          |
| os-extended-volumes:volumes_attached | []                                                         |
| progress                             | 0                                                          |
| status                               | ACTIVE                                                     |
| tenant_id                            | d777abbbd2e841588973968eced882f7                           |
| updated                              | 2016-01-07T16:56:13Z                                       |
| user_id                              | f8bdc4368291485e952d10e1a75c9a2a                           |
+--------------------------------------+------------------------------------------------------------+
```

GUIs supplied with the appliance

- The CMA has two different GUIs associated with it:
- XCAT - https://xcat_mn_addr/xcat
- OpenStack Horizon Dashboard - https://xcat_mn_addr/dashboard/admin/

What now?

- The basic building block in OpenStack is an image (like a Linux .iso file), so you can now follow the instructions for capturing an image in Chapter 6 of Enabling z/VM for OpenStack.

OR

- Use the new discovery function to import your existing guests into OpenStack.
  - Detailed instructions are available in the “Discovering Systems” section of Chapter 4 of Systems Management Application Programming.

Discovery

- Ideally all instances in OpenStack would be created from captured images.
- However, that would mean rebuilding all the guests in your shop, to OpenStack specs
- As a bridge, we provide a function in the CMA that allows you to import existing guests as instances, with a ***limited*** OpenStack functionality.
  - Stop/start and pause/resume
  - Reboot
  - Add disk or networking (with the caveat that any existing disks or networking will not be able to be viewed/created/destroyed by OpenStack)
  - Live Migration
- Anything related to images cannot be done with discovered guests
  - Snapshot
  - Resize

Discovery

- Discovery is started from the xCAT GUI.
- First, you need to make your guests accessible to the CMA, by adding the CMA's key to your guests
- Go to Nodes->Nodes and select the xcat machine of the host system on which you want to discover instances
- Configuration->Unlock

Discovery

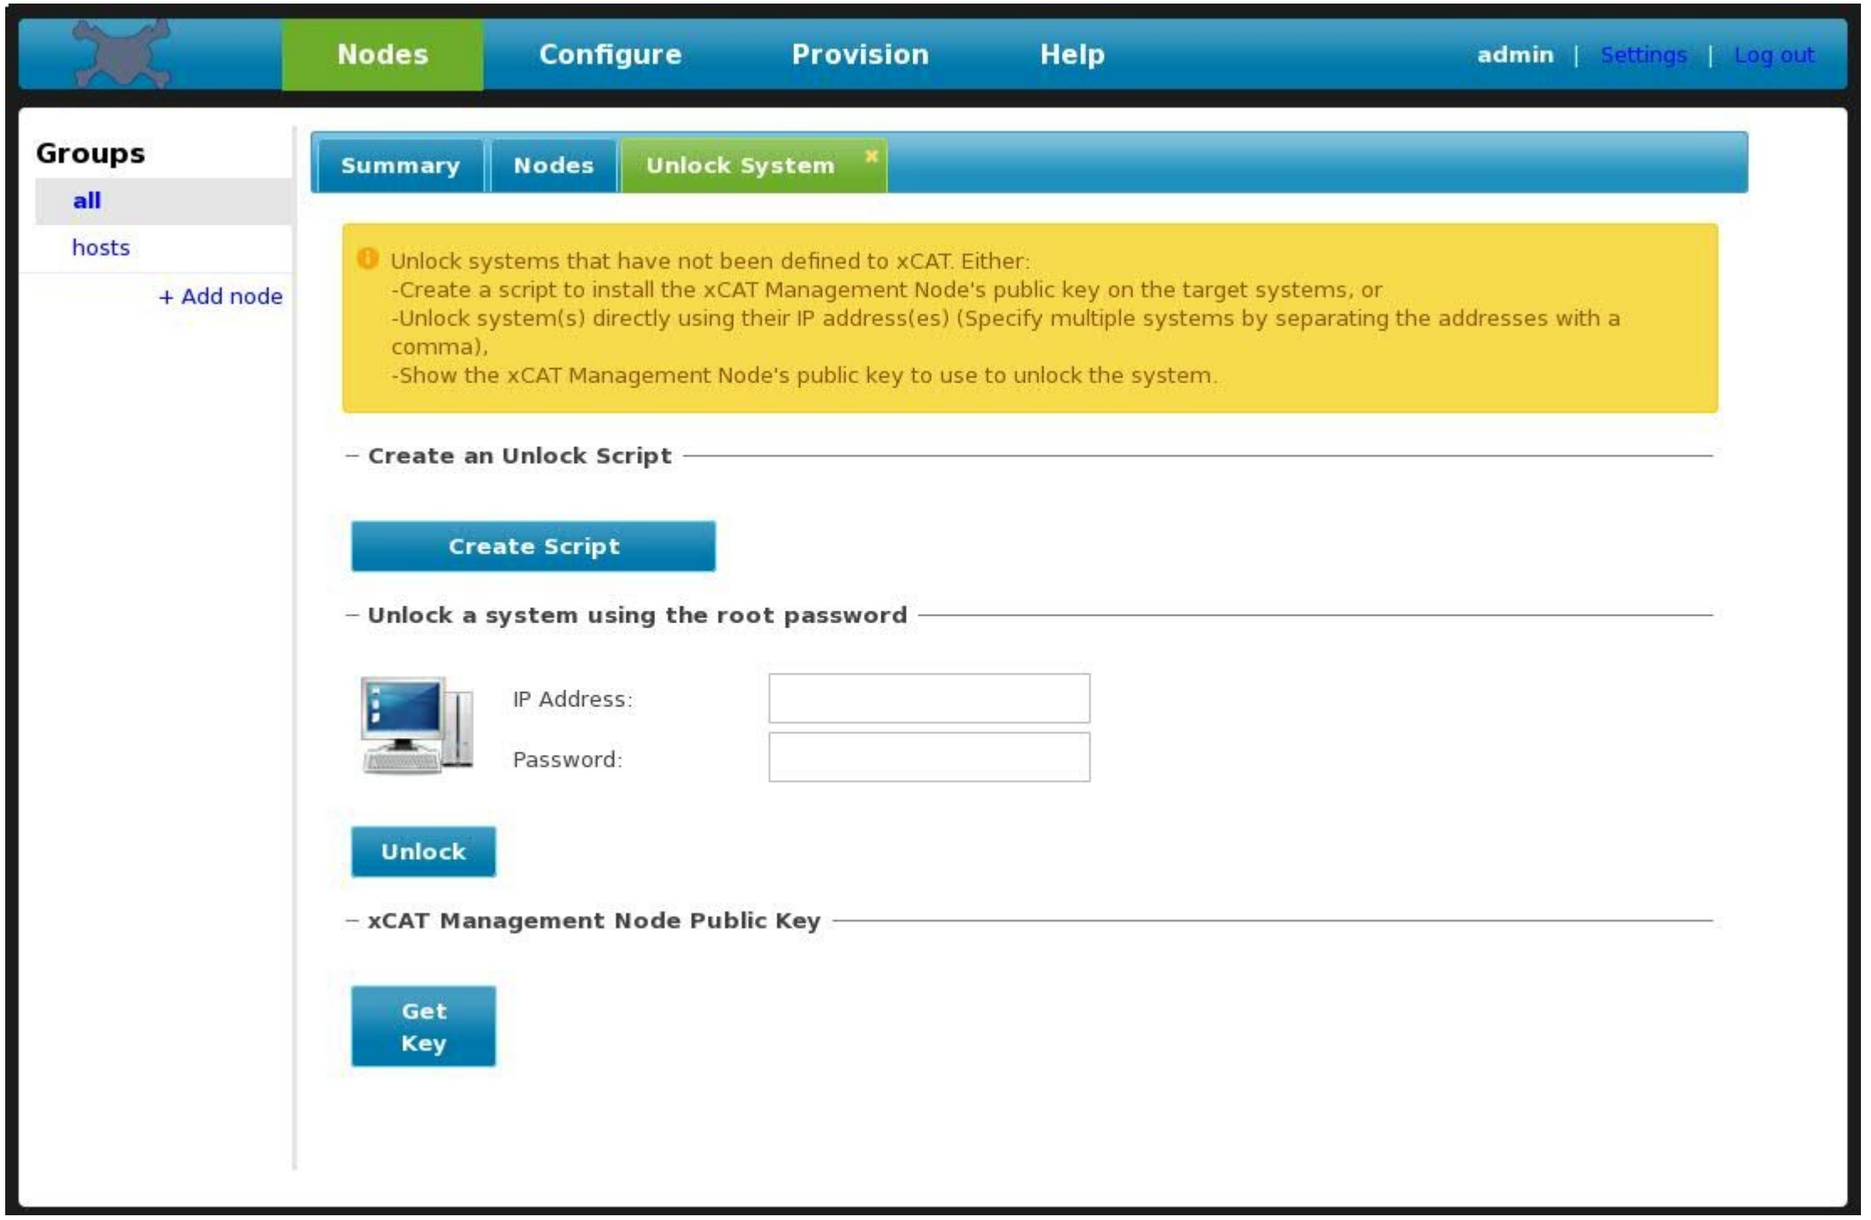

Discovery

- There are three choices for unlock:
  - Create a script that you can then run on each of your to-be-discovered Linux guests
  - Give xCAT the IP address(es) and root password(s) of the guest(s) you want to unlock and xCAT will go in and add the CMA's key to those guests
  - Display the xCAT public key so you can manually add it to your guests
- After a successful unlock, you should see:

```
/usr/bin/ssh setup is complete.
return code = 0
```

Discovery

- Now that your guests are unlocked, you can
- Go to Nodes->Hosts and select the host on which you want to discover instances.

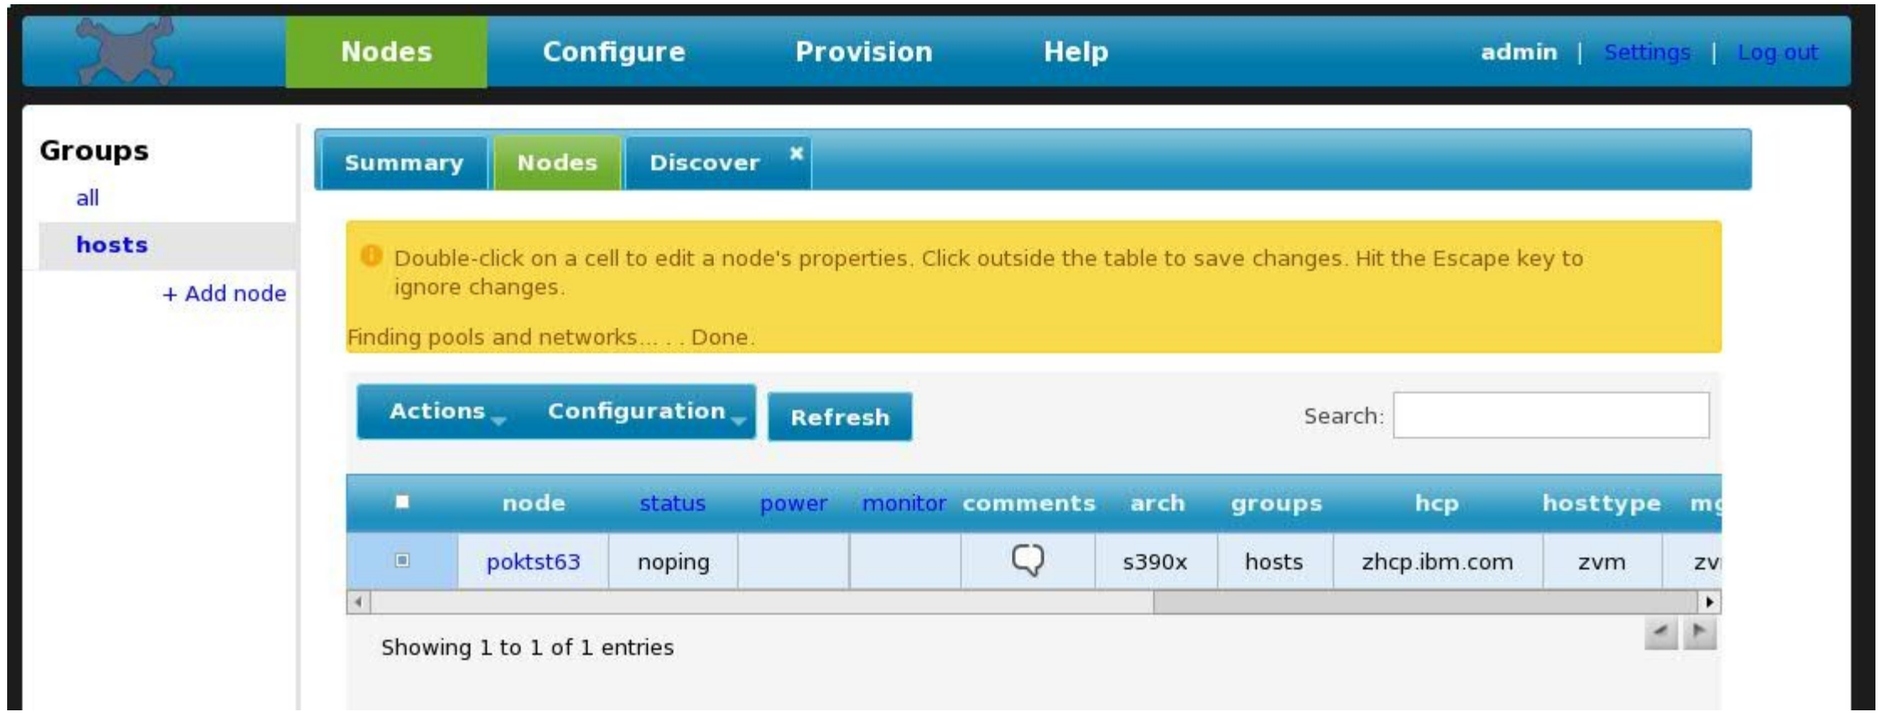

- Then go to Configuration->Discover Systems

Discovery

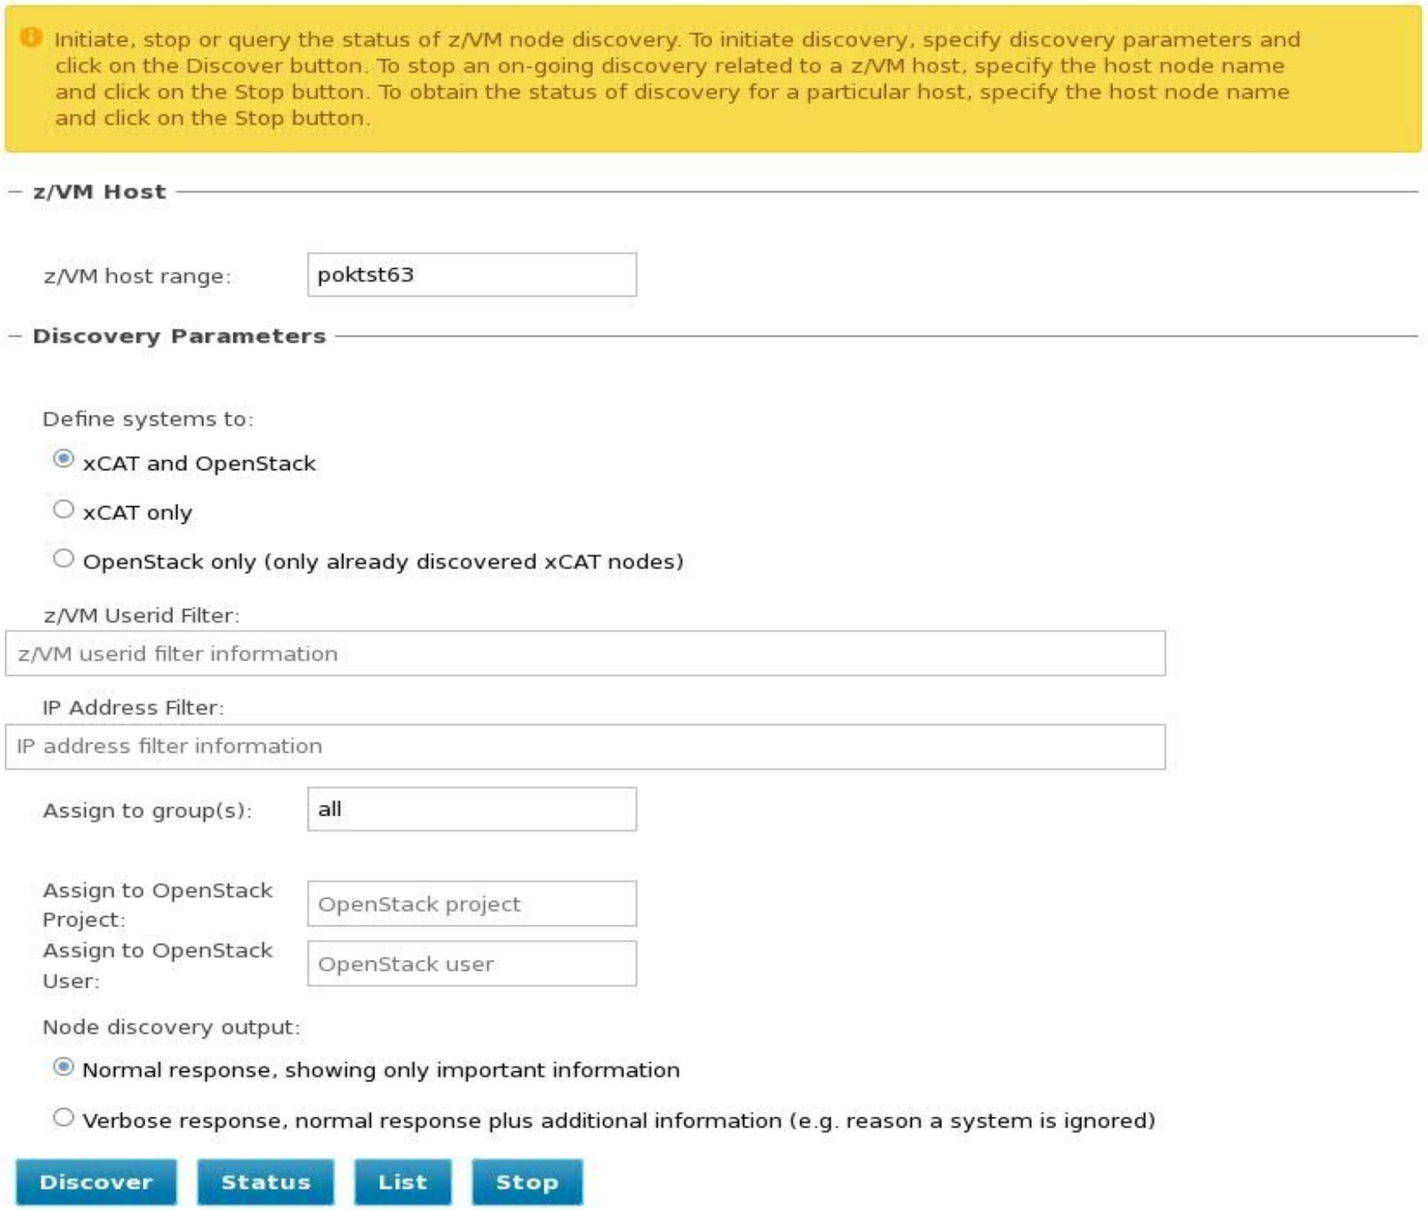

Discovery

- Example of verbose output:

Starting node discovery...

If node discovery is a short running task then its response will follow. If, however, the time it takes to complete discovery exceeds the http request timeout of a few minutes then the discovery response will not be returned to the browser. The status and list buttons can be used to obtained status on the discovery and see what systems have been discovered.
Processing: nodediscoverstart zvmhost=poktst63 defineto=both --verbose groups=all
z/VM discovery started for poktst63
For userid gpok198, 2 adapters were detected.
0700: Not active
0600: 6 MACs with 5 associated IP address(es)
Passing osp00012 to OpenStack for userid gpok198 on z/VM poktst63 with arguments: --memory 2GB --ipaddr 9.60.18.189 --hostname gpok189.endicott.ibm.com --cpucount 2 --guestname osp00012 --os rhel6.5 --verbose 1 --zvmhost poktst63 --uuid 6584fbb4-902e-491a-9037-5e3107e74f9a
Args: ['--memory', '2GB', '--ipaddr', '9.60.18.189', '--hostname', 'gpok189.endicott.ibm.com', '--cpucount', '2', '--guestname', 'osp00012', '--os', 'rhel6.5', '--verbose', '1', '--zvmhost', 'poktst63', '--uuid', '6584fbb4-902e-491a-9037-5e3107e74f9a']
Host IP addr: 9.60.18.197
Admin_token: f98238b7424bd0e9cfb9
Endpoint: http://9.60.18.197:35357/v2.0
No discovery project found, creating new.
No discovery user found, creating new.
Project id: d777abbbd2e841588973968eced882f7
User id: f8bdc4368291485e952d10e1a75c9a2a
Date/time used: 2016-01-07 10:21:18
Short hostname: gpok189
Memory in MB: 2048
Old guest name: osp00012
New guest name: osp00012
Instance UUID: a1bec88f-45b9-4197-a915-39722dd6cc8d
Request ID: req-89d17017-d4e0-407a-88dc-7106d80a373d
Reservation ID: r-jus10sse
Res1 ID: 6ec479d3-44b3-4a06-9ac7-73e2d567d59b
Res2 ID: e9407bc0-894d-40c8-8034-3d8bc462a577
Res3 ID: 5021cb4b-a8e9-4496-805b-ead0201ea0e6
Hypervisor Stats: {, {"num_task_None": 2, "io_workload": 0, "num_instances": 2, "num_vm_stopped": 1, "num_os_type_None": 2, "num_proj_e99aaeb16ebc43b7af314c5e3f284eef": 2}
Node created: osp00012
z/VM discovery is being stopped for poktst63.
Discovered 1 nodes running on poktst63.
NODE z/VM USERID
osp00012 gpok198

Discovery - results

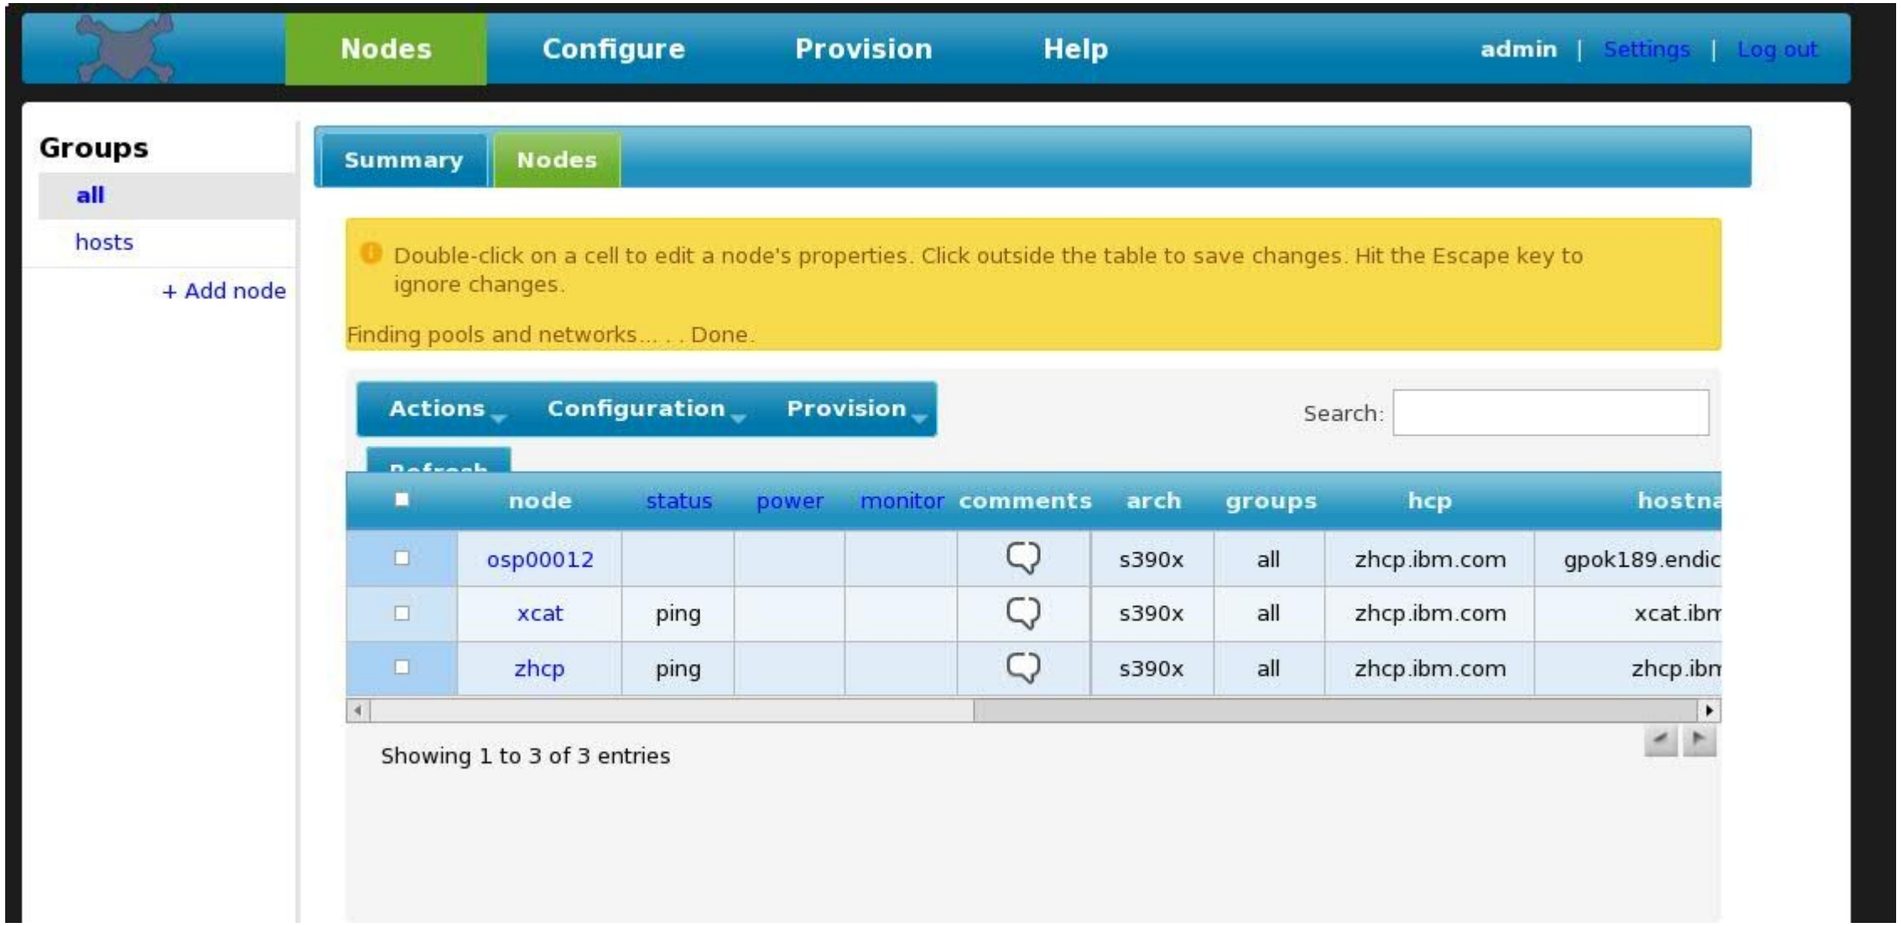

Discovery - results

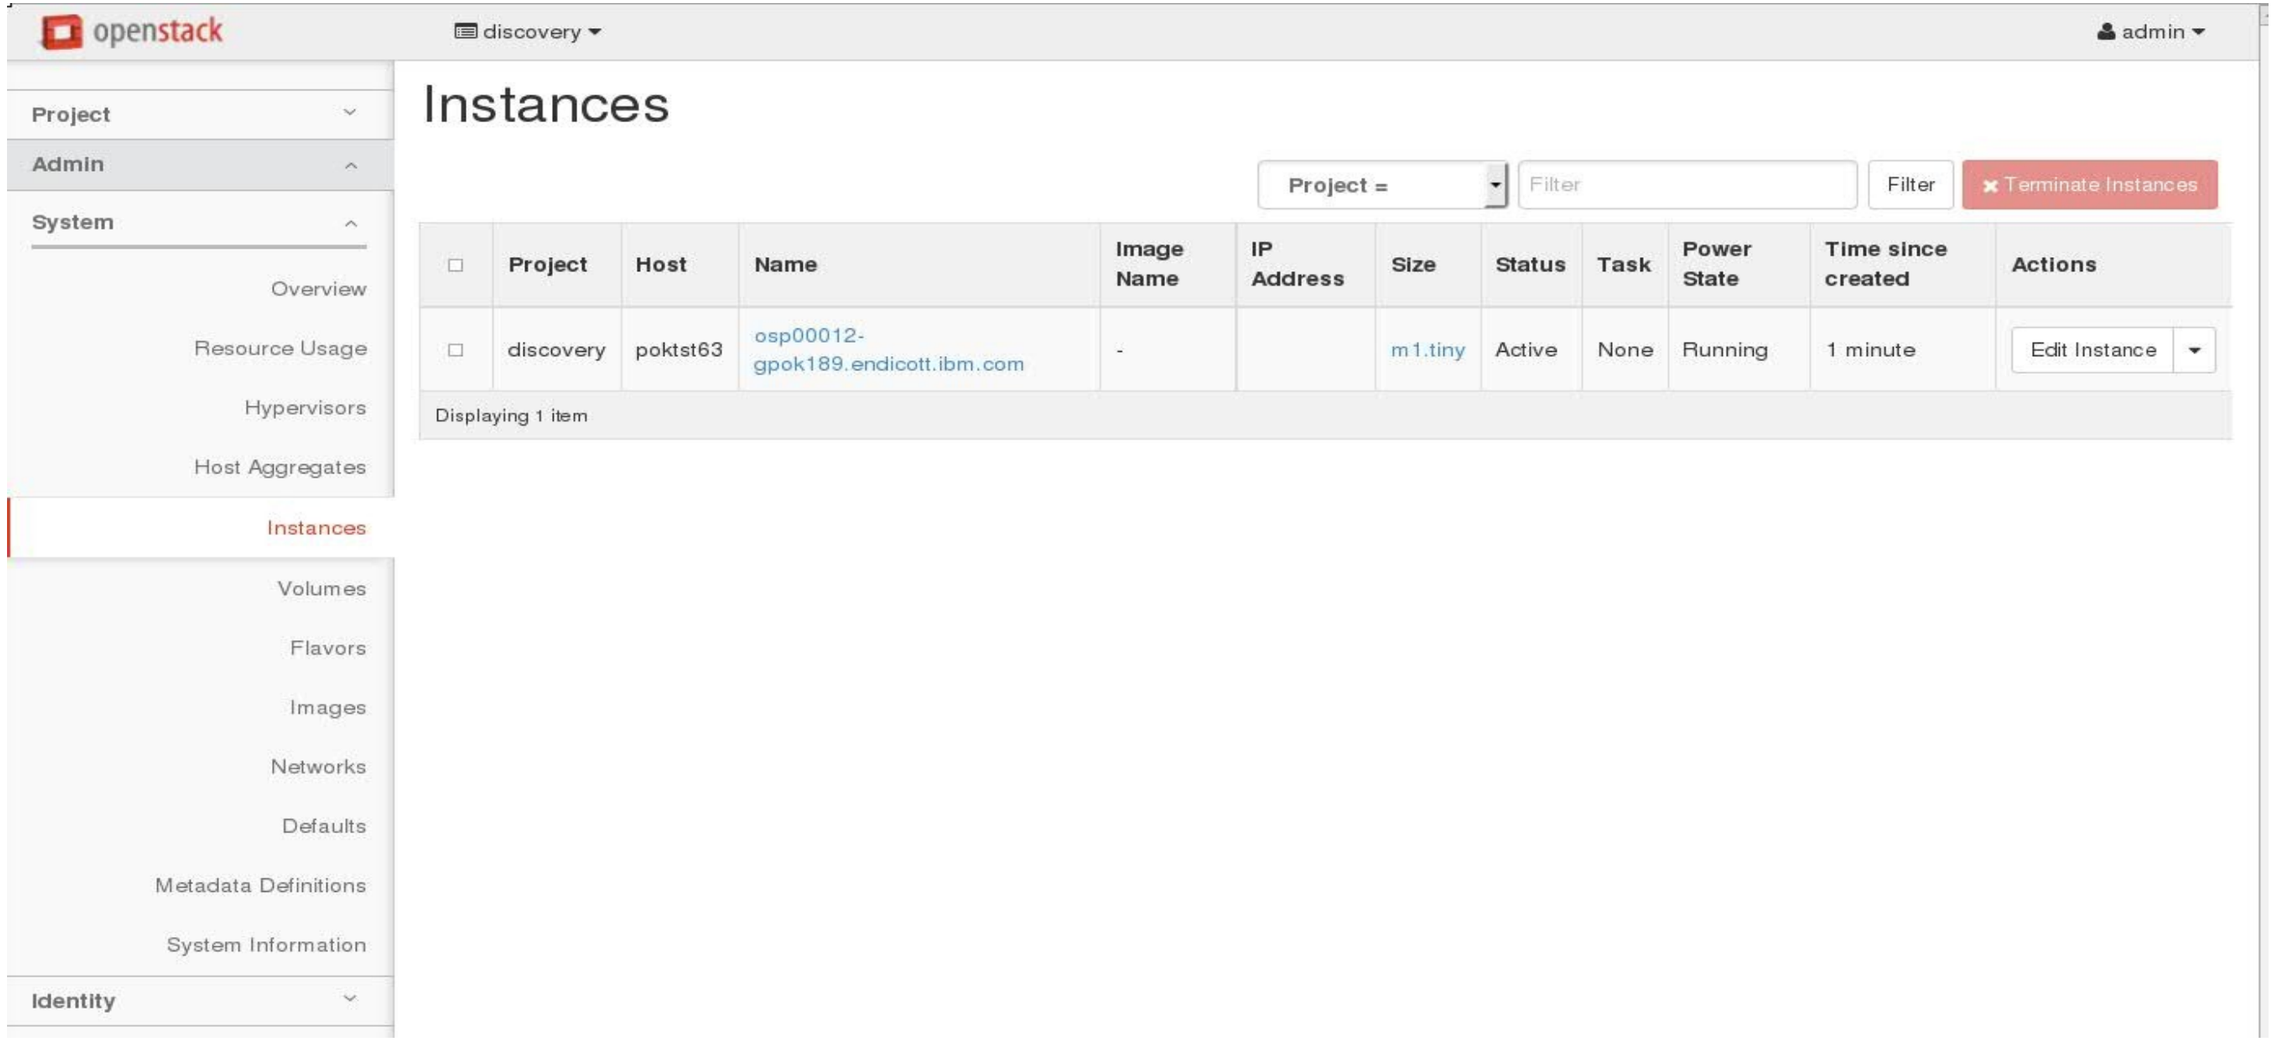

Discovery - results

- Discovered guests must:
  - Be logged on
  - Be Linux guests of a supported distribution (RHEL7 or SLES12, currently)
    - Unsupported distributions are discoverable, and some basic functions (power off/on) will work, but other functions may not
  - Have an IPv4 interface
  - Be accessible from the xCAT MN they're being discovered to
  - Be running on a hypervisor associated with a CMA controller/xCAT MN they're being discovered to

Things to know

- If you're migrating from CMA Juno to CMA Liberty, there are migration scripts to help you. See the CMA120 FILE on the MAINT.400 disk for more information.
- If you're migrating an xCAT only installation to CMA xCAT, there are also migration scripts to help you, see Appendix K of the SMAPI book for more information.

What's next

- CMA Newton will be our next major release
  - Expect fixpacks for our Liberty support between now and then
  - Our drivers have a Mitaka level, that we'll support as the OpenStack Mitaka release is supported
- We're working on a "continuous integration" system that will run z/VM driver tests on every patch in the community
  - So far it's a few x86 systems pointing jobs back to 1 z/VM LPAR

What's next

- We're working on a "continuous integration" system that will run z/VM driver tests on every patch in the community
  - So far it's a few x86 systems pointing jobs back to 1 z/VM LPAR
  - It takes about 2 hours 40 minutes to run each full set of OpenStack Tempest tests
  - We have to respond to each patch in Nova within 4 hours.

How can you help?

- Non x86 hypervisors don't have much visibility in the community
- Install/start playing around with OpenStack
  - On the mainframe (see my next session), or on your own
  - Get involved with the community
  - https://wiki.openstack.org/wiki/Getting_Started
- Interact with our drivers
  - Submit fixes, open bugs, etc
  - https://github.com/openstack/nova-zvm-virt-driver
  - https://github.com/openstack/networking-zvm
  - https://github.com/openstack/ceilometer-zvm
- Help us find the right balance between improving the appliance and improving our drivers

How can you help?

- Let us know about your experiences with OpenStack
  - If you can, please consider writing a blog post about your experiences
  - We need greater visibility within the OpenStack community

Conclusion

- OpenStack is a new way of managing resources in a cloud environment. It's not just a GUI for z/VM.
- It's backed by a vibrant, but x86/KVM-centric, community.
- z/VM has OpenStack drivers in the community and needs your help to make them thrive
- z/VM includes an OpenStack appliance to get you started with using OpenStack
- The appliance allows us to put in additional features not available in the community
  - Easier install
  - Discovery
- We need your feedback as to what's important and what can make OpenStack succeed in your shop!

Thanks!

Emily Hugenbruch
IBM
z/VM
Endicott, NY

 **@ekhugen**

Resources

- Blog posts from John Arwe (one of our team leads)
  https://www.ibm.com/developerworks/community/blogs/looselycoupled?tags=openstack&lang=en
- IBM Cloud Manager Appliance Information page
  http://www.vm.ibm.com/sysman/osmntlvl.html
- Blog posts from Emily Hugenbruch (more coming)
  https://developer.ibm.com/opentech/author/ekhugenbruch/
- See how this looks from your end user developer perspective
  http://www-03.ibm.com/systems/linuxone/
- Liberty Announcement
  http://mainframeinsights.com/zvm-key-cloud-infrastructure-component-open-stack-enablement/