Virtualiron

Enterprise-Class Virtualization with Open Source Technologies

Alex Vasilevsky
CTO & Founder
Virtual Iron Software
June 14, 2006

Virtualization Overview

Traditional x86 Architecture

- Each server runs single OS
- Typical “one server, one app” model
- Can contribute to low CPU utilization and server sprawl

With Virtualization

- Aimed at delivering
  - Higher CPU utilization
  - Reduction in server sprawl
  - Strong fault and security isolation, VM encapsulation

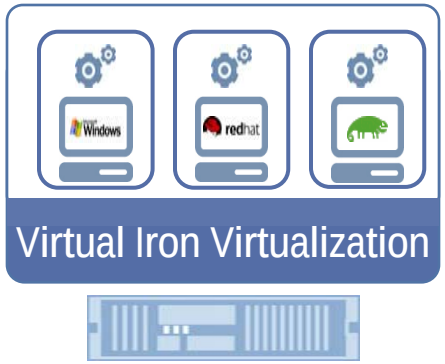

Server Virtualization Technology

Single OS Image: Virtuozzo, Vservers, Zones

- Group User processes into resource containers
- Hard to get strong isolation

Para-Virtualization: Denali, UML, Xen

- Create new “x86” virtual hw – close to physical, but not identical
- Run multiple modified guest OSes ported to new “x86” virtual hw

Full Virtualization: VMware, VirtualPC, QEMU

- Create identical x86 virtual hw
- Run multiple unmodified guest OSes – no porting required
- x86 is notoriously difficult to virtualize
- Emulation with binary rewrite (VMware)

Native Virtualization: Xen

- Newer x86 CPUs include hardware-assisted virtualization
  - VMX, VTx, VTi from Intel and AMD-V from AMD
- Create identical x86 virtual hw by leveraging hardware-assist
- Hardware support makes full virtualization possible without binary rewrite

Key Benefits of Server Virtualization

Server Consolidation

- Combine underutilized machines without conflict
- Support multiple platforms and environments through software development life cycle

Server Provisioning

- Standardize and quickly provision
- Reduce operational costs

Utility Computing

- Utilize hardware resources as needed – on demand
- Manage hardware like data – cut the ties between software and hardware

Virtualization In the Enterprise

Improve resource utilization, get more out of today's fast industry-standard hardware

Quickly and cheaply set up development, test, and production environments

Recover from failures quickly, reliably and cost-efficiently

Match workloads with available capacity to optimize efficiency and manage SLA's

Automate to reduce manual intervention, human errors, time and labor costs

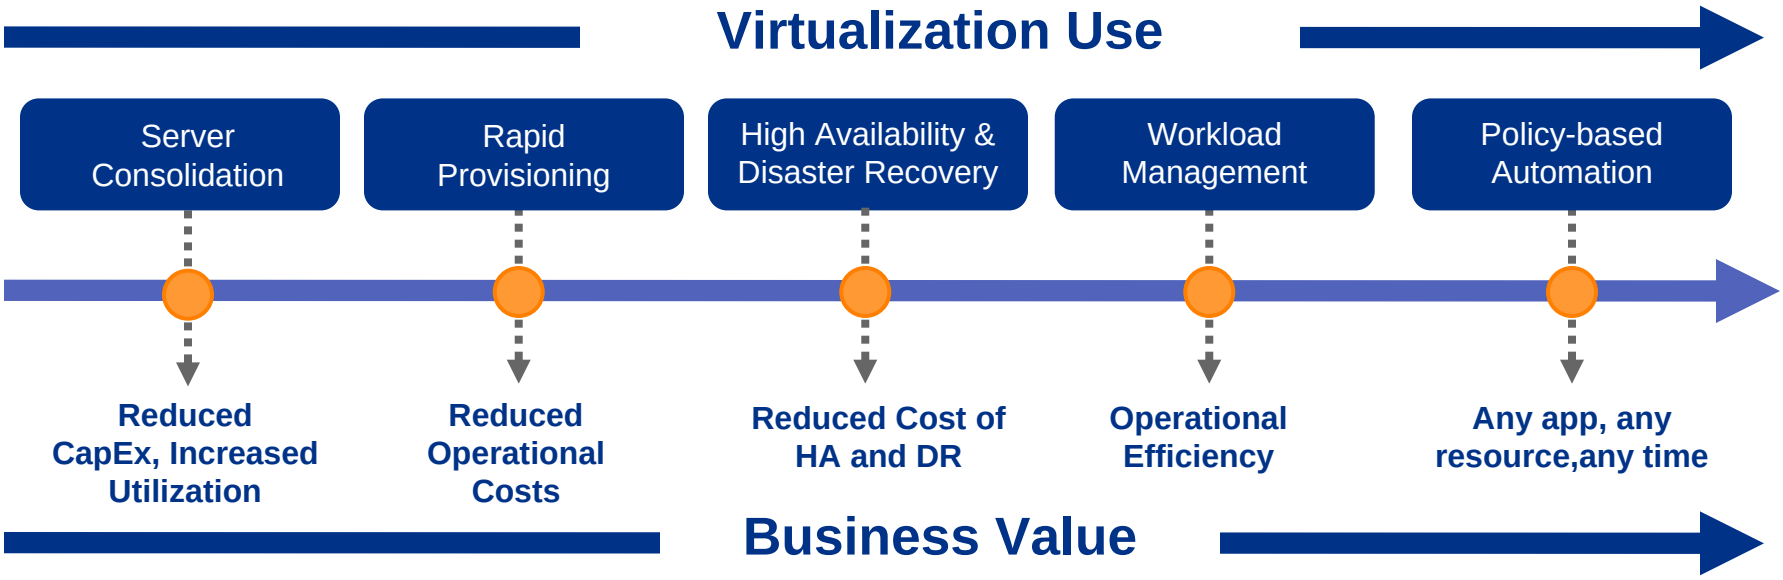

Virtual Iron Software

Founding Premise - Virtualized “Dynamic Infrastructure” for Industry-Standard Servers

- Reduce cost and complexity involved with managing and operating the data center
- Improve manageability, maximize utilization, and increase business agility

Enabling Server Virtualization Based Upon Open Source Technologies

- Open industry standards – Xen hypervisor
- Open source catalyzes adoption of server virtualization

Delivering Enterprise-class Virtualization

- Leverage advance virtualization and management as key enabler for consolidation, utility computing and service-oriented architectures

Investors

HIGHLAND CAPITAL PARTNERS

MATRIX PARTNERS

Goldman Sachs

intel

SAP

Virtual Iron's Solution - Virtual Infrastructure

Any Workload

Hardware Independence

Any Hardware

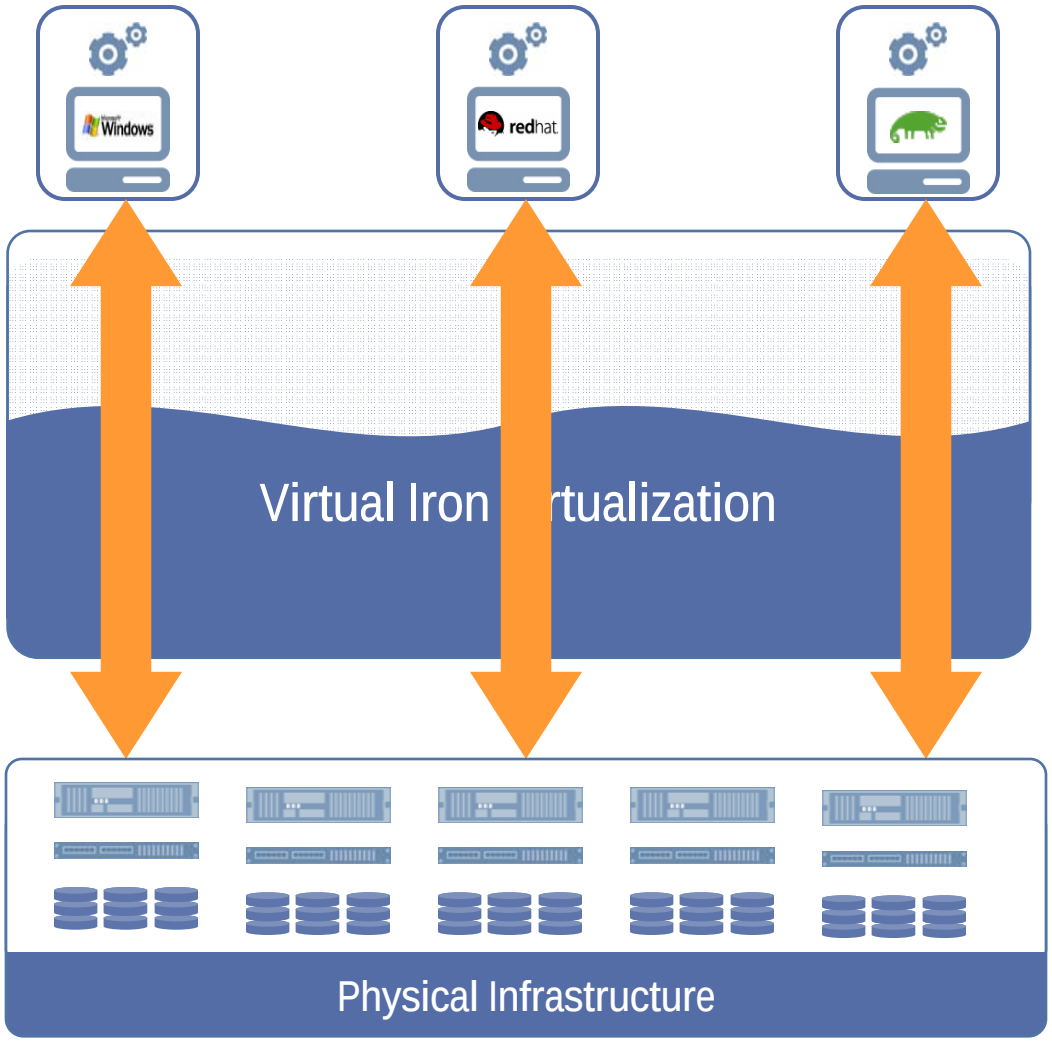

Move Live Workloads

Dynamic Mapping of Resources to Business Needs

Manage Services not Servers

Virtual Infrastructure Architecture

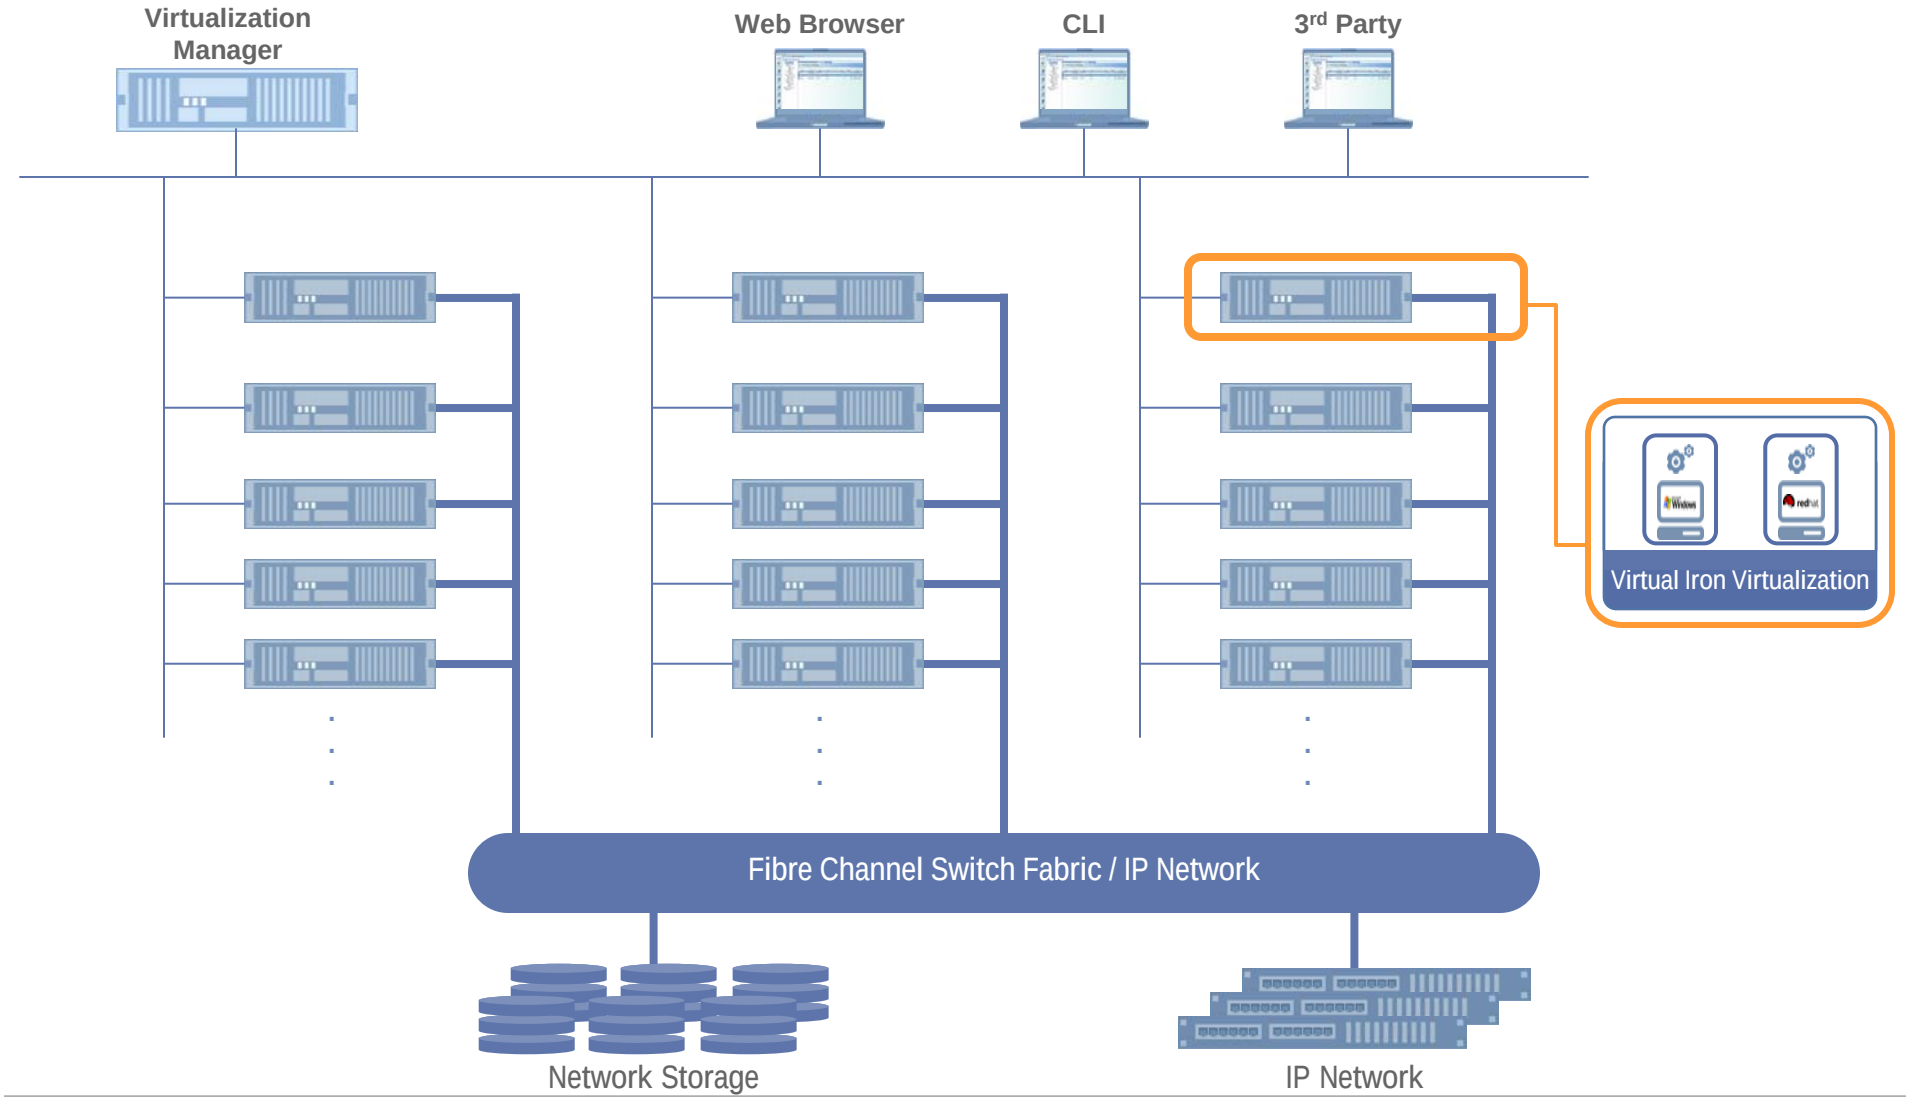

Virtual Infrastructure Architecture

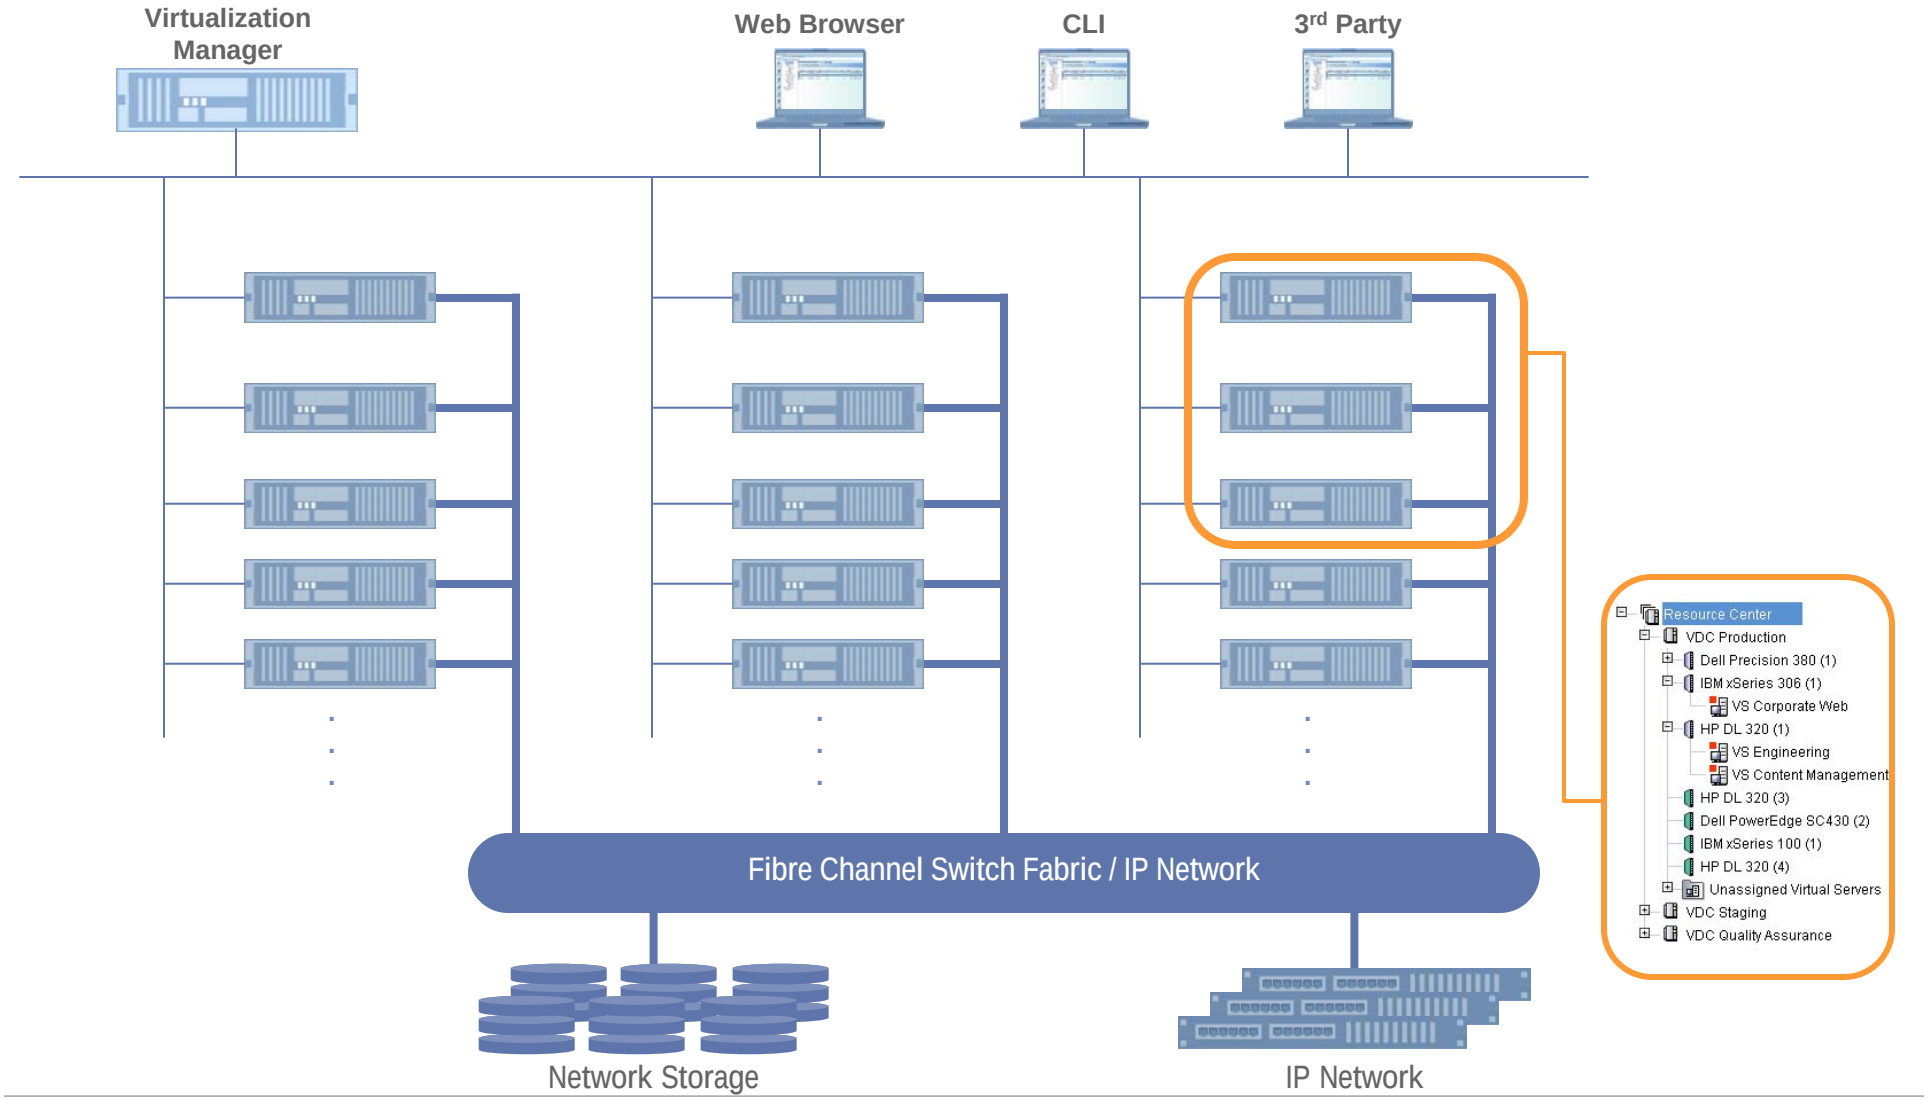

Virtual Infrastructure

- Create and manage “Software-based” infrastructure instead of “Static, Hardwired”
- Centrally manage virtual and physical resources as a shared pool
- Policy-based automation
- Self-configurable, with auto discovery and provisioning

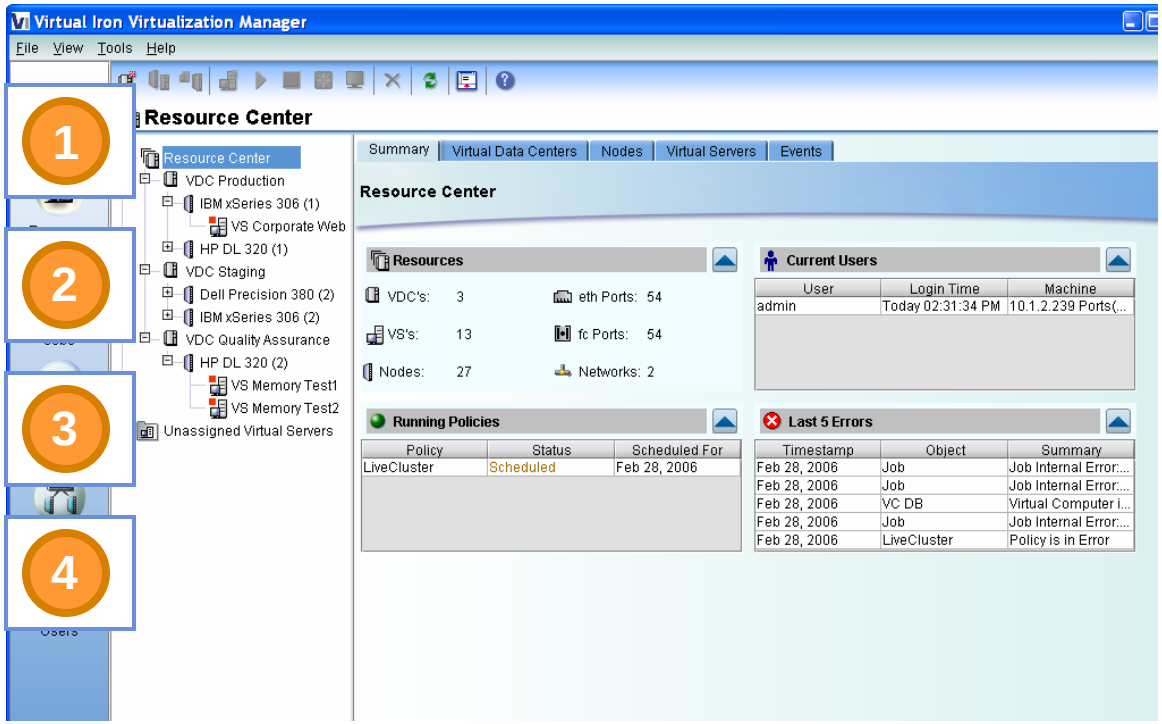

Virtual Infrastructure

Automated Resource Management

- Eliminate need for manual intervention
- Dynamic scaling
- Maintain high level of utilization
- Auto-recovery from failures

Policy-Based Capacity Management

- Monitor system utilization or application performance
- Exceeded thresholds trigger “capacity reconfiguration” rules

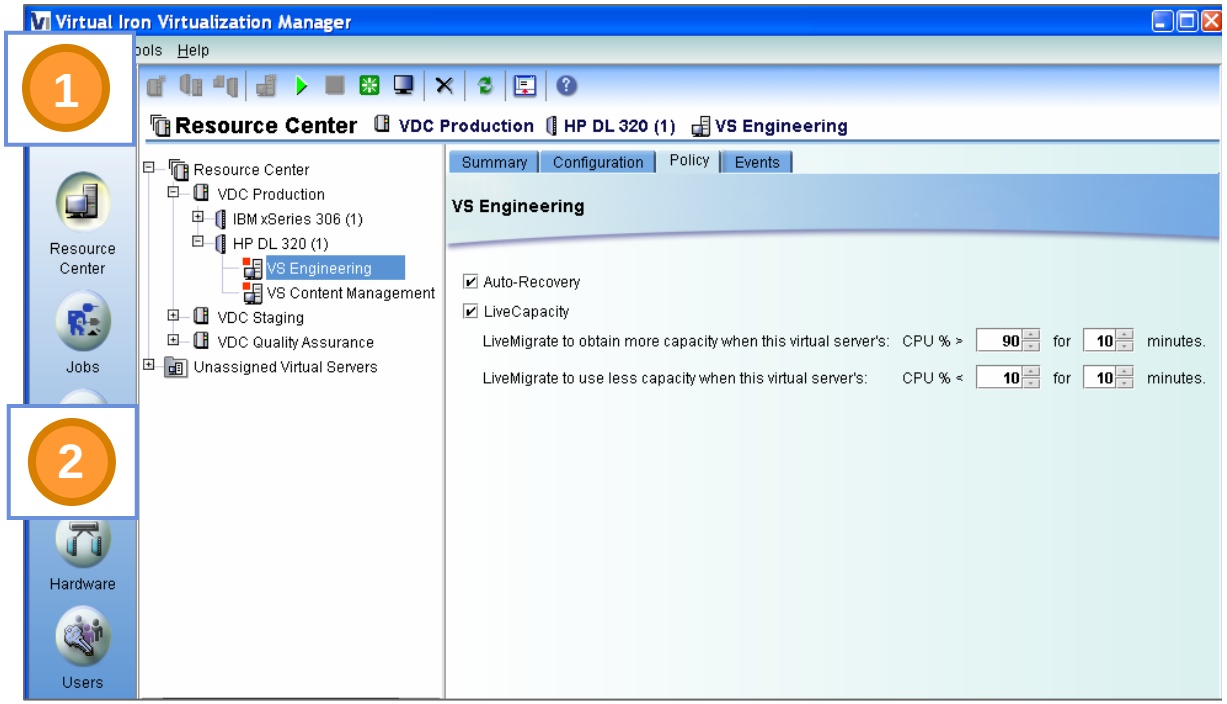

Virtualiron

Xen-based Virtualization

Virtualization Architecture

Quick Xen 3.0 Overview

Open source hypervisor

- Licensed under the GPL
- Provides secure isolation, resource control and QoS

Supports multiple processor families

- x86, x86-64, ia64 and PPC in varying degrees of maturity

Supports SMP guests

- Not for Hardware Virtual Machines (HVM) domains (Intel VT, AMD-V)

Supports VS migration

- Not for HVM domains

Widespread hardware support

Xen Architecture

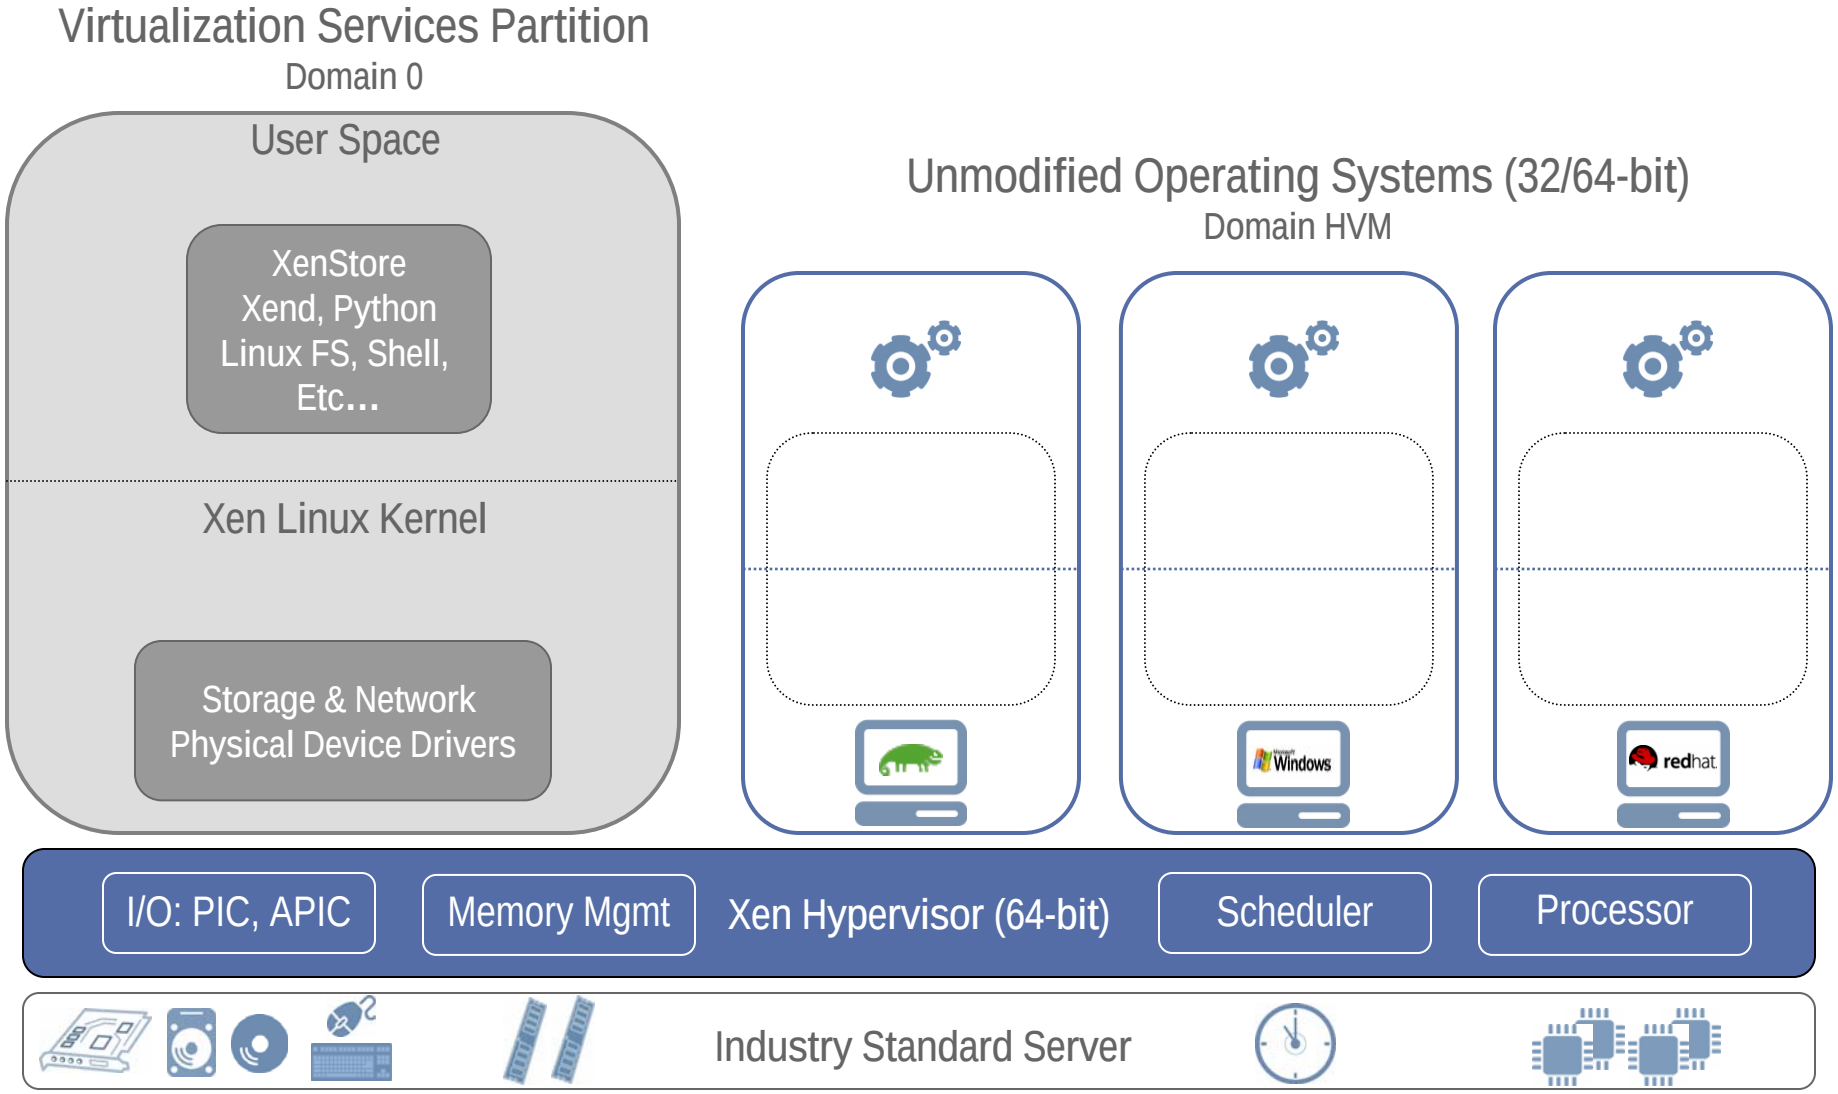

I/O Architecture

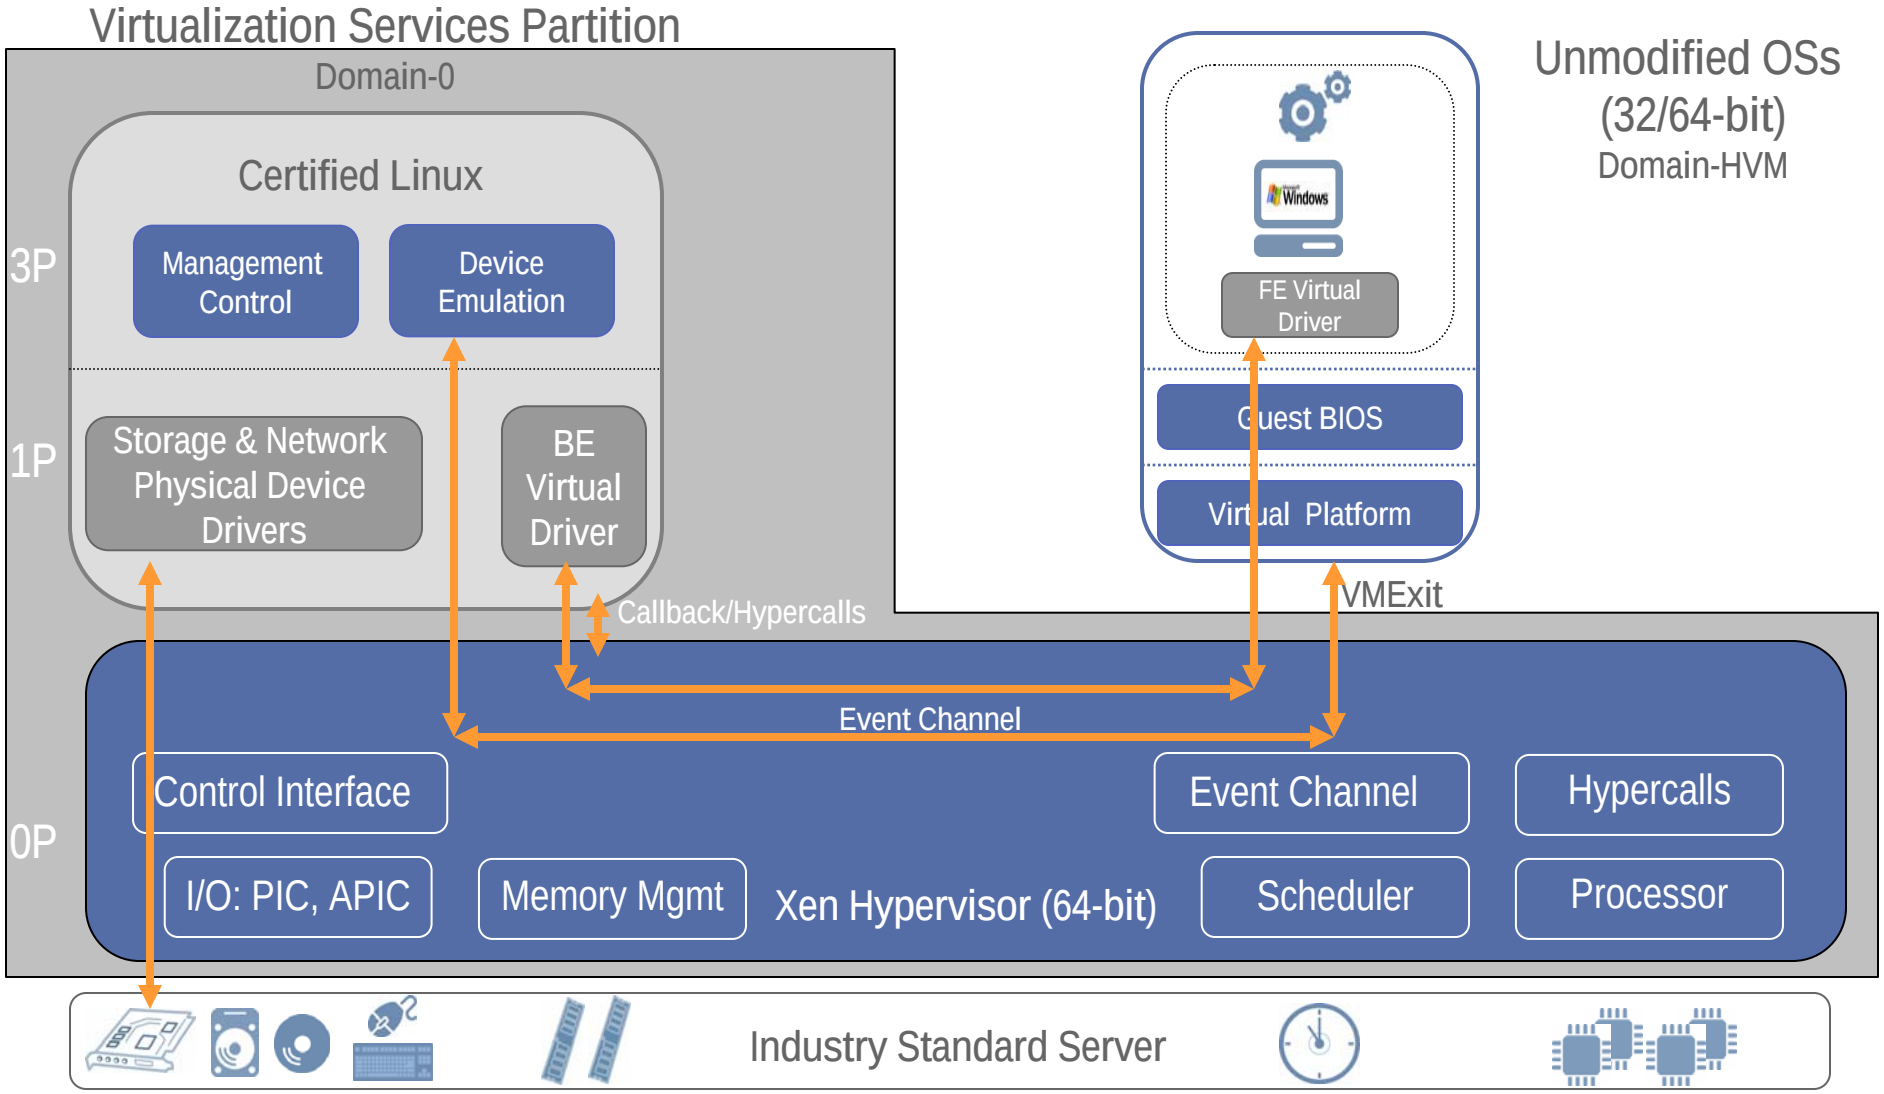

I/O Architecture – Front-ends and Back-ends

Domain-0 runs the back-end of the device, which is connected to each domain's front-end for that device

- netback, netfront for network devices (NICs)
- blockback, blockfront for block devices

Back-ends and front-ends communicate at a high level device abstraction - block class, network class, etc.

The domain doesn't care what kind of block device it's talking to, only that it looks like a block device

Ultimately, all communication between front-ends and back-ends happens in memory

Virtualization & Management Stack

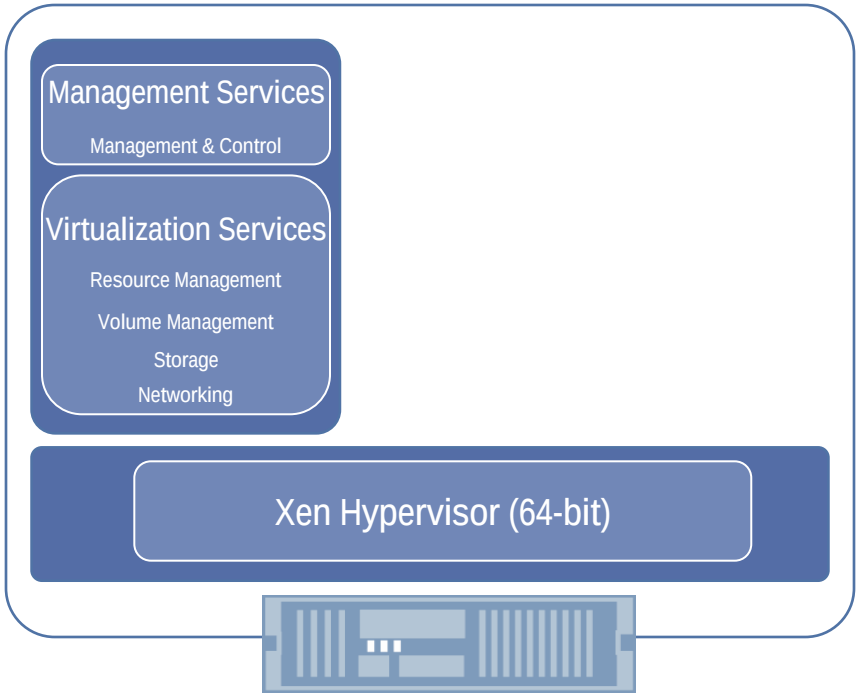

- Virtualization & Management Services
  - Full virtualization (on VTx hardware)
  - Server resource management
  - Create, control virtual servers
  - Virtual devices: networks, disks, console
  - Logical Volume Manager
  - 32 and 64 bit unmodified guest OS's
  - Dynamic virtual server migration
  - Secure, small, stateless
  - Based on certified Xen-enabled Linux kernel
- Xen Hypervisor (64-bit)
  - x64 Support (Xeon, Opteron)
  - Virtualization assist in hardware
  - Broad platform support

Virtual Servers

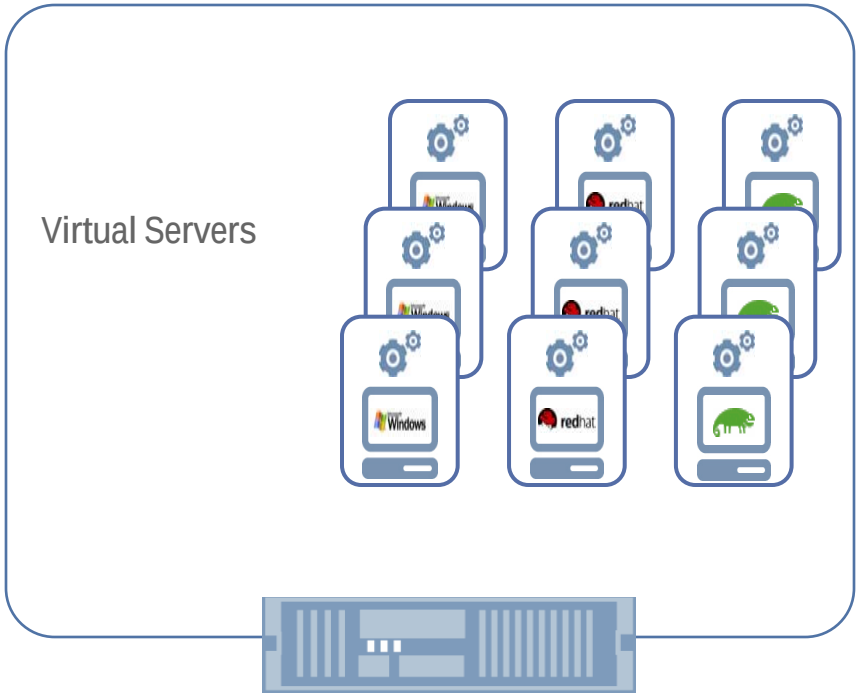

- Logical Partitioning
  - Up to 32 GOS's on a physical server
  - Fine grained resource controls
  - Full virtualization with VT & AMD-V
  - Unmodified 32 & 64-bit OS's
  - Fault containment and security isolation
- Dynamic Capacity
  - More Virtual CPUs than physical
- OS Install
  - CDROM, ISO Images, Network
- OS Boot
  - Disk, PXE, CD-ROM
- Graphical Console
  - SVGA/Keyboard/Mouse

Virtual Server Networking Options

Virtual Server Networking Options

- Isolated (no Virtual NICs)
- 1 to 8 Virtual NICs per virtual server
- Virtual NICs can be plugged into a Virtual Ethernet Switch

Virtual Ethernet Switch

- Connects Virtual Servers to external connectivity
- Connects VSs to other VSs via a private internal virtual switch

Physical Devices

- 1 to 8 physical NICs per server

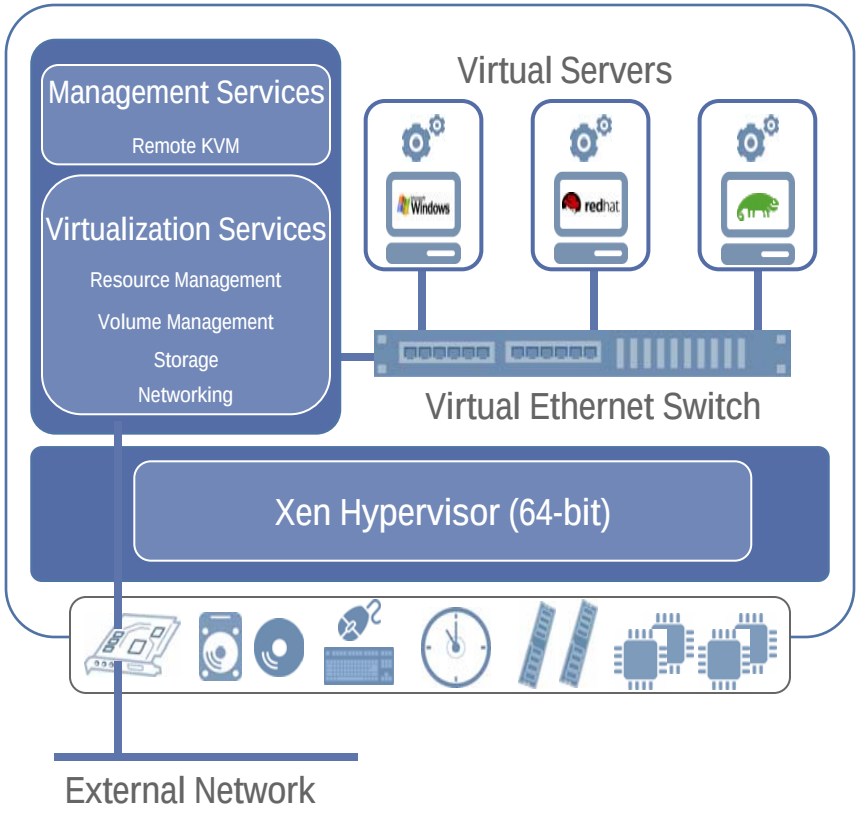

Virtual Server Storage Options

Raw Block Devices

- Access to raw LUNs
- "Pass-through" block device
- Physical to Virtual Clustering
- Clustering of Virtual Servers
- Enables array-based software operations in the guest
  - SCSI commands are passed through
- LiveMigrate is enabled for VSs with raw LUNs
  - Target server must be able to see the underlying raw LUN
- Excellent Performance

Virtual Disks

- File encapsulation (MS VHD format)
- Device transparency
- Storage type is transparent to guests
- Can be extended via LVM
  - Shrinking is not supported
- Can be easily copied, exported (just a file)

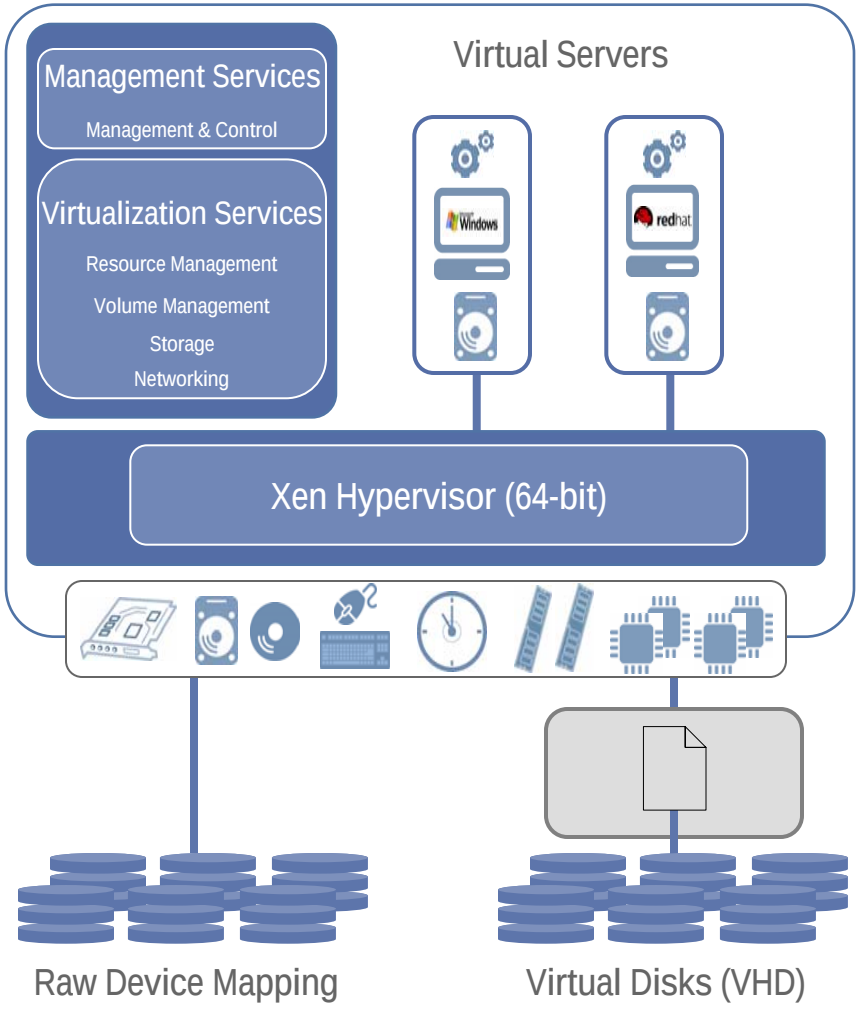

Virtualiron

Virtual Infrastructure

Virtual Infrastructure Architecture

Virtual Infrastructure Architecture

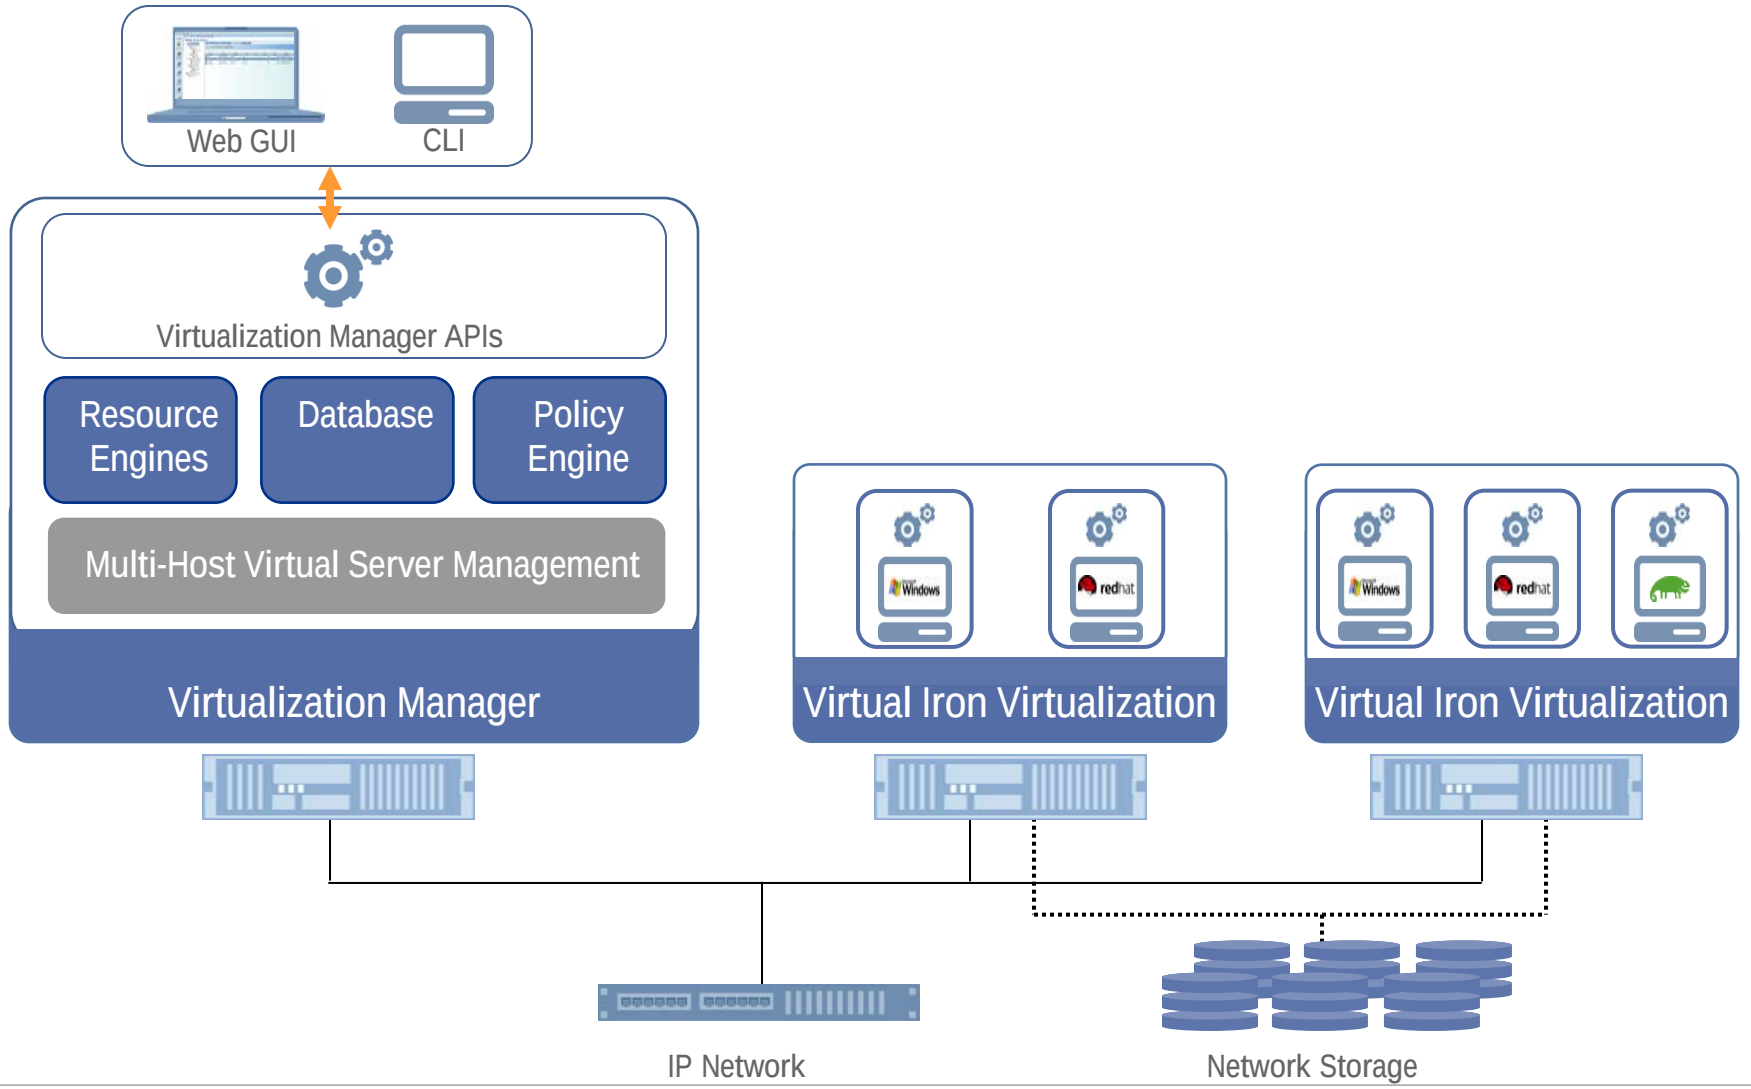

Virtualization Manager

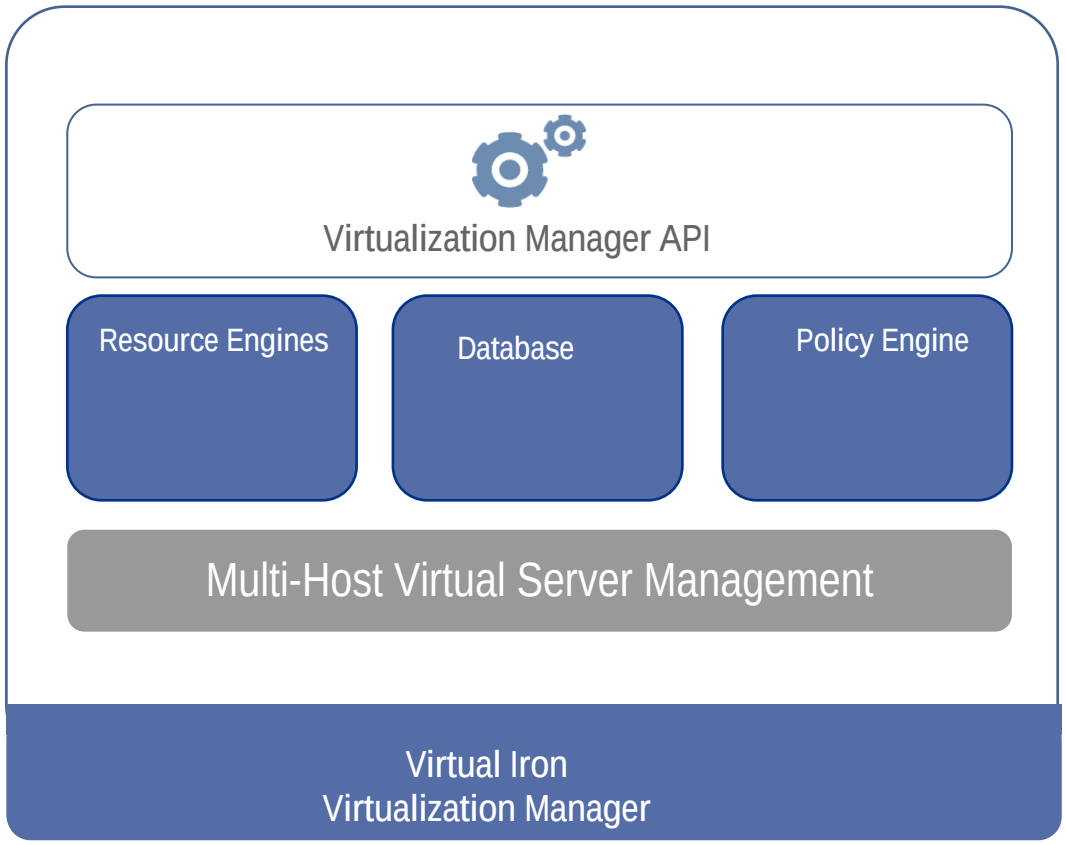

Virtual Infrastructure

- Virtual Environment Creation
- Visual Status Dashboards
- Access controls

Virtual Server Control

- Manage, Suspend, Migrate, LiveMigrate, LiveCapacity

Resource Management

- Bare metal provisioning
- Auto hardware discovery, monitoring

Policy-based Automation

- Capacity Management
- Availability services
- Rules Engine, Statistics, Event Monitor

System and Utilization Reports

- Resource utilization
- System events

Programmable (API, Scripts, Rules)

Virtual Infrastructure – LiveMigrate

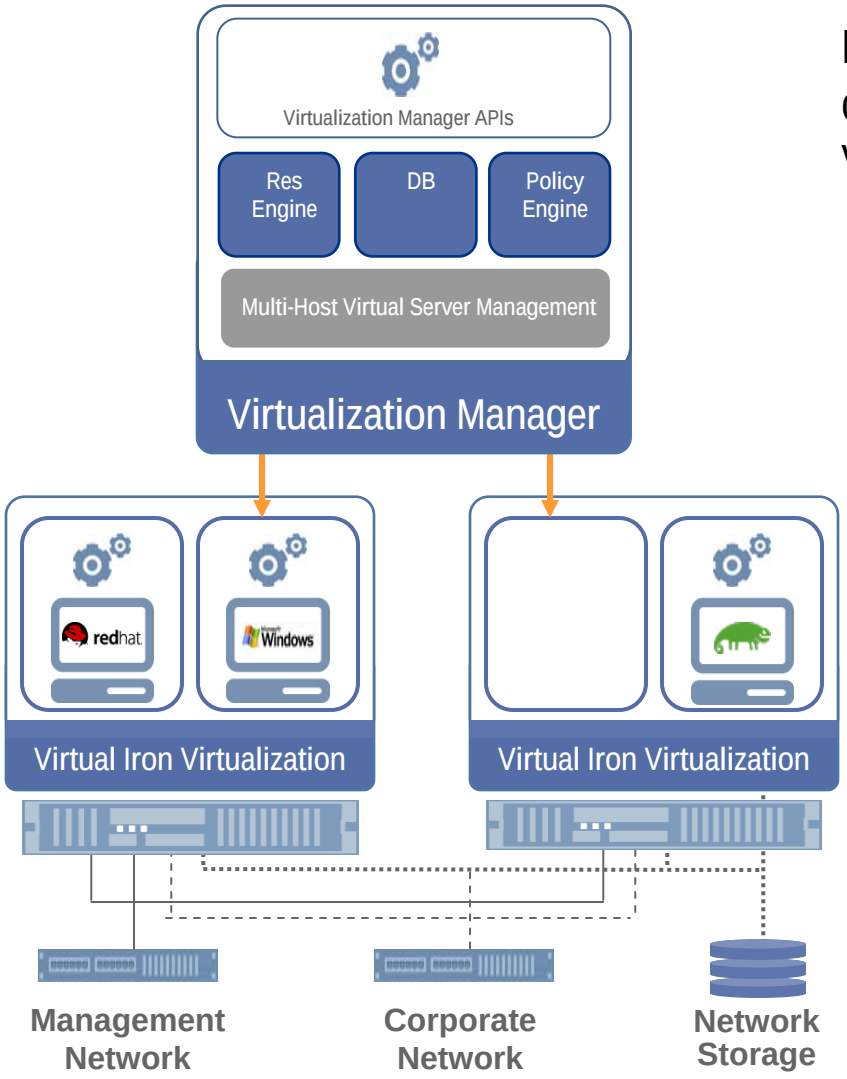

LiveMigrate moves a running virtual server from one physical server to another, co-ordinated by the Virtualization Manager

Step 1: Virtual server resources reserved on destination server

Duration: Microseconds

Step 2: Copy virtual server memory over Ethernet; disk remains on network storage; QoS of server unaffected

Duration: Minutes

Step 3: Stop original server; copy remaining memory; start destination server

Duration: Milliseconds

Virtual Infrastructure – LiveCapacity

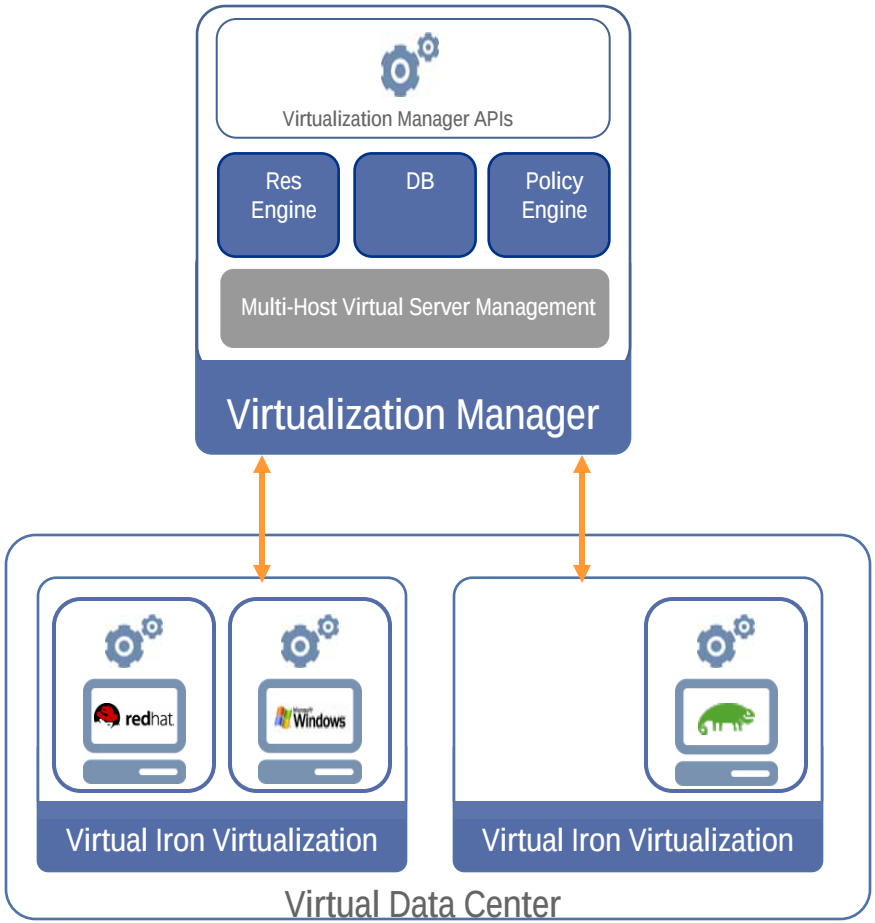

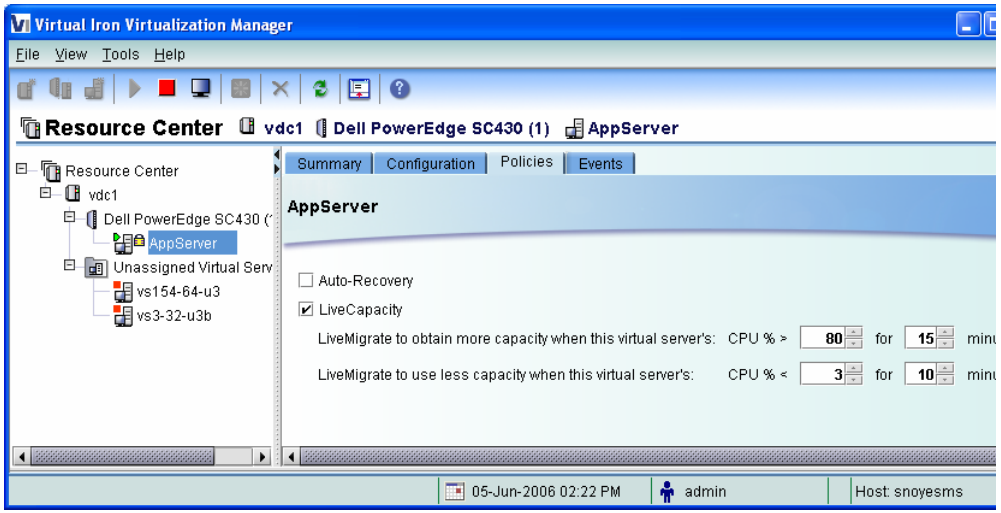

LiveCapacity

- Scheduling across a shared pool of resources
- Continuously samples performance data from every server and every VS in the VDC
- Policy driven scheduling decisions
- Uses LiveMigrate to continuously optimize
- Reacts to adding or removing servers from VDC

Virtual Infrastructure – LiveRecovery

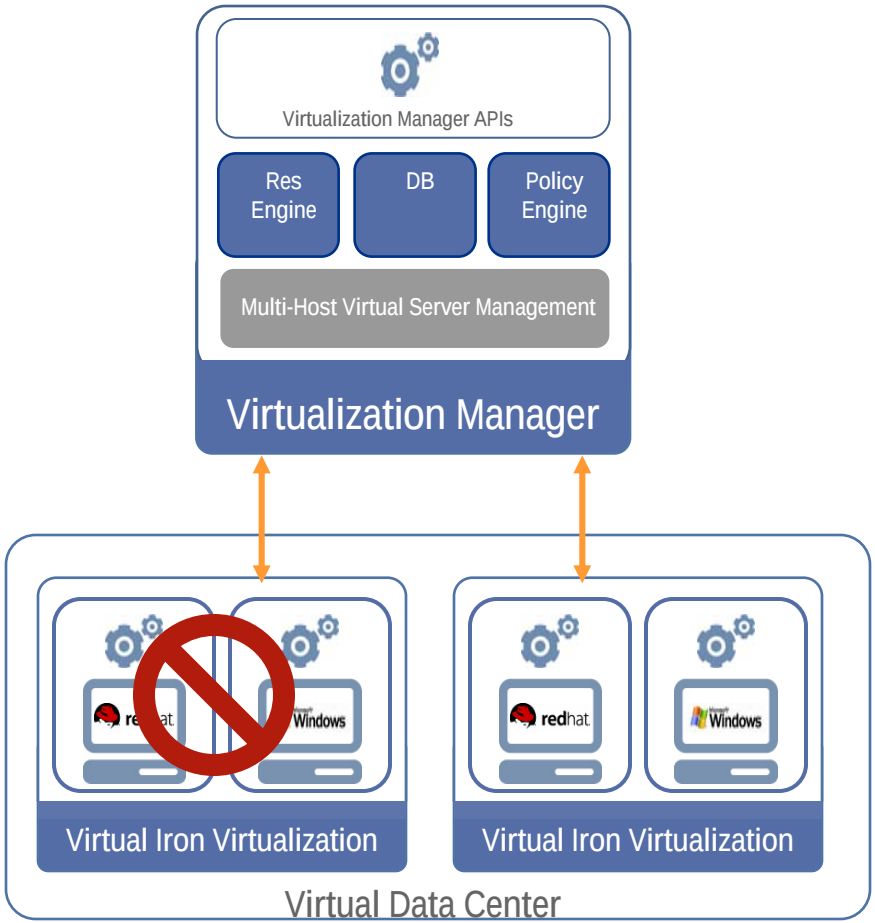

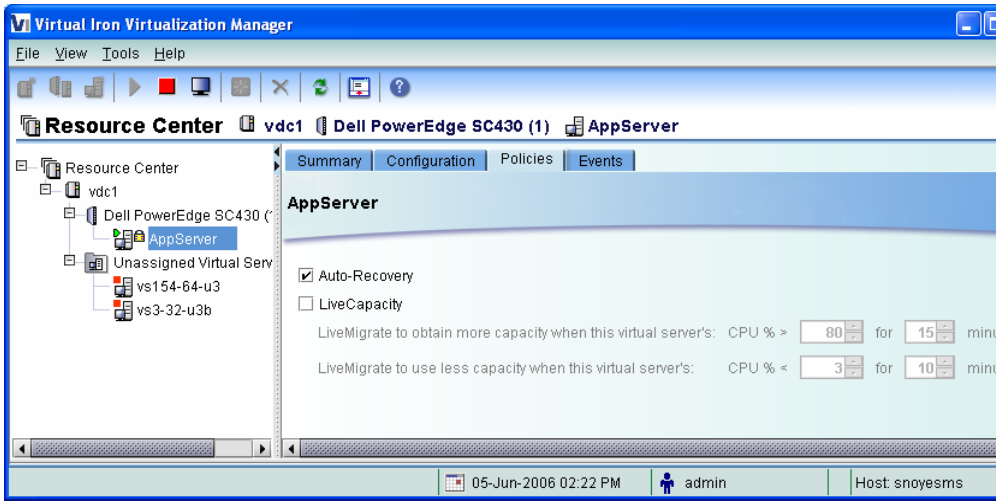

LiveRecovery - high availability for all your servers

- Losing a server in a VDC means fewer resources, not lost virtual servers
- Impacted virtual servers are relocated & restarted on remaining servers
- Placement automatically optimized
- Virtualization Manager handles all set-up and configuration
- With none of the cost and complexity of clustering