

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT

**LEARN. NETWORK.
EXPERIENCE OPEN SOURCE.**

www.theredhatsummit.com

RHEV-M Architecture

Itamar Heim

Consulting Software Engineer, Red Hat

June 23rd, 2010

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



Agenda

- .RHEV Overview
- .Platform and Architecture
- .RHEV Manager
- .Host
- .Guest
- .Storage
- .Network
- .VDI

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



Disclaimers

- .Any forward looking statements are for the sake of technical discussion rather than product roadmap
- .This session will cover the high level architecture of RHEV Manager - details may be abstracted for simplicity

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



Red Hat Enterprise Virtualization

• Advanced KVM hypervisor technology and enterprise grade centralized management

• Leading performance and scalability for enterprise applications

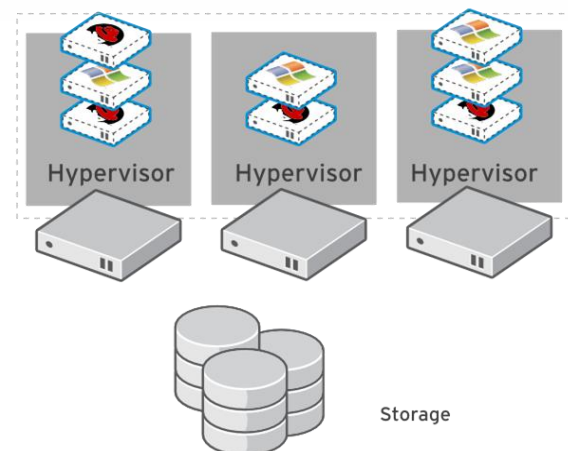
• Highest level of kernel level **security** and isolation with SE-Linux

• Large **eco-system** of thousands of hardware and software vendors

• Up to 70% cost savings relative to VMware

RED HAT ENTERPRISE VIRTUALIZATION MANAGER

Name	Cluster	Host	IP Address	Memory	CPU	Network	Display	Status
DNS-RHEL4	Default	rhev-1.rhev.redhat.com		0%	0%	0%	VNC	Down
RHEL3	Default	rhev-1.rhev.redhat.com		0%	12%	0%	VNC	Up
RHEL4	Default	sandbox.rhev.redhat.com		0%	0%	0%	VNC	Up
RHEL5	Default	rhev-3.rhev.redhat.com		0%	0%	0%	VNC	Up
RHEL5_SAP	Default	sandbox.rhev.redhat.com		0%	0%	0%	VNC	Up
RHEL5_webserver	Default	sandbox.rhev.redhat.com		0%	0%	0%	VNC	Up
RHEL5_Desktop	Default			0%	0%	0%		Down
winxp3	Default	rhev-3.rhev.redhat.com	172.16.31.4	42%	0%	0%	VNC	Up
winxp8	Default	rhev-1.rhev.redhat.com	172.16.31.5	45%	31%	0%	Spice	Up
Windows7	Default	rhev-1.rhev.redhat.com		0%	0%	0%	Spice	Up
Winxp-1	Default	rhev-2.rhev.redhat.com		0%	0%	1%	Spice	Up
Winxp-2	Default			0%	0%	0%		Down
Winxp-3	Default	rhev-2.rhev.redhat.com	172.16.31.4	38%	0%	1%	Spice	Up
Winxp-4	Default	rhev-2.rhev.redhat.com		0%	0%	1%	Spice	Up
Winxp-5	Default	rhev-2.rhev.redhat.com		0%	0%	1%	Spice	Up



SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



RED HAT ENTERPRISE VIRTUALIZATION OVERVIEW

The screenshot displays the RHEV web interface. At the top, it shows the user is logged in as 'rhevadmin'. Below the navigation tabs, a table lists several virtual machines:

Name	Cluster	Host	IP Address	Memory	CPU	Network	Display	Status	Uptime	Logged-in User
RHEL3	Default	beta		0%	0%	0%	VNC	Up	5 min	Unknown
RHEL4	Default	alpha		0%	0%	0%	VNC	Up	5 min	Unknown
RHEL5	Default	beta		0%	0%	0%	VNC	Up	5 min	Unknown
RHEL6_Beta1	Default	beta		0%	0%	0%	VNC	Up	5 min	Unknown
Win7Machine	Default	beta	169.254.184.t	14%	0%	0%	Spice	Up	3 min	None
WinXP1	Default	beta	169.254.56.1	17%	1%	0%	VNC	Up	2 days	Administrator@i

Below the table, a detailed view for a VM is shown with the following specifications:

- Name: RHEL3
- Description: RHEL3
- Template: Blank
- Defined Memory: 2048 MB
- Operating System: Red Hat Enterprise Linux 3.x x64
- CPU Cores: 2 (2 Sockets, 1 Cores per Socket)
- Default Display Type: VNC
- Origin: RHEV
- Highly Available: No
- Priority: Low

The interface also includes a 'RHEV Manager' icon and a status message: 'Last Message: May 18, 16:20 VM Win7Machine started on Host beta'.

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



Admin Console

ENTERPRISE VIRTUALIZATION
Logged in user: Administrator | Sign out | Configure | About

Search: Vms: [x] [★] [GO]

Bookmarks: [New] [Edit] [Remove]

- Desktops
- Servers
- Windows_VMs
- RHEL_VMs
- Host_over_80%_cpu
- Host_over_80%_network
- Running_VMs
- All_Hosts

Navigation: Data Centers | Clusters | Hosts | Storage | Virtual Machines | Pools | Templates | Users | Events | Monitor

Actions: [New Server] [New Desktop] [Edit] [Remove] [Migrate] [Make Template] [Restore] [Backup] [Move] [Custom Actions] [Guide Me]

Name	Cluster	Host	IP Address	Memory	CPU	Network	Display	Status
DNS-RHEL4	Default			0%	0%	0%		Down
RHEL3	Default	rhev-1.rhev.redhat.com		0%	0%	1%	VNC	Up
RHEL4	Default	rhev-1.rhev.redhat.com		0%	12%	0%	VNC	Up
RHEL5	Default	sandbox.rhev.redhat.com		0%	0%	0%	VNC	Up
RHEL5_SAP	Default	rhev-3.rhev.redhat.com		0%	0%	0%	VNC	Up
RHEL5_webserver	Default	sandbox.rhev.redhat.com		0%	0%	0%	VNC	Up
RHEL5-Desktop	Default			0%	0%	0%		Down
win2k3	Default			0%	0%	0%		Down
win2k8	Default	rhev-3.rhev.redhat.com	172.16.31.4	42%	0%	0%	VNC	Up
Windows7	Default	rhev-1.rhev.redhat.com	172.16.31.5	69%	51%	0%	Spice	Up
WinXP-1	Default	rhev-2.rhev.redhat.com		0%	0%	1%	Spice	Up
WinXP-2	Default			0%	0%	0%		Down
WinXP-3	Default	rhev-2.rhev.redhat.com	172.16.31.4	38%	0%	1%	Spice	Up
WinXP-4	Default	rhev-2.rhev.redhat.com		0%	0%	1%	Spice	Up
WinXP-5	Default	rhev-2.rhev.redhat.com		0%	0%	1%	Spice	Up

Details for win2k3:

- General: Operating System: Windows 2003, Description: , Template: Blank, Number of CPUs: 2
- Users: ,
- Network Interfaces: ,
- Virtual Disks: ,
- Snapshots: ,
- Applications: ,
- Defined Memory: 512 MB, Domain: , Time Zone: Eastern Standard Time

Search driven UI makes managing thousands of objects easy

Bookmarks – dynamic filters and other advanced GUI functions

Cross-correlate information across users, VMs, hosts, applications, status etc.

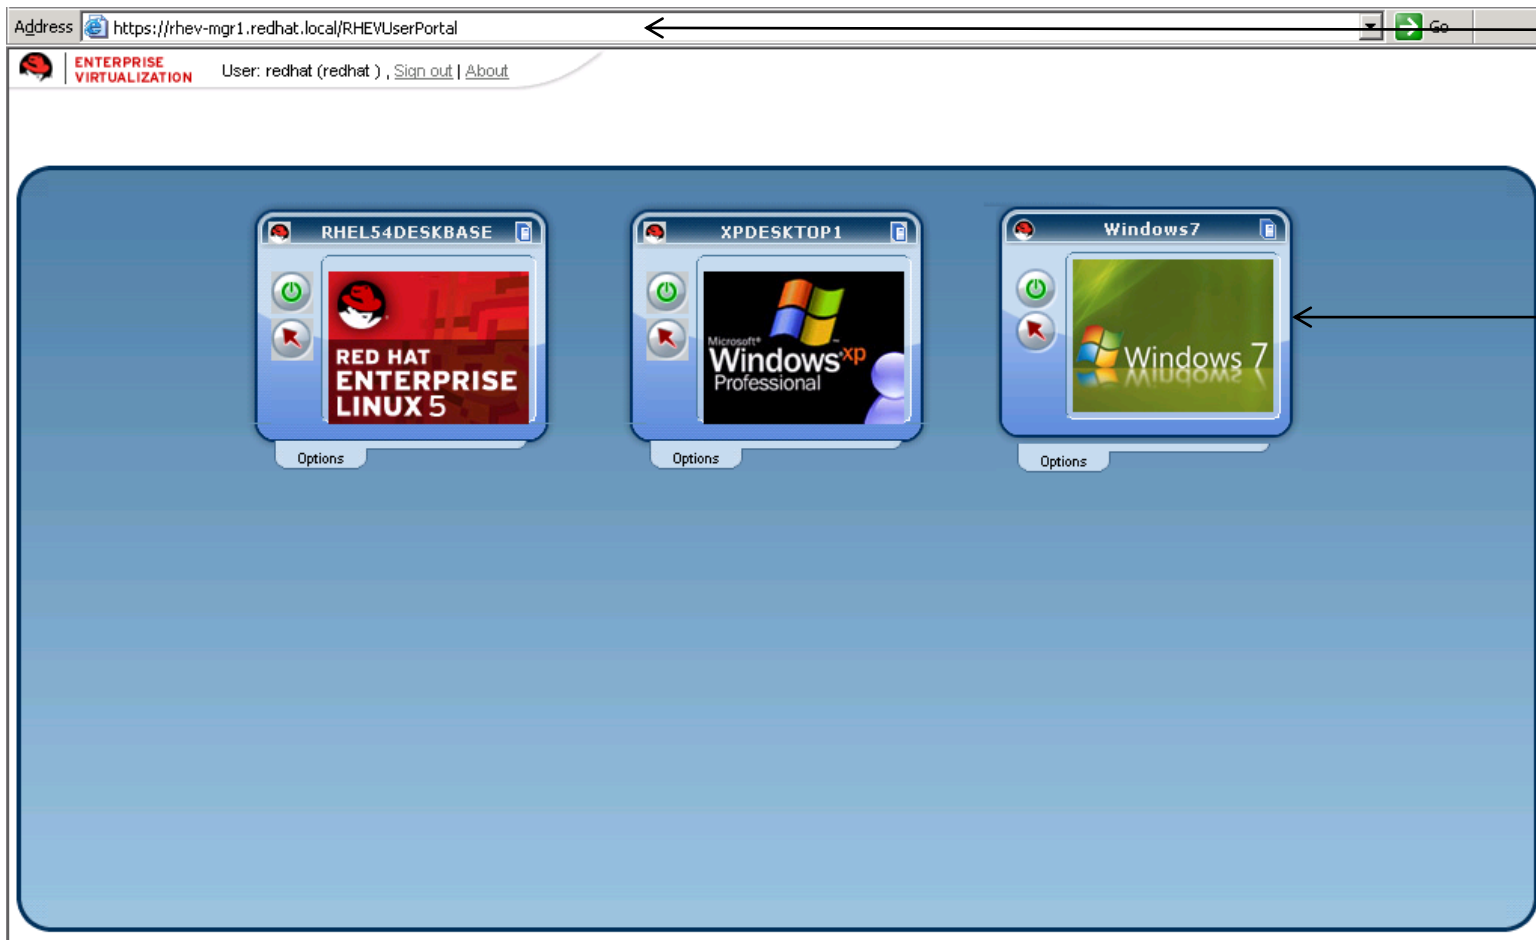
SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



User Portal



Access from Internet Explorer or Firefox on PC or thin client

User can start and stop the VM, access the console and select remote protocol – SPICE or RDP

Or the portal can automatically launch a single VM for the user

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT




PowerShell

Select all VM's consuming CPU

```
PS E:\> select-vm -searchtext "status = up and cpu_usage > 1"

VmId           : a6720e93-6f9a-45b1-82c0-4b4f558b82ad
Name           : 2008R2Russian
Description    :
HostClusterId  : 1
TemplateId     : 00000000-0000-0000-0000-000000000000
Session        : Unknown
MemorySize     : 1024
Priority        : 1
Domain         :
OperatingSystem : Windows2008R2x64
CreationDate   : 6/8/2010 6:05:00 PM
NumOfMonitors  : 1
NumOfSockets   : 1
NumOfCpusPerSocket : 1
DefaultHostId  :
NiceLevel      : 0
FailBack       : False
DefaultBootSequence : C
VmType         : Server
HypervisorType : KVM
OperationMode  : FullVirtualized
```



Script and automate all administration

All functions accessible via PowerShell

Shutdown any VM which is consuming 100% cpu

```
PS E:\> select-vm -searchtext "status = up and cpu_usage = 100" | shutdown-vm_
```

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



Windows? Java? Linux

- .Let's get this off the table 😊
- .RHEV 2.1 and 2.2 run on Windows
 - C# Backend code
 - Mature codebase acquired from Qumranet (SolidICE)
- .Next version will have a Java/JBoss backend
- .Following version will have full Linux support

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



Moving to Java

- A porting project is usually risky business
- Need to balance features and porting effort
- Used automatic conversion approach
- Allowed weekly merges from C# to Java
- 100% pass on Java backend for 2.2 automatic tests (thousands of tests based on automation)
- Similar approach used for porting the DB & the UI

<http://lpeer.blogspot.com/2010/04/switching-from-c-to-java.html>

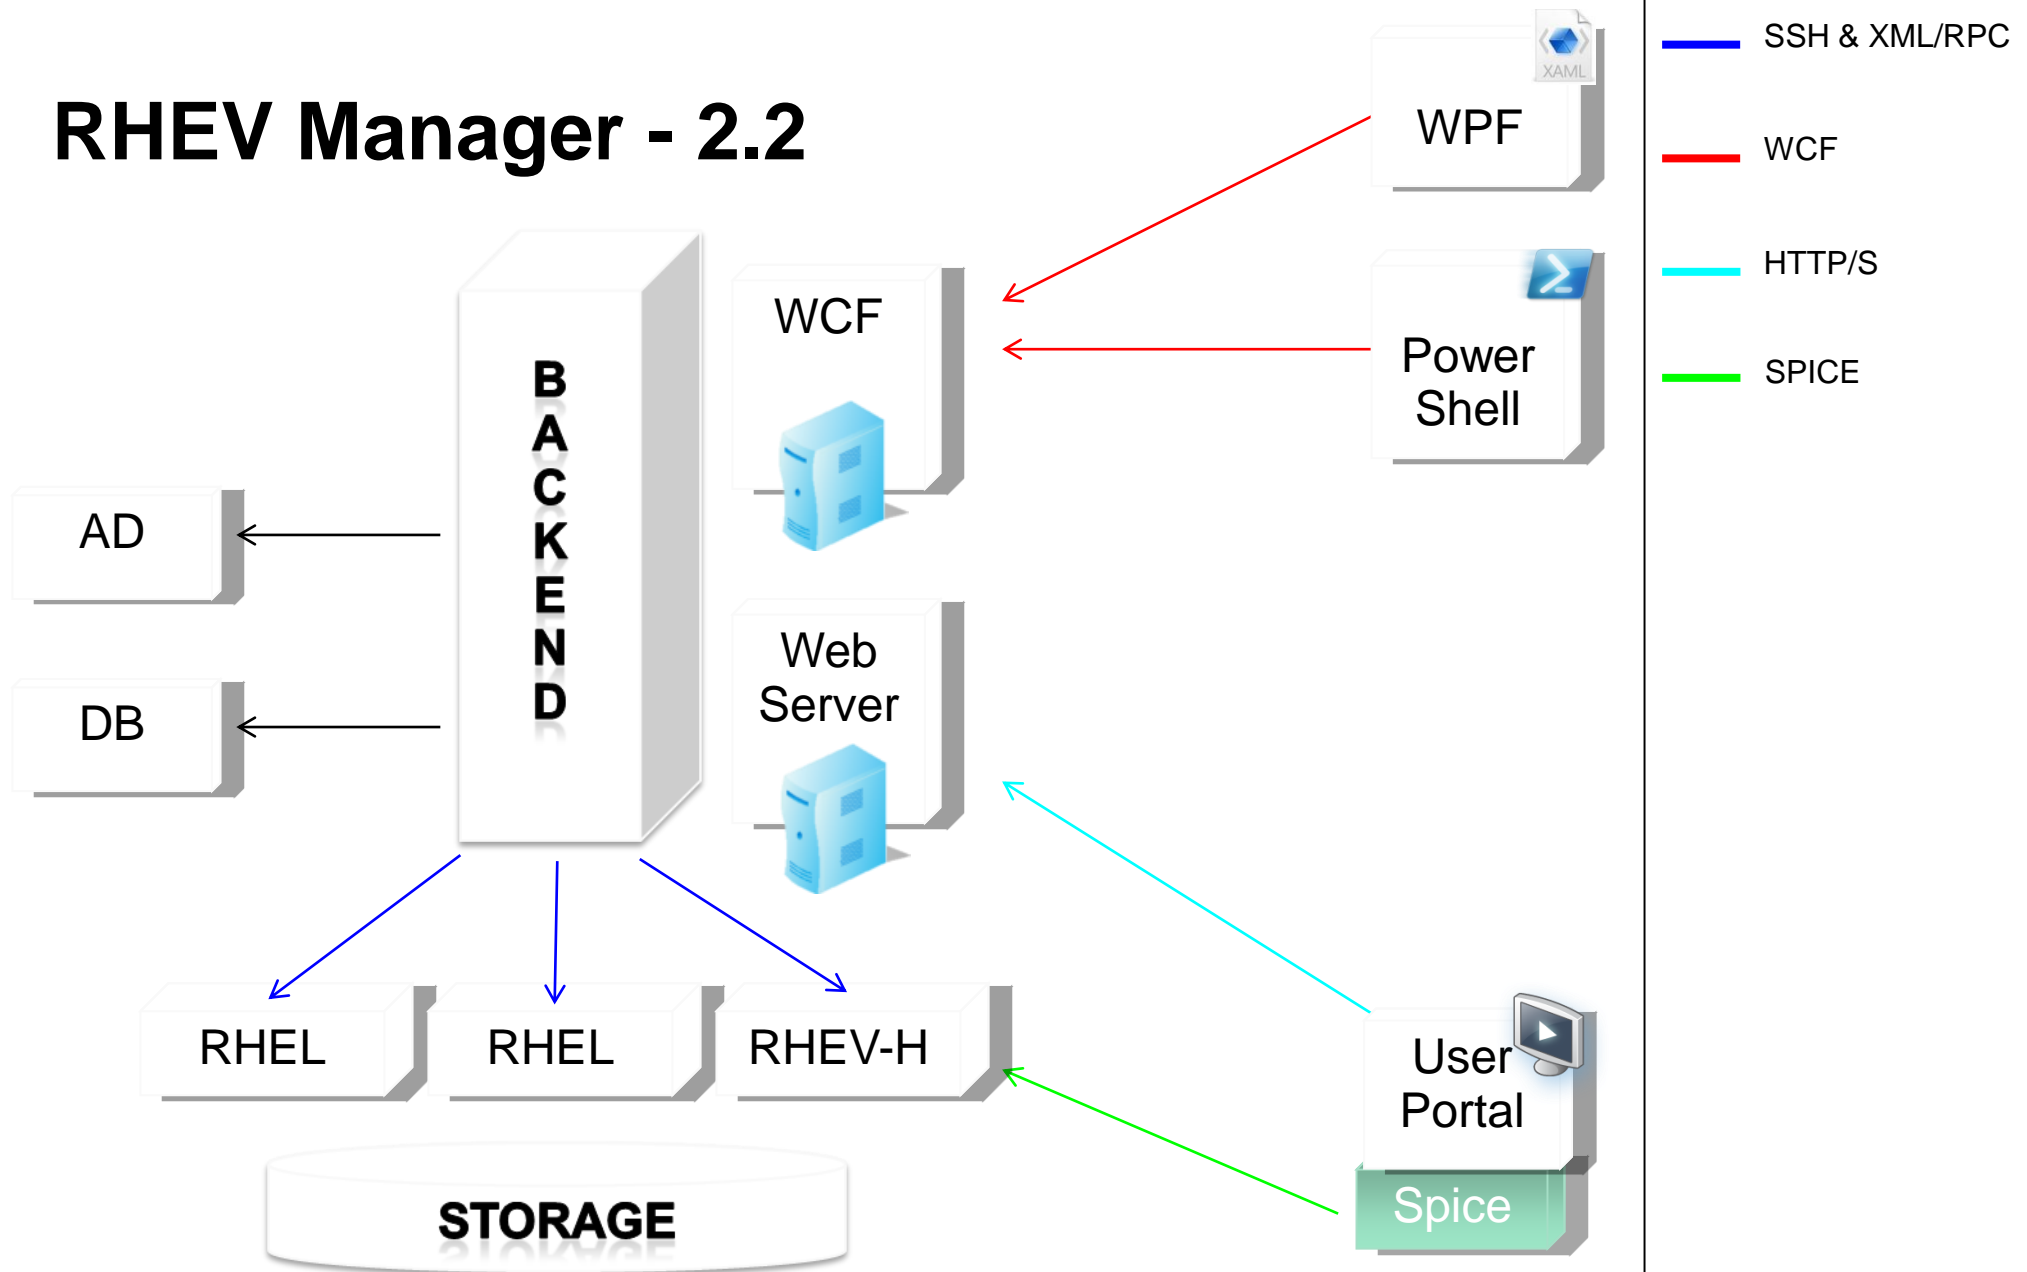
SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



RHEV Manager - 2.2



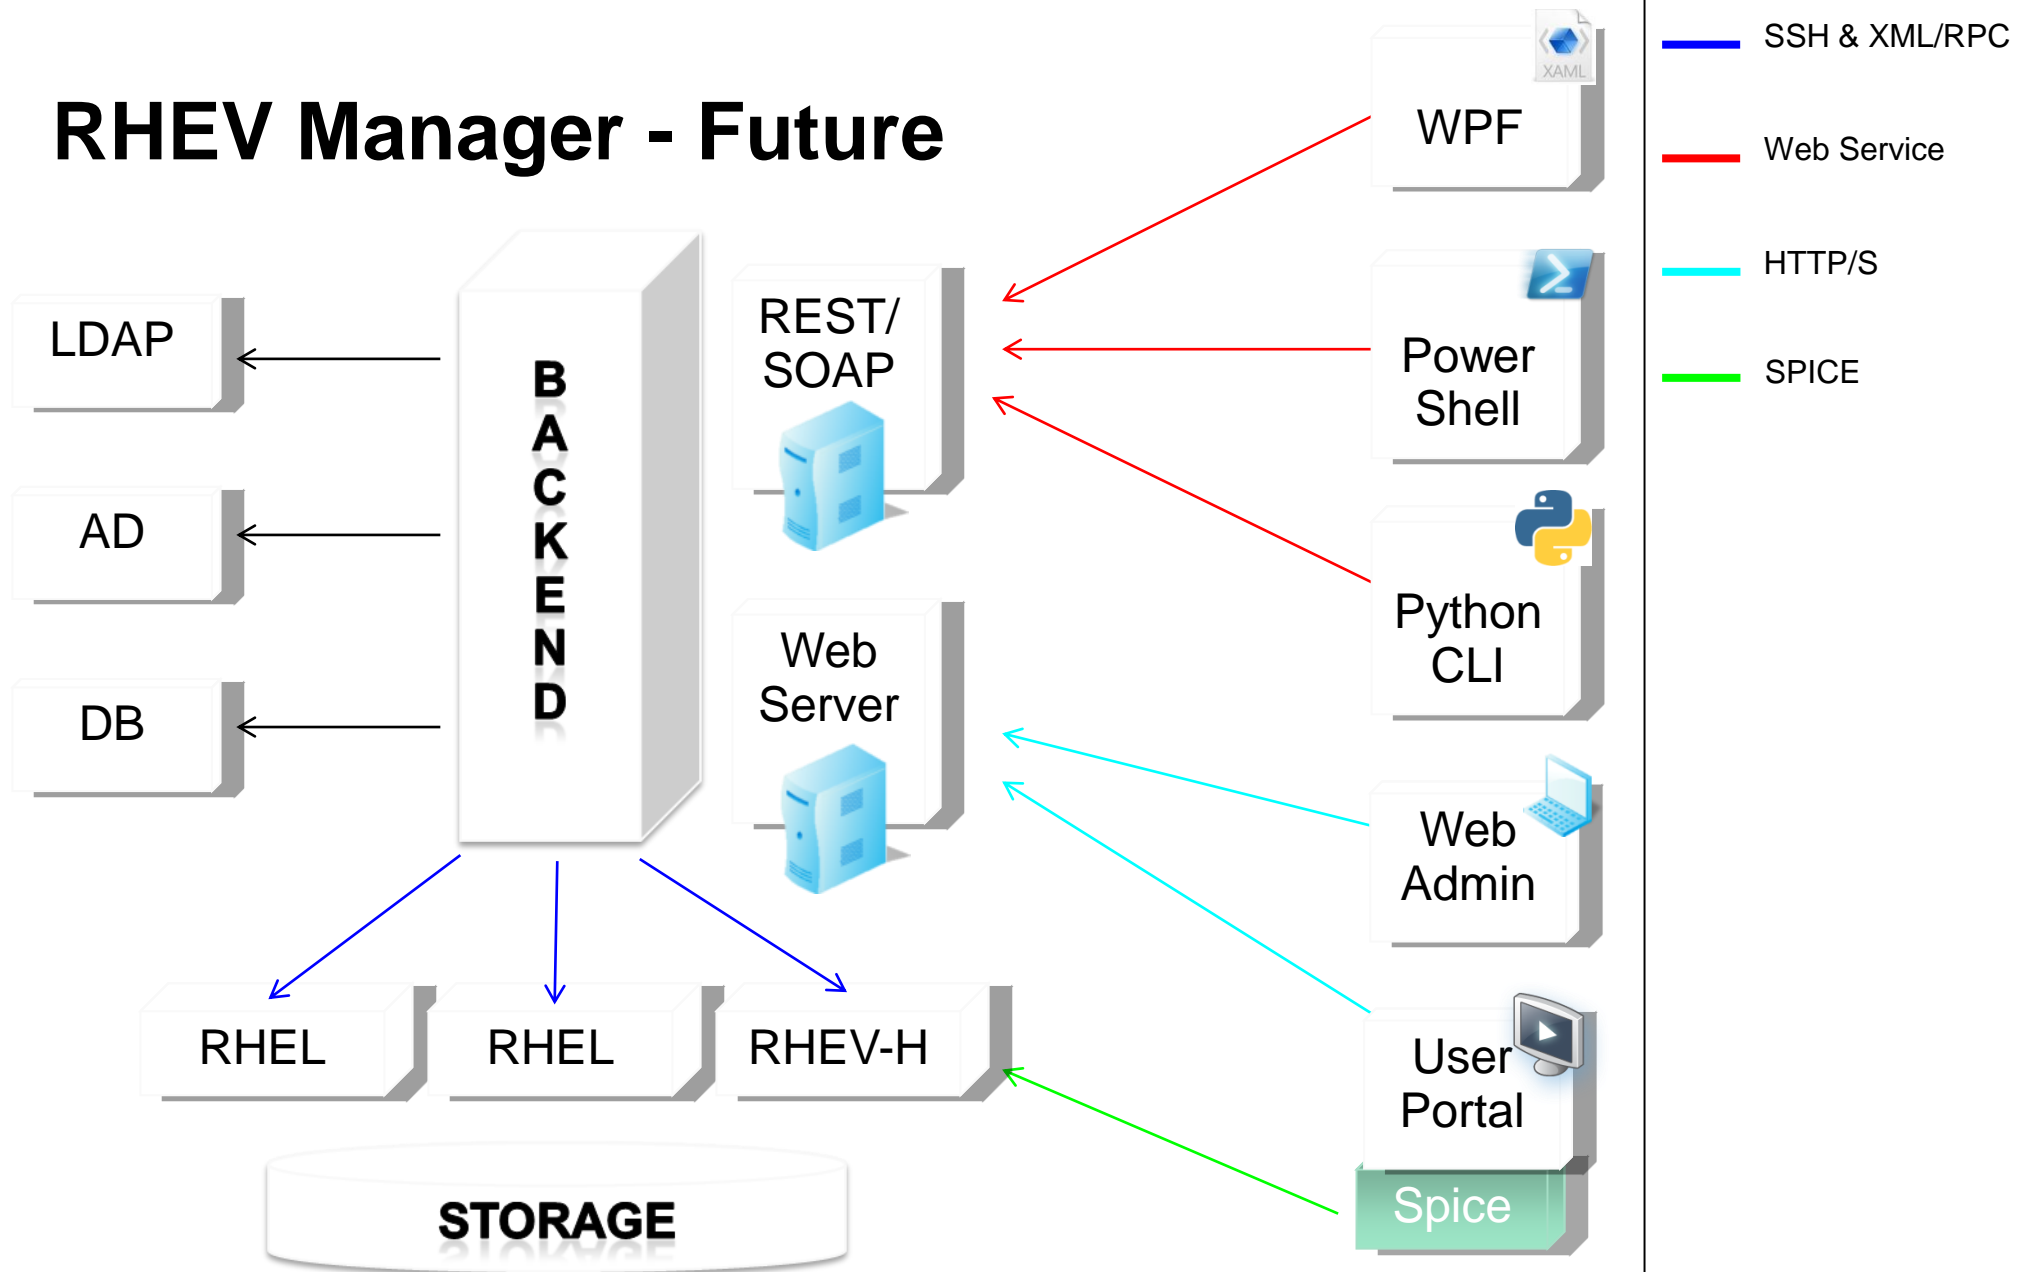
SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



RHEV Manager - Future



SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



Packages

.RHEV-M

.Installer

.Windows guest

.Guest Tools iso

.virtio-win/virtio-drivers.vfd

.Windows client

.Spice ActiveX (in rhev-m installer)

.USB client

.Host

.VDSM for RHEL

.RHEV-H

.Linux guest

.VirtIO drivers

.Spice guest driver

.Linux client

.spice-xpi, spice-client

.spice-usb-redirector

.V2V rpm's

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



RHEV Manager Server

- .Mainly a C# Service with an internal SQL Server DB
- .Also includes
 - .Admin console (WPF)
 - .User portal web site (includes SPICE ActiveX)
 - .Host registration site
 - .Components web site
 - .History DB & service, Notification service
 - .Log collector, ISO Uploader, Configuration tool
 - .NetConsole service

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



RHEV Manager - Host Services

- .Host registration / installation
- .Host configuration
- .Host monitoring
- .Host fencing
- .Storage cluster monitoring
- .Storage provisioning (consume VGs, or create from LUNs)
- .Network configuration management



RHEV Manager - VM Services

- .VM life cycle management
- .VM scheduling and load balancing
- .VM high availability
- .Time leased pool monitoring

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



RHEV Manager – Other Services

- Active Directory cache manager
- Audit log cleanup manager
- Asynchronous search handler

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



Host

- .RHEL or RHEV-H
- .Centrally managed via RHEV-M
- .VDSM installation for RHEL
- .Upgrade for RHEV-H
- .Network
- .Storage

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT




RHEV-M Host Agent (VDSM)

- Covers all functionality required by RHEV-M
- Tailored to RHEV-M flows
- Supports long async task, transactions
- In RHEL 6, will use libvirt for VM life cycle operations



Host – RHEV-M

- XML/RPC
- Net Console
- HTTP/S
- SSH



SUMMIT

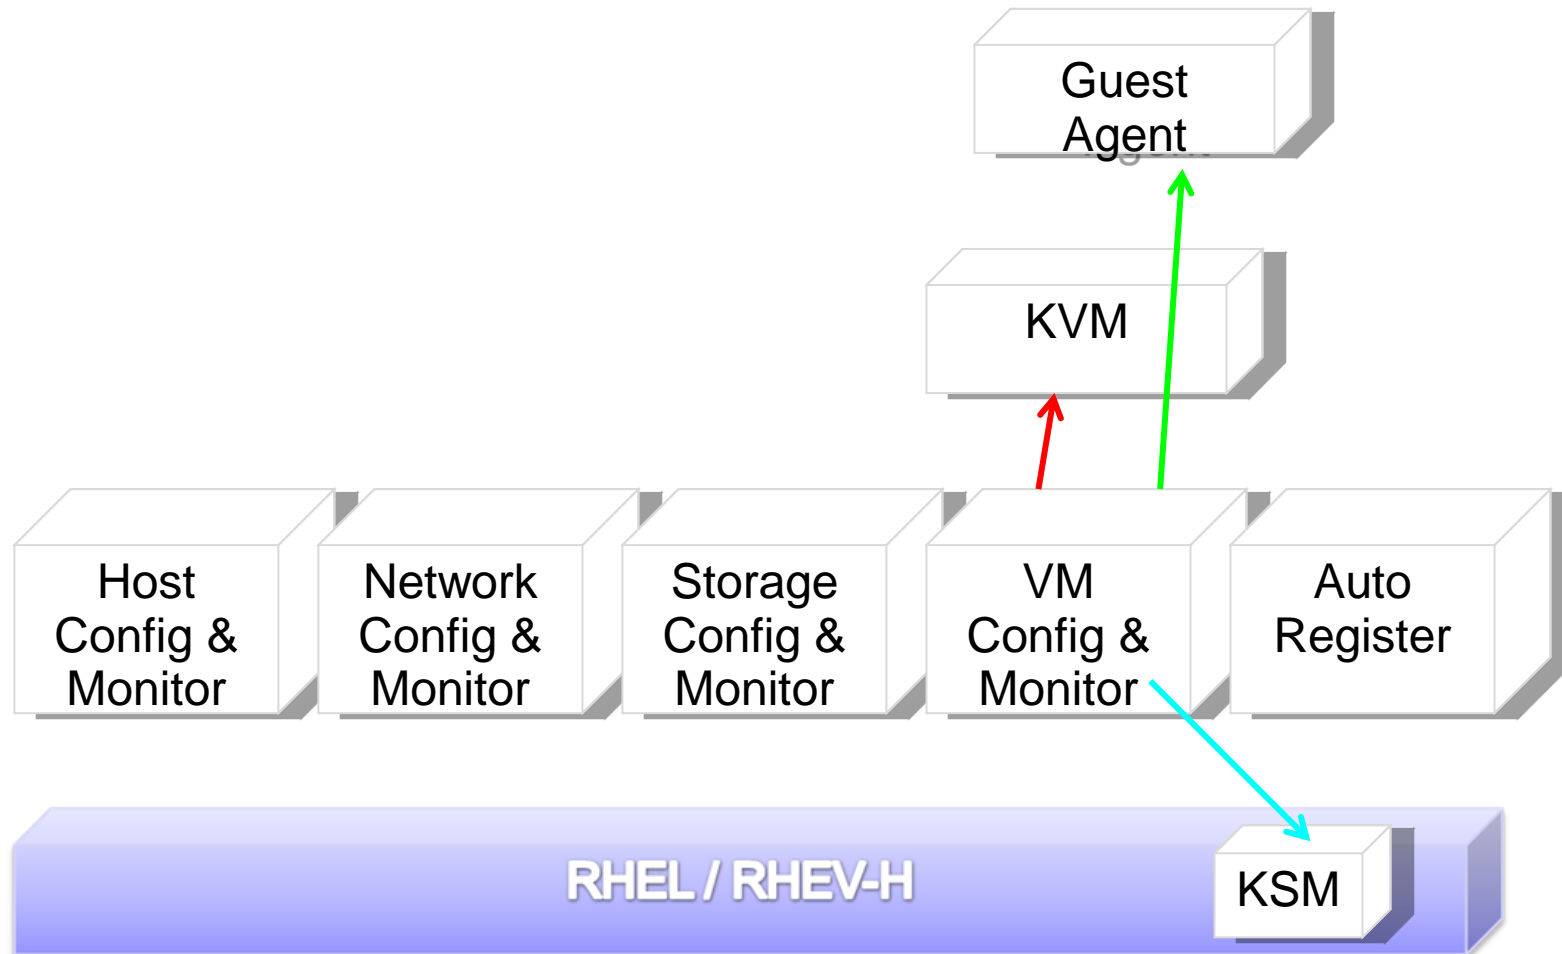

JBoss
WORLD

PRESENTED BY RED HAT

*RHEV-H only



Host - VM



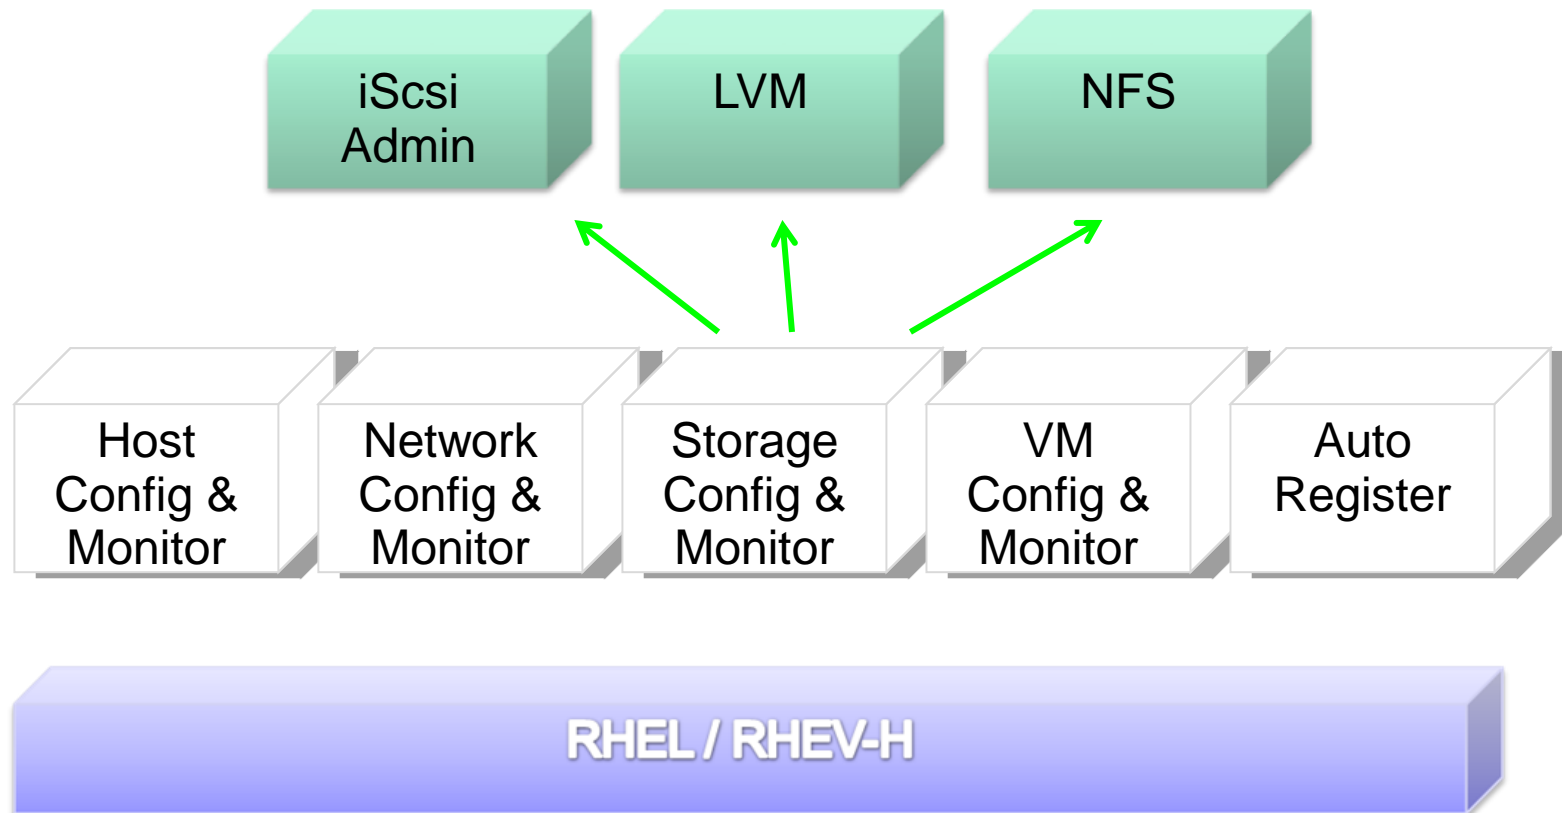
SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



Host - Storage



SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



RHEV Storage

- VDSM manages a Storage Pool, comprised of Storage Domains
- **Storage Pool** - a VM repository that contains meta data about storage domains, storage tasks, VMs, locks, etc.
- **Storage Domain** - a disk image repository
- **Disk Image** – a collection of volumes (chain of snapshots)
- **Volume** - stored as files in NFS, and as Logical Volumes for FC/iScsi
- Thin provisioning for SAN supported (storage mailbox based)



RHEV Storage “Clustering”

- A Storage Pool is implemented as a managed cluster
- Manager is RHEV-M, running on a node external to the hosts using the storage pool
- Heartbeats and fencing (power, and scsi in the future) are used in case of node failures
- Storage based leased locks used as another layer of protection
- Clustering wise - VMs are mostly single reader/writer – locks mostly needed to handle failures
- Can easily create a cluster of >100 nodes

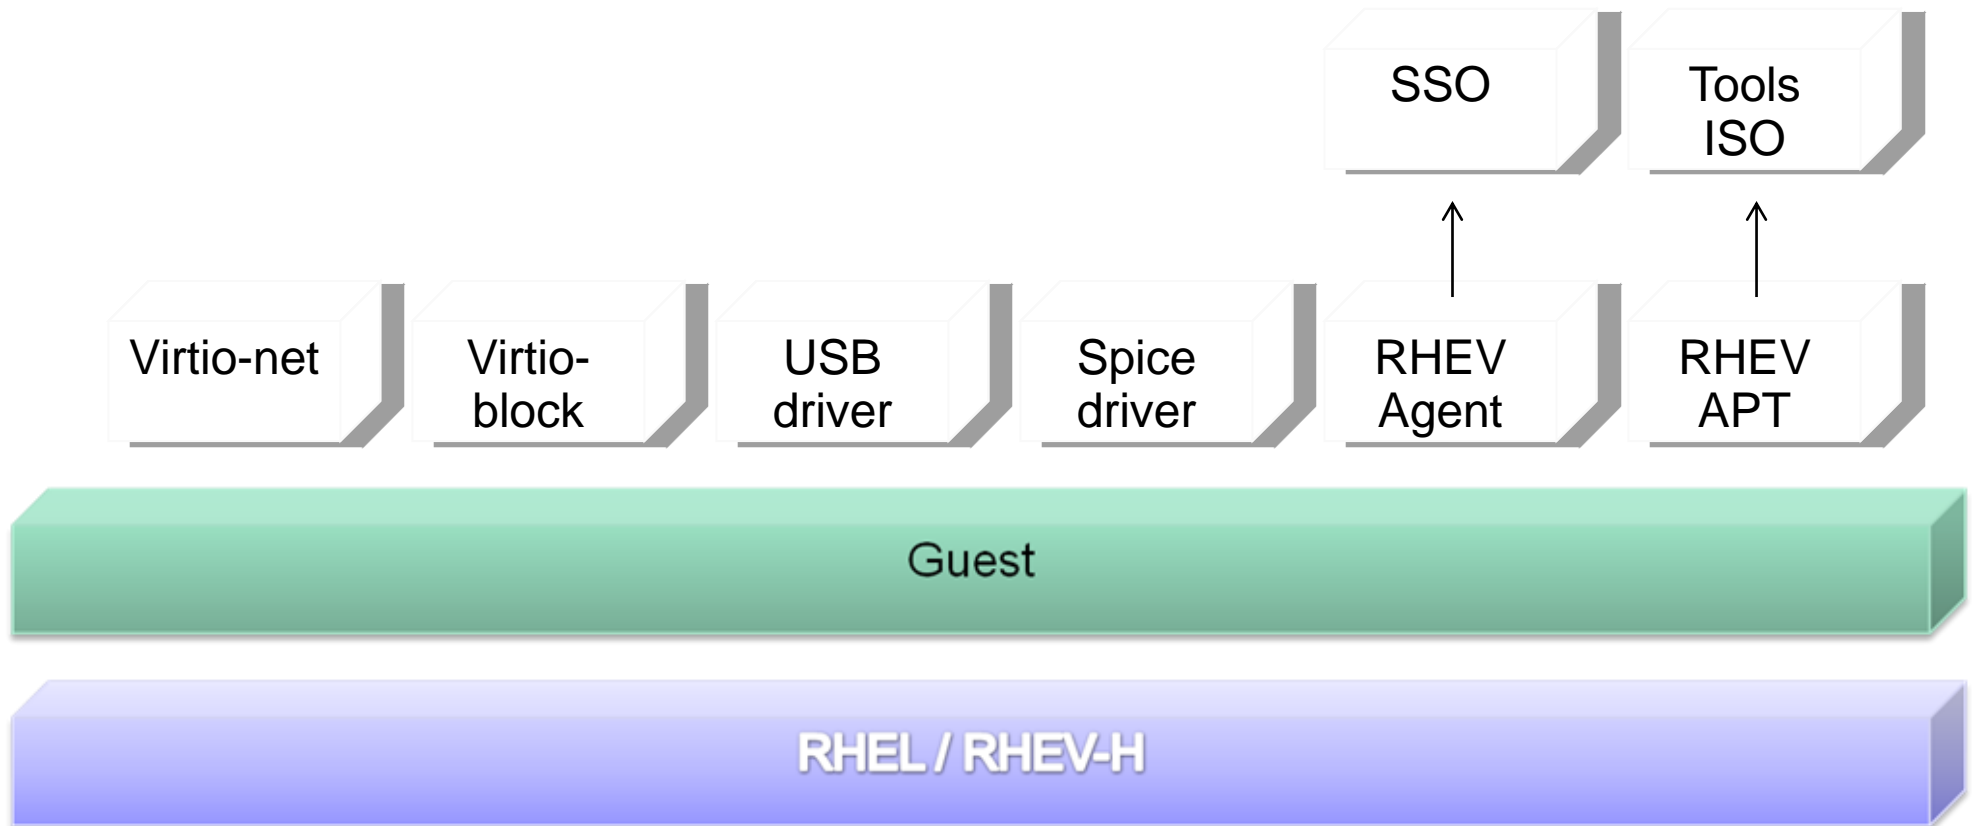


RHEV Network

- Managed by RHEV-M - no host level configuration
- Data Center → Host Cluster → Host
- Physical → [bond] → vlan/bridge (logical)
- VM Logical Network



Guest



Not all components are available for all guests

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



RHEV Guest Tools

- .Guest Agent
- .Guest stats
- .Application list
- .IP
- .Logged in user name
- .Internal shutdown and shutdown notification
- .SSO
- .RHEV Application Provisioning Tool (APT) – allows automatic install and updates of windows guest tools

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



Virtual Desktop Infrastructure (VDI)

- .User portal
- .Different use cases
 - . Non IT
 - . Various clients (win/linux/thin clients)
 - . user experience
- .Connection broker
- .Scale
- .ActiveX / XPI Spice support
- .VM Pools

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



Support Matrix

.Hosts

.RHEL/RHEV-H 5.4/5.5/6.0

.Clients

.Windows XP, Windows 7 32/64

.Internet Explorer 6-8

.RHEL 5.4/5.5/6.0

.Guests

.Windows XP/2003, Windows 7/2008/R2 32/64

.RHEL 3/4/5/6(coming) 32/64

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



Scalability

- .Tested with 100 physical hosts / 500 guests*
- .Simulated with 200 virtual hosts
- .Tested 100 guests on a single host
- .User portal workload of 500 concurrent users
- .>100 nodes in a single storage cluster

- .Already heard customers want more (2000 hosts)...

*Using SQL Server Express

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



Summary of Expected Changes

- .Backend to JBoss
- .Web Admin
- .PostgreSQL support
- .RHEL support
- .VDSM over libvirt

SUMMIT

JBoss
WORLD

PRESENTED BY RED HAT



Questions?

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT



FOLLOW US ON TWITTER

www.twitter.com/redhatsummit

TWEET ABOUT IT

[#summitjbw](https://twitter.com/summitjbw)

READ THE BLOG

<http://summitblog.redhat.com/>

SUMMIT

**JBoss
WORLD**

PRESENTED BY RED HAT

