



KVM for IBM z Systems and LinuxONE

Tony Gargya (gargya@de.ibm.com)
Lead Architect, KVM for IBM z

Boeblingen Development Lab
Germany



Trademarks

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

BigInsights	DFSMSdss	FICON*	IMS	RACF*	System z10*	zEnterprise*
BlueMix	DFSMSshm	GDPS*	Language Environment*	Rational*	Tivoli*	z/OS*
CICS*	DFSORT	HyperSwap	MQSeries*	Redbooks*	UrbanCode	zSecure
COGNOS*	DS6000*	IBM*	Parallel Sysplex*	REXX	WebSphere*	z Systems
DB2*	DS8000*	IBM (logo)*	PartnerWorld*	SmartCloud*	z13	z/VM*
DFSMSdfp						

* 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 [OpenStackwebsite](#).

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.



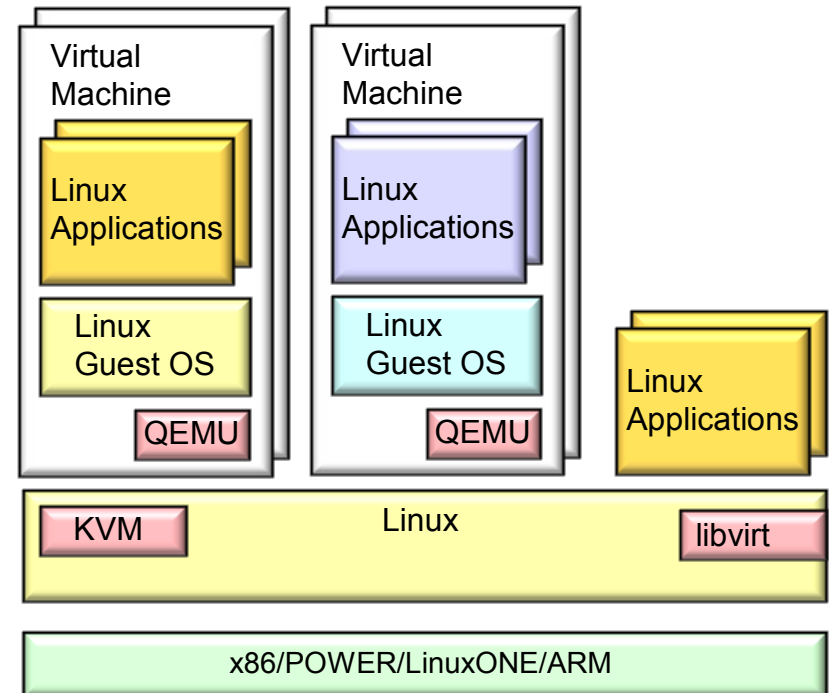
Agenda

- What is KVM
- What is KVM for IBM z Systems and LinuxONE
 - What is in 1.1.1
- What is new in 1.1.2
- What could come thereafter
- Existing Systems Management Tooling



Kernel Based Virtual Machine (KVM)

- **An open source hypervisor based on Linux**
 - Linux provides the base capabilities
 - KVM turns Linux into a hypervisor
 - QEMU provides I/O device virtualization and emulation
- **Provides flexibility in technology choices**
 - Open
 - Scalable
 - Economical





LinuxONE Virtualization Options



LinuxONE has three strategic virtualization platforms

- KVM for LinuxONE
- IBM z/VM
- IBM Processor Resource/System Manager (PR/SM)



KVM for LinuxONE provides an open source choice for LinuxONE virtualization for Linux workloads. Best for clients that are not familiar with z/VM and are Linux centric admins.

z/VM

Proprietary Server Virtualization that is deeply integrated into System z. Complete hardware awareness. Supported on all IBM z Systems and LinuxONE servers. z/VM will continue to be enhanced to support Linux Workloads.

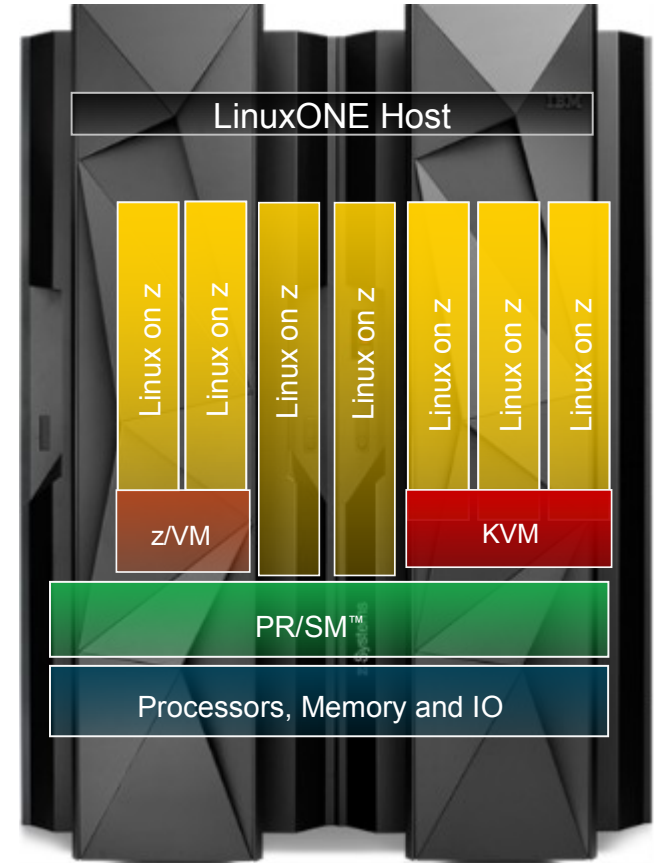
PR/SM

Divide one physical server into up to 85 logical partitions (LPAR) running a mix of multiple z/VM, Linux and KVM for LinuxONE instances isolated and secured in parallel. Share resources across LPARs or dedicated to a particular LPAR. Instances are isolated and secured.



Standards based virtualization with KVM for LinuxONE

- **Standard** Linux management and operation controls leading to greater operational efficiencies
- **Standard** KVM interfaces allow for quick startup for clients who are familiar with x86 Linux
- KVM-based virtualization LinuxONE allows businesses to reduce costs by deploying fewer systems to run more workloads, sharing resources, and improving service levels to meet demand
- KVM open source solution for running virtual servers on LinuxONE enables cloud deployments and big data solutions while reducing complexity and cost



A new hypervisor choice for z Systems



KVM Positioning

Target Customers for **z/VM** : *Linux Clients that ...*

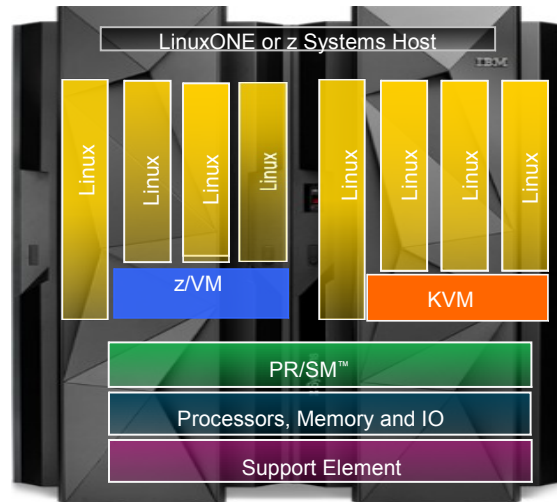
- Already use z/VM for Linux workloads
- Skilled in z/VM and prefer proprietary model
- Invested in tooling for z/VM environment
- Require technical capabilities in z/VM (e.g. I/O pass-through, HiperSockets™, SMC-R, ...)

Target Customers for **KVM** : *New Linux Clients that ...*

- Sold on Open Technologies, Open Source Oriented
- x86 centric – familiar with KVM
- Linux admin skills
- Need to integrate into a distributed Cloud environment, using standard interfaces

z/VM

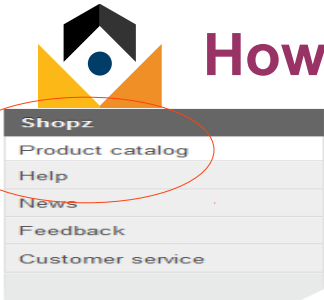
- World class quality, security, reliability - powerful and versatile
- Extreme scalability creates cost savings opportunities
- Exploitation of advanced technologies, such as:
 - 📁 Shared memory (Linux kernel, executable, communications)
- Highly granular control over various resources
- Improves productivity by hosting non-Linux workloads such as z/OS, z/VSE, and z/TPF



KVM

- Standard skills, terminology and technology makes configuration and operation of server virtualization easier/faster
- Leverage common Linux administration skills to administer virtualization
- Flexibility and agility leveraging the Open Source community
- Provides an Open Source virtualization choice

How to order the product / obtain Fixpacks



Product catalog

NOTE: You must login to Shopz and click on the 'My downloads' page to display Standalone product fixes.

Catalog view (Products in this view: 1)

Country/Region:

Package type:

Group:

Language:


Search field: Product ID Product description

Search for:

Sort by: Product ID Product description

KVM for IBM z: Operating Systems

	Product	Description	Version	Language	Notes
◆	[5648-KVM]	KVM for IBM z	1.01.01	English (US)	



SELECT FIXES

Virtualization software, KVM for IBM z Systems (1.1.1.3, All platforms)

Select fixes

The following results match your request. Select the fixes you want to download.

[Share this download list](#)

- To try a different query, go to the [Identify fixes](#) page.

[Show fix details](#) | [Hide fix details](#)

1-1 of 1 results

- 1. **fix pack:** [KVMIBM-1.1.1.4-20160614-s390x](#) → Jun 17, 2016
KVM for IBM z Systems 1.1.1.4 Updates
Platforms:
Applies to versions: 1.1.1, 1.1.1.1, 1.1.1.2, 1.1.1.3
Upgrades to: 1.1.1.4
Severity:
Categories:
Abstract:
Restrictions: The fix pack contains cumulative fixes.
license
[README](#) [RESTRICTIONS](#)

Related links
- Go to Fix Central mobile
- Fix Level Recommendation Tool (FLRT)

Change your selection

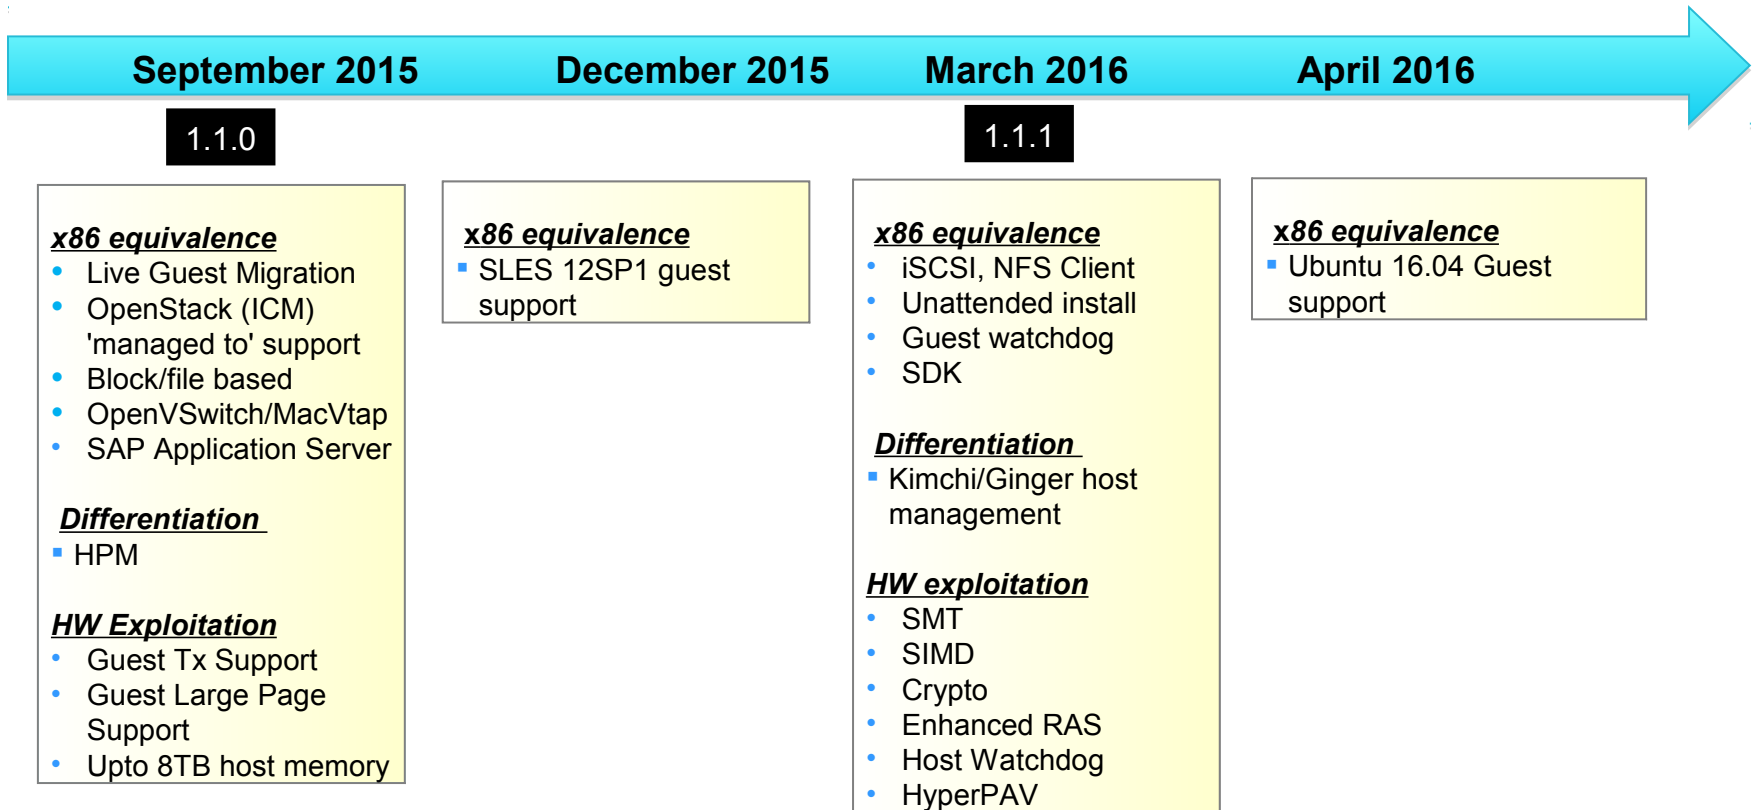
Product selector

Installed Version

Platform



KVM for IBM z Functionality – looking back





KVM for IBM z Systems (1.1.2) supports

Servers	IBM z13™ IBM z13s™ IBM LinuxONE Rockhopper™ IBM LinuxONE Emperor™ IBM zEnterprise® zEC12 IBM zEnterprise® zBC12
Guest Operating systems supported	SUSE Linux Enterprise Server (SLES12 SP1) Ubuntu 16.04 LTS
Networking features supported (NICs)	IBM OSA-Express5S IBM OSA-Express4S
Crypto Coprocessor supported	Crypto Express4S Crypto Express5S
Storage devices are supported	ECKD™ DASD <ul style="list-style-type: none">▪ DS8870/DS8880® (FICON®-attached) FCP SCSI disks: <ul style="list-style-type: none">▪ DS8870/DS8880® (FCP - attached)▪ XIV®▪ Storwize® V7000, V5000, V3700, V3500, SAN Volume Controller▪ FlashSystem™▪ iSCSI▪ NAS

Note: Refer to the KVM for IBM z Systems: Planning and Installation Guide (SC27-8236-02) for the most current information



KVM for LinuxONE

Optimized for z Systems and LinuxONE architecture

Multiple KVM hypervisors on one box

Coexists with z/VM virtualization environments, Linux on IBM z, z/OS, z/VSE, z/TPF

Multi-tenancy

Industry standard KVM hypervisor enables single cross-platform virtualization to help simplify systems management

Eliminate downtime by dynamically modifying I/O device configuration for virtual servers so business applications remain active

Live virtual server workload migration for minimal impact to your business while workloads are relocated

Save on storage cost with copy-on-write virtual disks by not needing full disks until used

Policy-based goal-oriented monitoring and management (HPM) of virtual server CPU resources so critical workloads receive priority

Memory and CPU overcommit to achieve higher VM density per host



KVM for LinuxONE



Tx Support

Security

- SELinux
- firewalld
- Auditd

Availability Monitoring

- Nagios plugins
- ELK: Elasticsearch, Logstash and Kibana (remote)

Performance monitoring

- ELK (remote)
- Sysstat Sar and kSar(remote)

Guest I/O

- Block-based and File-based (raw, qcow2)
- Networking Virtualization via OpenVSwitch and MacVTap



KVM for LinuxONE



Support new analytics workloads with Single Instruction Multiple Data (SIMD) for competitive advantage

Deliver higher compute capacity with support for Simultaneous Multithreading (SMT) to meet new business requirements

Secure and protect business data with Crypto exploitation that leverages hardware acceleration for cryptographic functions – increases entropy

Provide clients with choices for flexibility based upon their storage environment

- Connect a variety of peripherals, especially storage devices drives, with Internet Small Computer System Interface (iSCSI)
- Access files on remote hosts exactly the same way a user would access any local files with Network File System (NFS) which works across a variety of server and host architectures



KVM for LinuxONE

Unattended installation of the KVM hypervisor simplifies administration

Supported by DPM simplified interface for platform management. KVM for IBM z is the only supported hypervisor. Modify system resources without disrupting running workloads

Upgrade tool easy upgrade from 1.1.1 to 1.1.2

Software Development Kit (SDK) enables clients and IHV/ISVs to extend the hypervisor with their applications. For example they can build management agents which need to run on the hypervisor.

Single Hypervisor Management GUI (Kimchi) manage attached storage and networks without deep system z knowledge





Kimchi UI

Wok 9:15Z 151.63 root Help

Dashboard Network **Storage** System Modules Administration

- Volume Group
- File Systems
- Swap Devices
- Storage Devices
 - Actions
 - Format ECKD
 - Remove
 - Refresh Add
 - Filter

	Multipath Links	Type	Size (MB)
<input type="checkbox"/> IBM.750000000DX111.0002.81	3	dasd	23034
<input type="checkbox"/> IBM.750000000DKRP1.0001.2a	3	dasd	2347
<input checked="" type="checkbox"/> IBM.750000000DX111.0001.2a	3	dasd	2347
<input type="checkbox"/> IBM.750000000DX111.0001.29	3	dasd	unknown
<input type="checkbox"/> IBM.750000000DKRP1.0001.2b	3	dasd	2347
<input type="checkbox"/> IBM.750000000DX111.0001.2b	3	dasd	46068
<input type="checkbox"/> IBM.750000000DKRP1.0001.29	3	dasd	46068

Showing 1 to 7 of 7 entries

- FCP Tape Devices
- SAN Adapters

Powered by: Wok Version: 2.2.0-0



KVM for IBM z and LinuxONE v 1.1.2 (GA 10/28/16)



Enhancements to text based installer

- add DASD dynamically
- optional encryption of page volumes
- support for http, https and nfs protocol support for installing packages

Easier and accessible Installer

an additional web browser based installer to simplify installation and fulfill IBM accessibility guidelines

Simplify management of resources on hypervisor

Kimchi/Ginger tool:

- user management
- kernel modules list/load/unload
- user activity log
- multi-cultural support

Monitor Virtual Servers Status via Nagios

additional Nagios plug-in (check_kvm)

Make virtualization management simpler and familiar

provides same virt-manager and virt-install tooling experience as on x86



virt-manager UI



New VM (on s38lp56) ×

Create a new virtual machine
Step 3 of 5

Choose Memory and CPU settings

Memory (RAM): MIB
Up to 29970 MIB available on the host

CPUs:
Up to 8 available

New VM (on s38lp56) ×

Create a new virtual machine
Step 4 of 5

Enable storage for this virtual machine

Create a disk image for the virtual machine
 GiB
13.1 GiB available in the default location

Select or create custom storage



KVM for IBM z and LinuxONE v 1.1.2 (GA 10/28/16)



Dynamic addition of vCPUs to Guests
CPU hotplug capability to virtual servers

Guest boot from ISO
Simplifies installation of guest operating systems

Support for IBM License Metric Tool (ILMT)
Customers can determine their full and sub-capacity PVU licensing for IBM software

Identify potential problems in hypervisor configuration
Recommendations made by Healthchecker / included plugins

Provides high availability of resources – especially guests
Includes industry standard Pacemaker and Corosynch and System z specific fencing agent



Pacemaker Overview - Structure

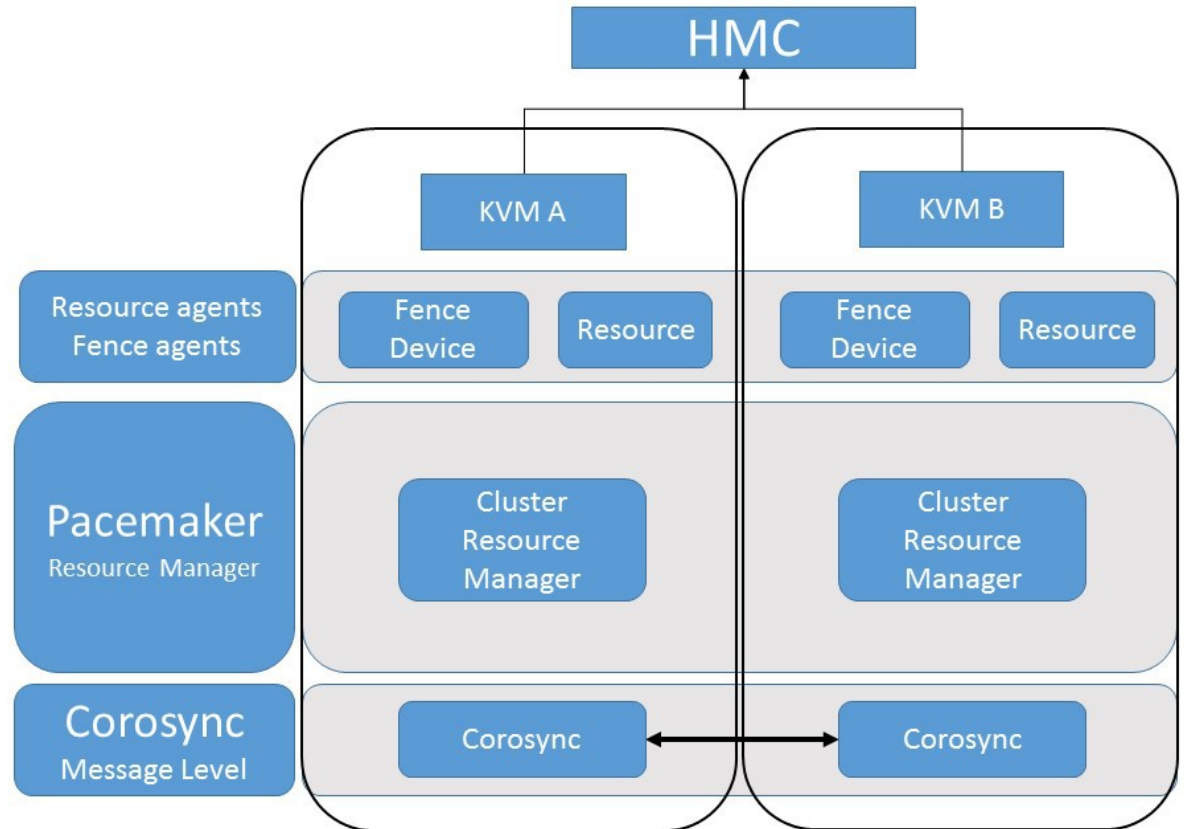


Supported resource agents:

1. VirtualDomain
2. MailTo
3. ethMonitor
4. Filesystem
5. symlink

Fence agent:

- Activates fence device when a node needs to be shutdown to prevent data loss
- fence_ibmz is the fence agent shipped in KVM for IBM z Systems





Candidates for future releases

There are **no** IBM commitments in place at this point to deliver these capabilities. We do not recommend basing any plans on these capabilities

- Ceph support
- Spectrum Scale (GPFS) support
- Hipersockets support
- Mellanox ConnectX (RoCE) Support
- Guest Network boot
- Nested Virtualization
- OpenStack compute node
- Kimchi guest management
- HPM memory management
- Zabbix agent
- GDPS support

We welcome YOUR input;
become a sponsored user !

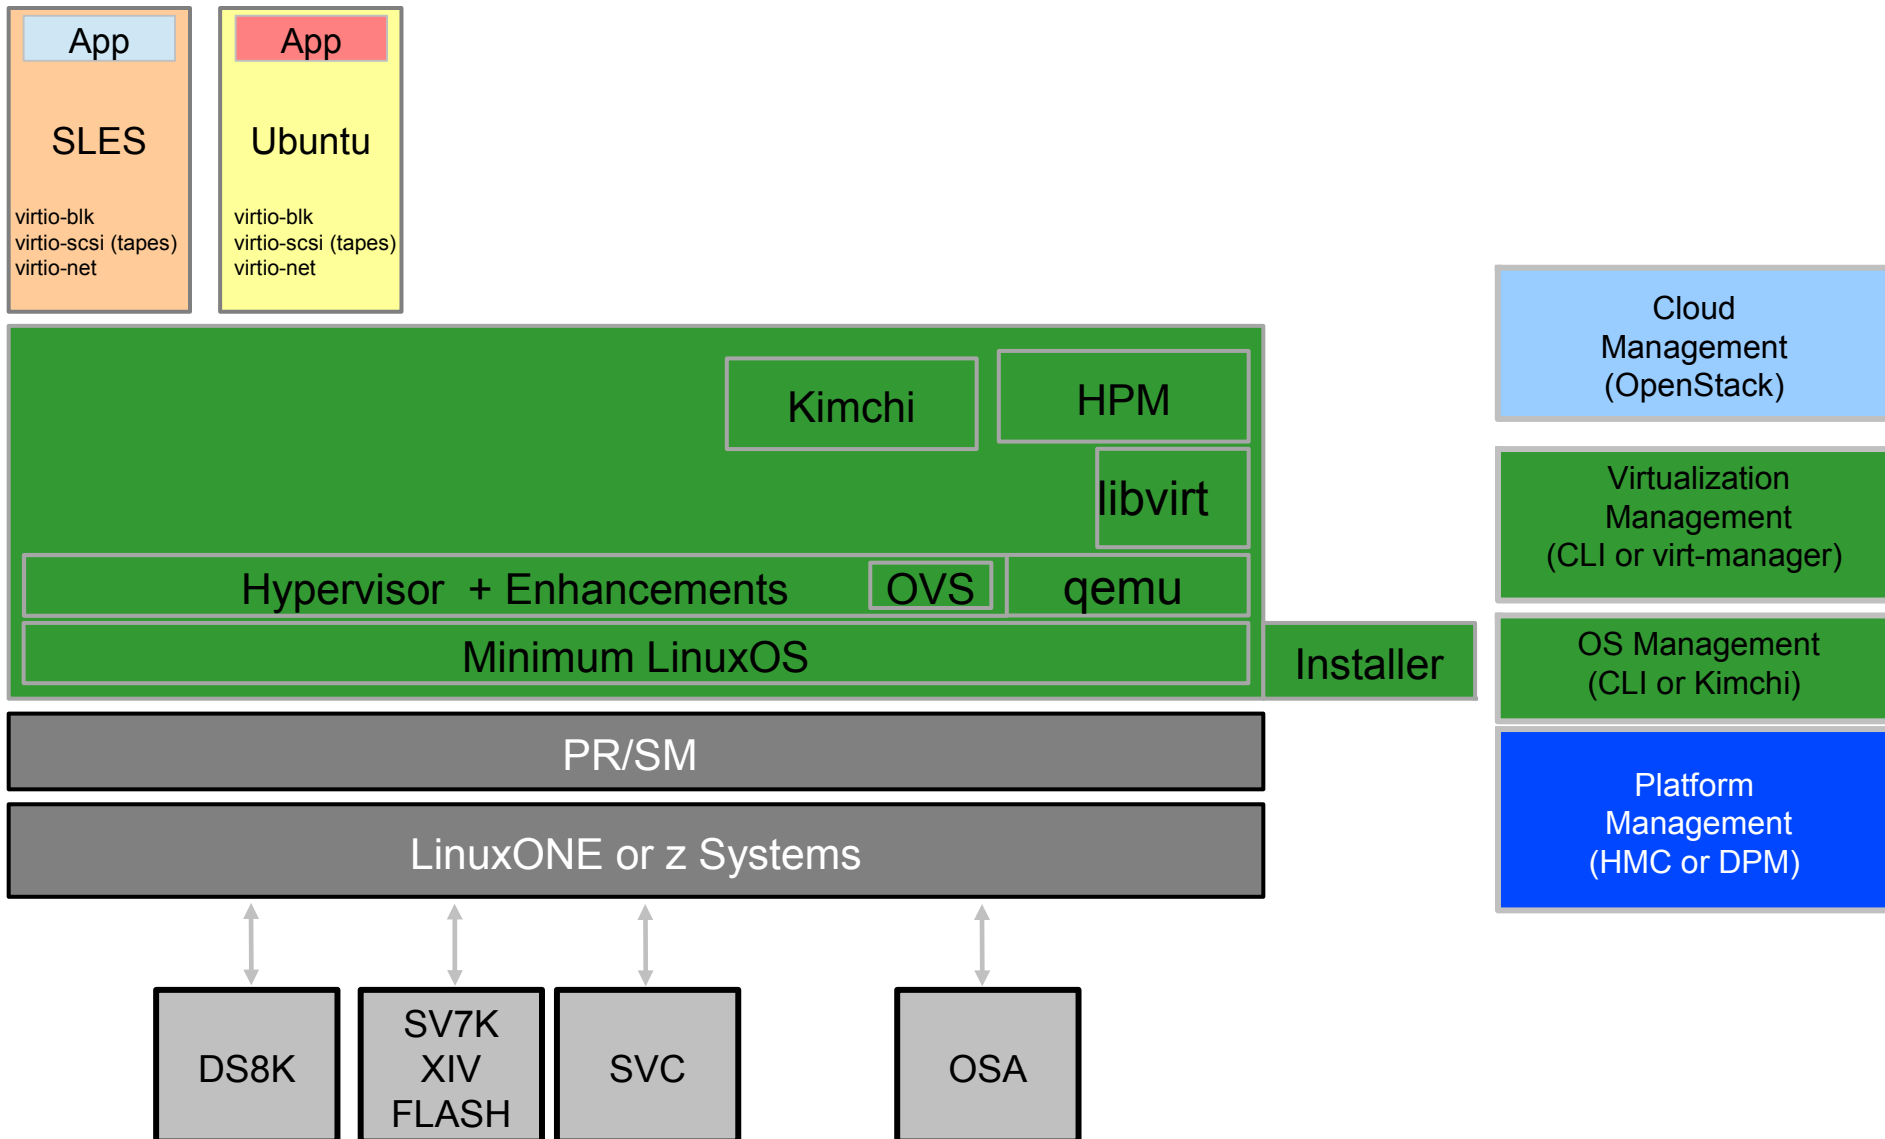


Agenda

- What is KVM
- What is KVM for IBM z Systems and LinuxONE
 - What is in 1.1.1
- What is new in 1.1.2
- What could come thereafter
- Existing Systems Management Tooling




Solution View





Standard Interfaces for Infrastructure Management

- Tasks performed by Linux HostOS/Hypervisor Administrator to manage a system
- Boot / Shutdown the Host operating system
- Setup Security
 - Firewalls, SELinux, PAM config
- Manage System Resources
 - configure systemd
 - automate system tasks
- Manage Users and Groups
- Configure Network
 - configure attached devices including bonding
 - focus on administering connectivity via libvirt between guest/host network
- Configure Storage
 - format/partition devices, configure attached devices including multipathing
 - manage file systems, LVM




- ▶ Use standard Linux CLIs and config files and/or
- ▶ Use Kimchi UI



Standard Interfaces for KVM Virtualization Management

- c-library to interact with hypervisors
 - **KVM**, Xen, LXC
 - Virtual machine management API
 - create, destroy, start, stop, suspend, resume VMs
 - basic support for static and live migration
 - Basic management of virtual networks and storage
 - Support for Linux Control Groups, i.e. fine grained resource management
 - SELinux Support with sVirt
 - Every KVM management application uses libvirt
- ▶ Use virsh as a command-line front-end to libvirt and/or
- ▶ Use virt-manager UI





Standard Interfaces for Cloud Management

SD Infrastructure APIs

- Services and Resources
- Server, Storage and Network
- Broad Ecosystem Forming



OpenStack API

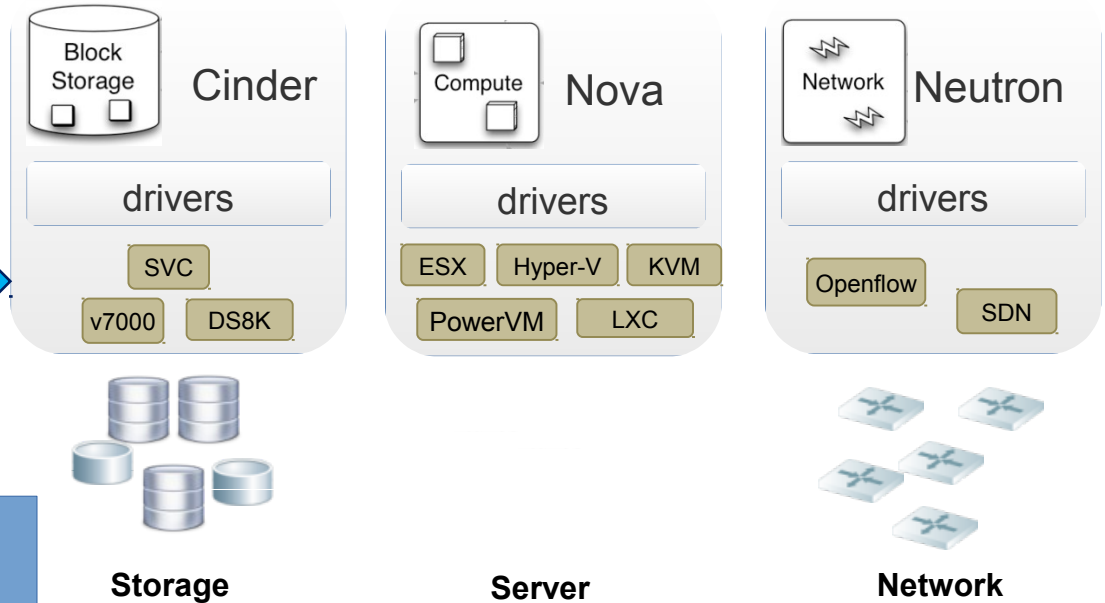
SD Infrastructure Services

- Software Image Services
- Infrastructure Patterns
- VM Placement Intelligence



Vendor Led Scalable Model

- Drivers provided by the vendors
- Broad Ecosystem
- Management standardization



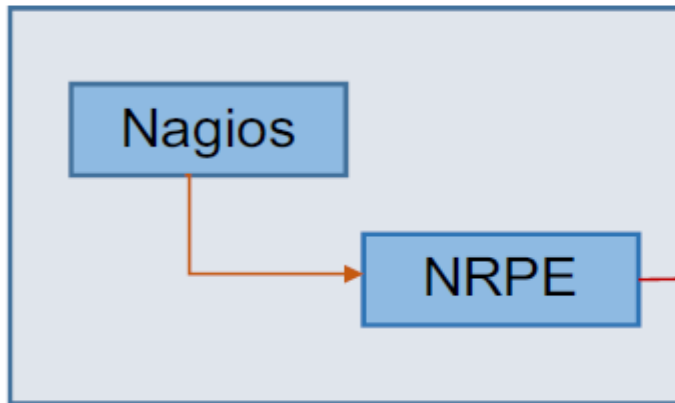
Code enabling KVM for IBM z is in the upstream code base since the OpenStack Kilo release*
 OpenStack distribution vendors can add support for KVM for IBM z based on that code

*<http://docs.openstack.org/developer/nova/support-matrix.html>

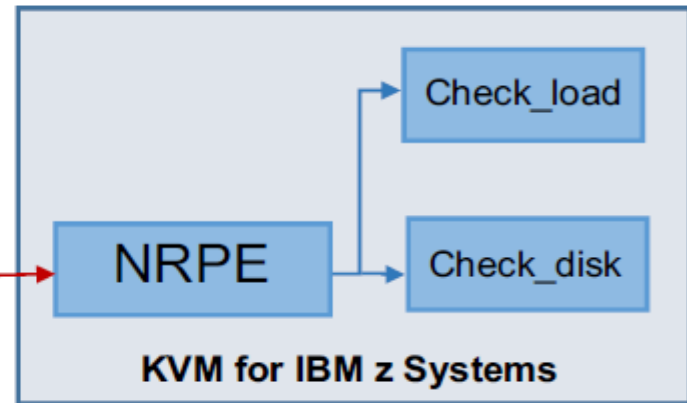


Availability Monitoring via Nagios

Nagios Server



Monitored Host



Nagios®

- General
 - Home
 - Documentation
- Current Status
 - Tactical Overview
 - Map (Legacy)
 - Hosts
 - Services
 - Host Groups
 - Summary
 - Grid
 - Service Groups
 - Summary
 - Grid
 - Problems
 - Services (Unhandled)
 - Hosts (Unhandled)
 - Network Outages
 - Quick Search:
- Reports
 - Availability
 - Trends (Legacy)
 - Alerts
 - History
 - Summary
 - Histogram (Legacy)
 - Notifications
 - Event Log

Current Network Status
 Last Updated: Thu Oct 1 11:15:12 EDT 2015
 Updated every 90 seconds
 Nagios® Core™ 4.1.1 - www.nagios.org
 Logged in as nagiosadmin

Host Status Totals

Up	Down	Unreachable	Pending
3	0	0	0
All Problems		All Types	
0		3	

Service Status Totals

Ok	Warning	Unknown	Critical	Pending
9	2	0	2	0
All Problems		All Types		
4		13		

View Host Status Detail For All Host Groups
 View Status Overview For All Host Groups
 View Status Summary For All Host Groups
 View Status Grid For All Host Groups

Service Status Details For All Host Groups

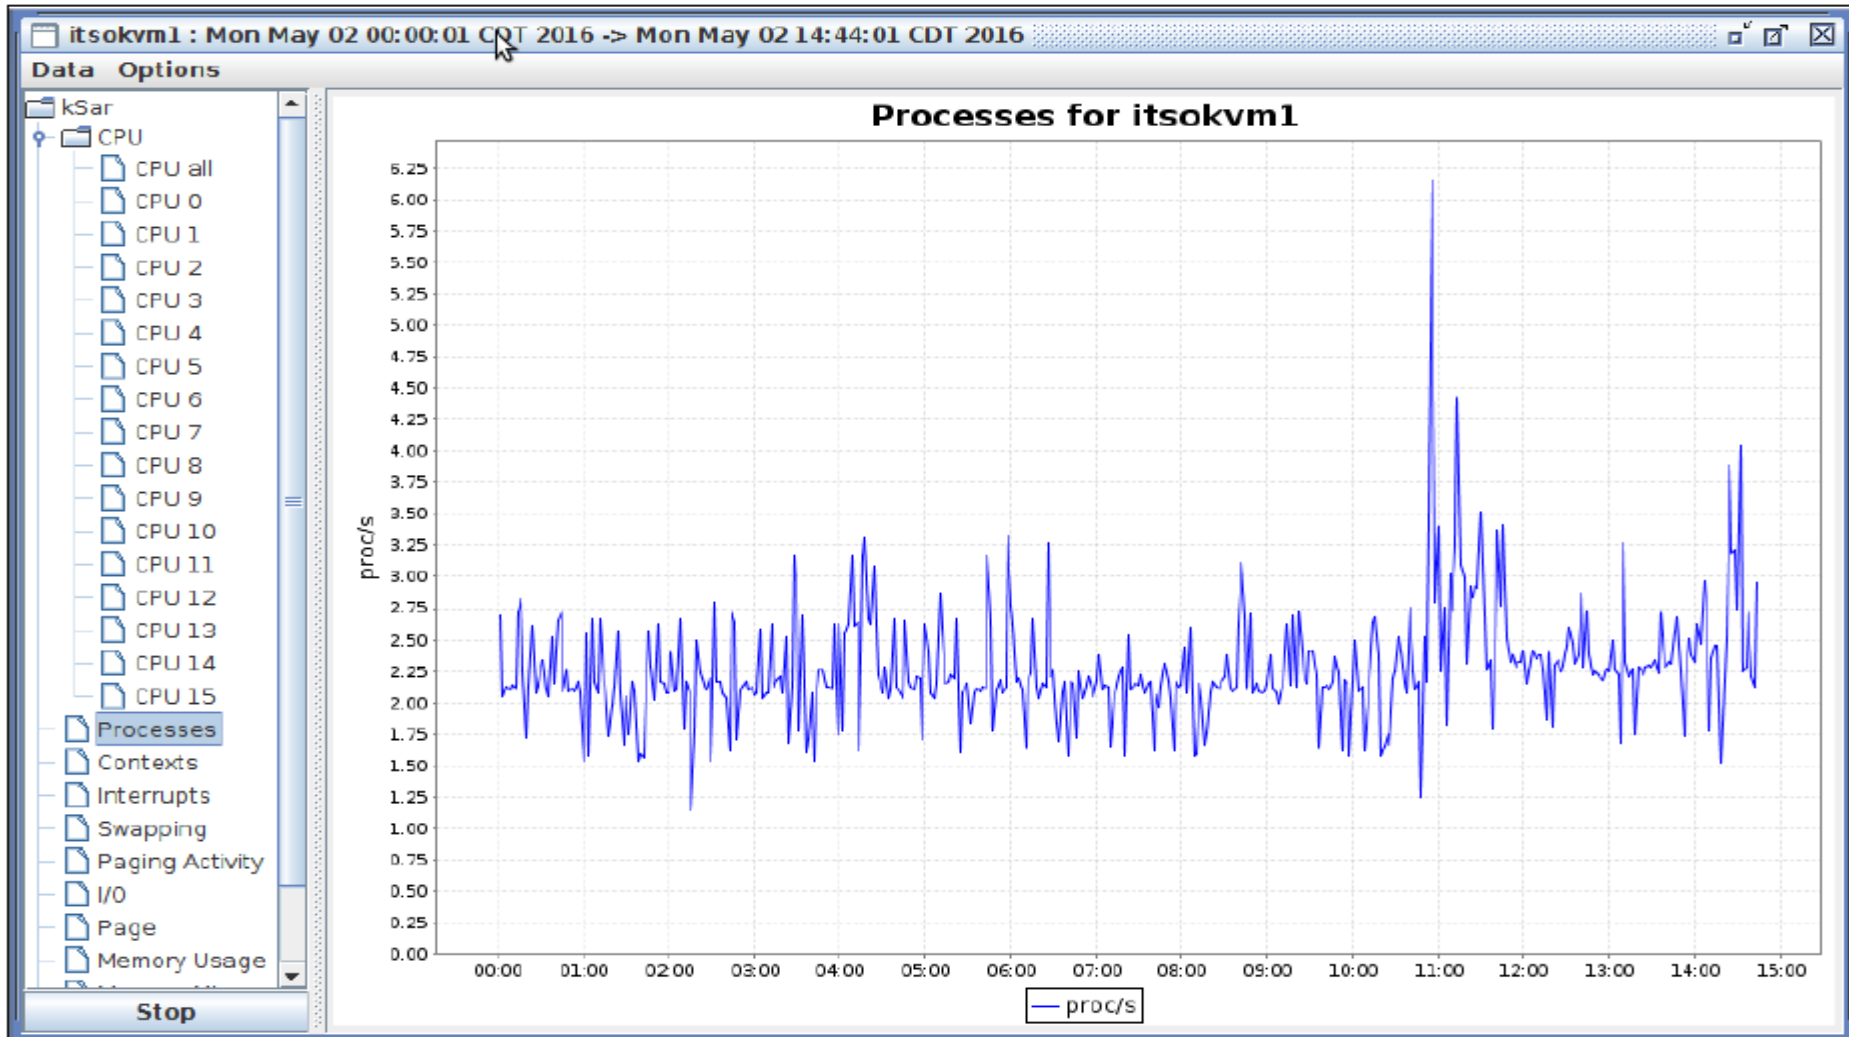
Limit Results: 100

Host	Service	Status	Last Check	Duration	Attempt	Status Information
ibszkvm1	CPU Load	OK	10-01-2015 11:11:37	1d 23h 23m 35s	1/3	OK - load average: 0.20, 0.20, 0.24
	Current Users	OK	10-01-2015 11:06:38	1d 23h 8m 31s	1/3	USERS OK - 3 users currently logged in
	FTP Monitoring	CRITICAL	10-01-2015 11:07:40	1d 23h 21m 27s	3/3	NRPE: Command 'check_procfs' not defined
	SSH Monitoring	CRITICAL	10-01-2015 11:06:47	1d 23h 20m 23s	3/3	NRPE: Command 'check_disk' not defined
	Total Processes	WARNING	10-01-2015 11:09:53	1d 23h 19m 19s	3/3	PROCS WARNING: 181 processes
localhost	Current Load	OK	10-01-2015 11:12:42	1d 23h 38m 54s	1/4	OK - load average: 1.99, 2.22, 2.32
	Current Users	OK	10-01-2015 11:11:57	1d 23h 38m 14s	1/4	USERS OK - 1 users currently logged in
	HTTP	WARNING	10-01-2015 11:12:37	1d 23h 37m 34s	4/4	HTTP WARNING: HTTP/1.1 403 Forbidden - 4284 bytes in 0.001 second resp
	PING	OK	10-01-2015 11:13:07	1d 23h 38m 59s	1/4	PING OK - Packet loss = 0%, RTA = 0.06 ms
	Root Partition	OK	10-01-2015 11:13:48	1d 23h 38m 24s	1/4	DISK OK - free space: / 32476 MB (73% inode=93%):
	SSH	OK	10-01-2015 11:14:24	1d 23h 35m 44s	1/4	SSH OK - OpenSSH_6.6.1 (protocol 2.0)
	Swap Usage	OK	10-01-2015 11:15:01	1d 23h 35m 4s	1/4	SWAP OK - 88% free (3358 MB out of 3815 MB)
	Total Processes	OK	10-01-2015 11:10:37	1d 23h 34m 32s	1/4	PROCS OK: 83 processes with STATE = RSZDT

Results 1 - 13 of 13 Matching Services

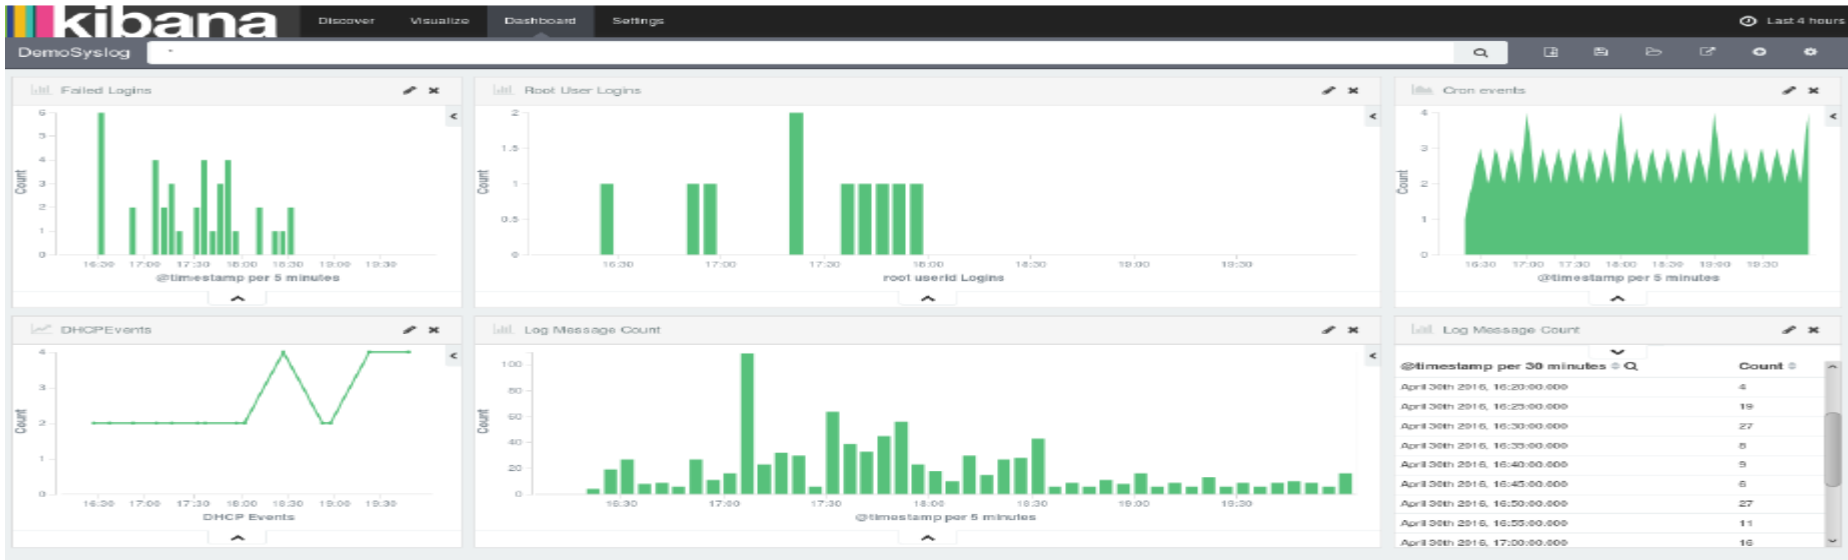
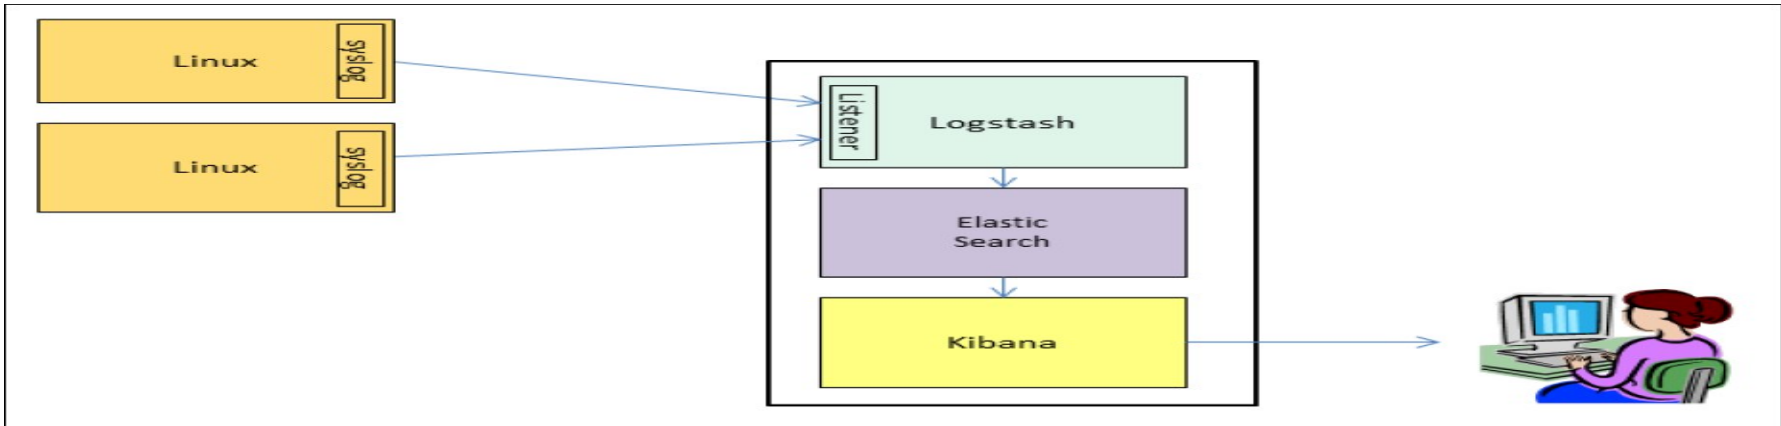


Performance Monitoring via Sar/kSar





Availability/Performance Monitoring via ELK





Hypervisor Performance Manager

- Supports policy-based goal-oriented monitoring and management of CPU resources
- Introduces the concept of workload resource groups, which define the context for monitoring and management
- Shipped as part of the KVM for IBM z delivery
 - Optionally enabled
- Scope of management is single KVM for IBM z instance
 - zHPM has no knowledge outside of its KVM for IBM z instance
- Controlled through RESTful Web Services APIs and CLI
 - APIs
 - Point of integration with higher-level virtualization management solutions
 - Support for scripting
 - Fully documented external interface
 - CLIs provide support for local administration



Where to get detailed Information

- **Product Documentation** at http://www-01.ibm.com/support/knowledgecenter/linuxonibm/liaaf/lnz_r_kvm.html
 - KVM for IBM z Systems: Planning and Installation Guide SC27-8236-02
 - KVM for IBM z Systems: Administration Guide SC27-8237-02
 - Linux on z Systems: Virtual Server Management SC34-2752-02
 - KVM Virtual Server Management Tools SC34-2763-00
 - Linux on z Systems: Virtual Server Quick Start SC34-2753-02
 - Linux on z Systems: Device Drivers, Features, and Commands for Linux as a KVM Guest SC34-2756
 - ...
- **Redbook: Getting Started with KVM for IBM z System**
<https://www.redbooks.ibm.com/abstracts/sg248332.html?Open>
- **Performance Data / Planning Tools**
 - Limits: <http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/PRS5331>
 - Large Systems Performance Reference (LSPR):
 - <https://www-304.ibm.com/servers/resourcelink/lib03060.nsf/pages/lspriTRKVMonZv110?OpenDocument>
 - zPCR
 - <http://www-03.ibm.com/support/techdocs/atmastr.nsf/WebIndex/PRS1381>
- **List of supported IBM SW:**
<http://www.ibm.com/software/reports/compatibility/clarity/productsOnVe.html>
- **Request for Enhancement (RFE) Community on developerWorks**
[https://www.ibm.com/developerworks/rfe/execute?
use_case=changeRequestLanding&BRAND_ID=0&PROD_ID=1532&x=17&y=18](https://www.ibm.com/developerworks/rfe/execute?use_case=changeRequestLanding&BRAND_ID=0&PROD_ID=1532&x=17&y=18)



Questions?

