Virtualization – Beyond The Hype

GTC Southwest 2009, Seminar S11

Jim Rymarczyk
IBM Fellow, Chief Virtualization Technologist
rymar@us.ibm.com

IT Complexity – the need for IT transformation

Major Forces Are Driving the Need For IT Transformation

Costs & Service Delivery

- ⇨ Rising costs of systems & networking operations
- ⇨ Explosion in volume of data and information
- ⇨ Difficulty in deploying new applications & services

Business Resilience & Security

- ⇨ Growing systems & applications availability needs
- ⇨ Security of your assets & your clients' information
- ⇨ Landslide of compliance requirements

Energy Requirements

- ⇨ Rising energy costs & rising energy demand
- ⇨ Power & thermal issues inhibit operations
- ⇨ Environmental compliance & governance mandates

Technology Advances

- ⇨ Service orientation
- ⇨ End-to-end service mgmt
- ⇨ Comprehensive virtualization
- ⇨ Converged networks
- ⇨ Solid state storage
- ⇨ IT appliances
- ⇨ Storage de-duplication
- ⇨ Many cores & threads per chip
- ⇨ Low-cost high-BW fiber optics
- ⇨ Petaflop supercomputers
- ⇨ Cloud computing services
- ⇨ Real-time data streams

The IT Infrastructure Complexity and Cost Problem

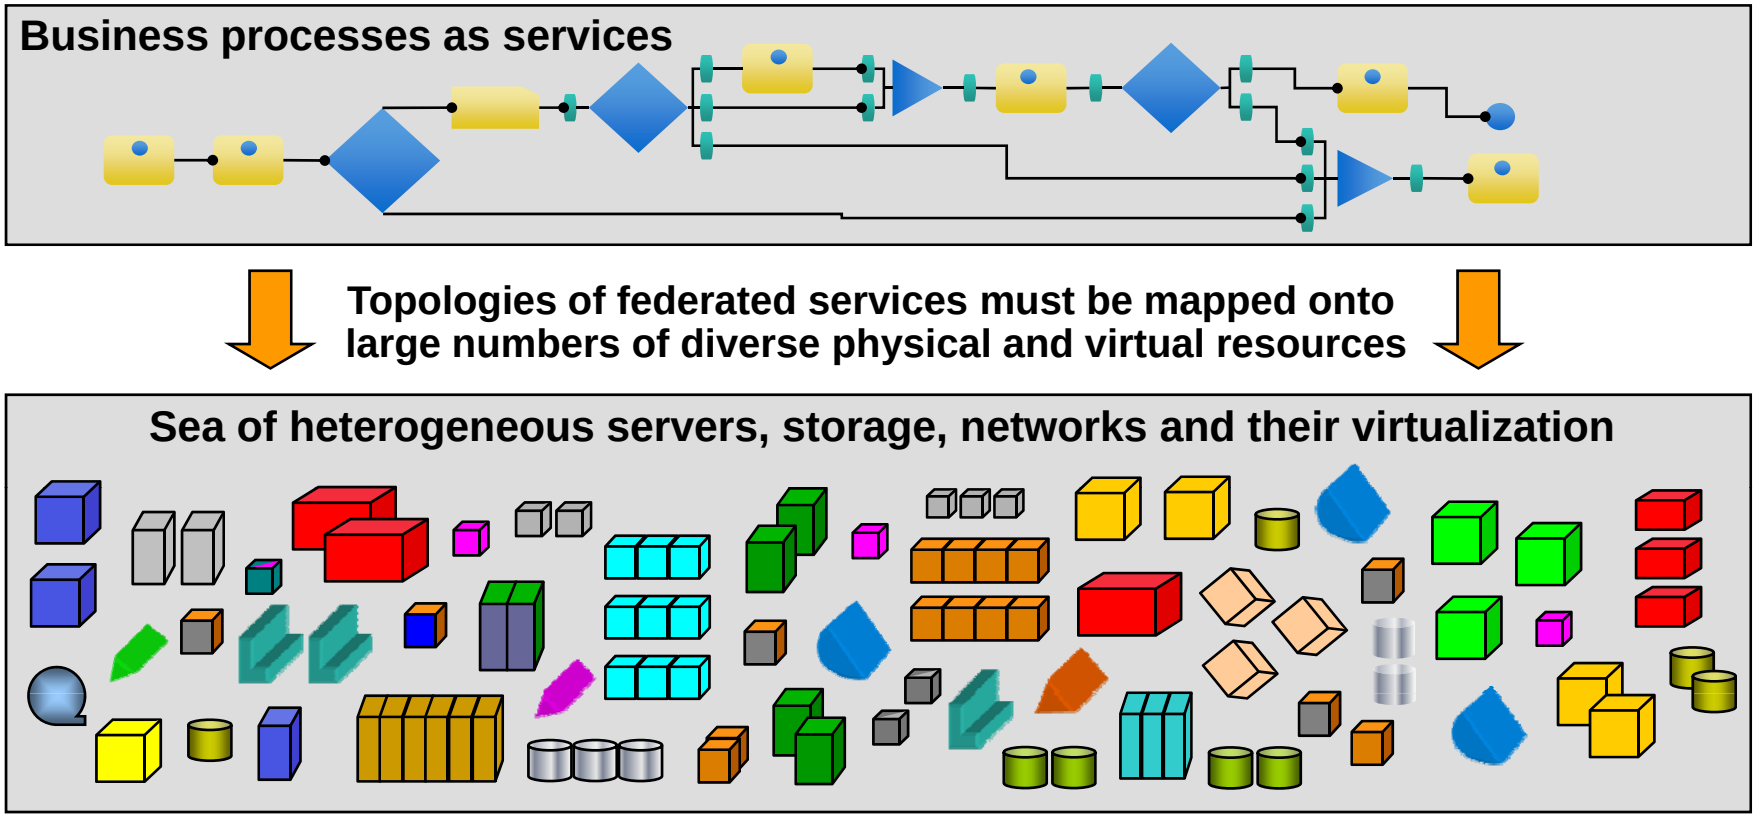

Businesses spend a large fraction of their IT budgets on data center resource management rather than on valuable applications and business processes

"Enterprises report that IT operational overhead = 70% of their IT budget and growing . . . leaving precious few resources for new initiatives."

– Forrester, 2007

Today's Enterprise Application Complexity

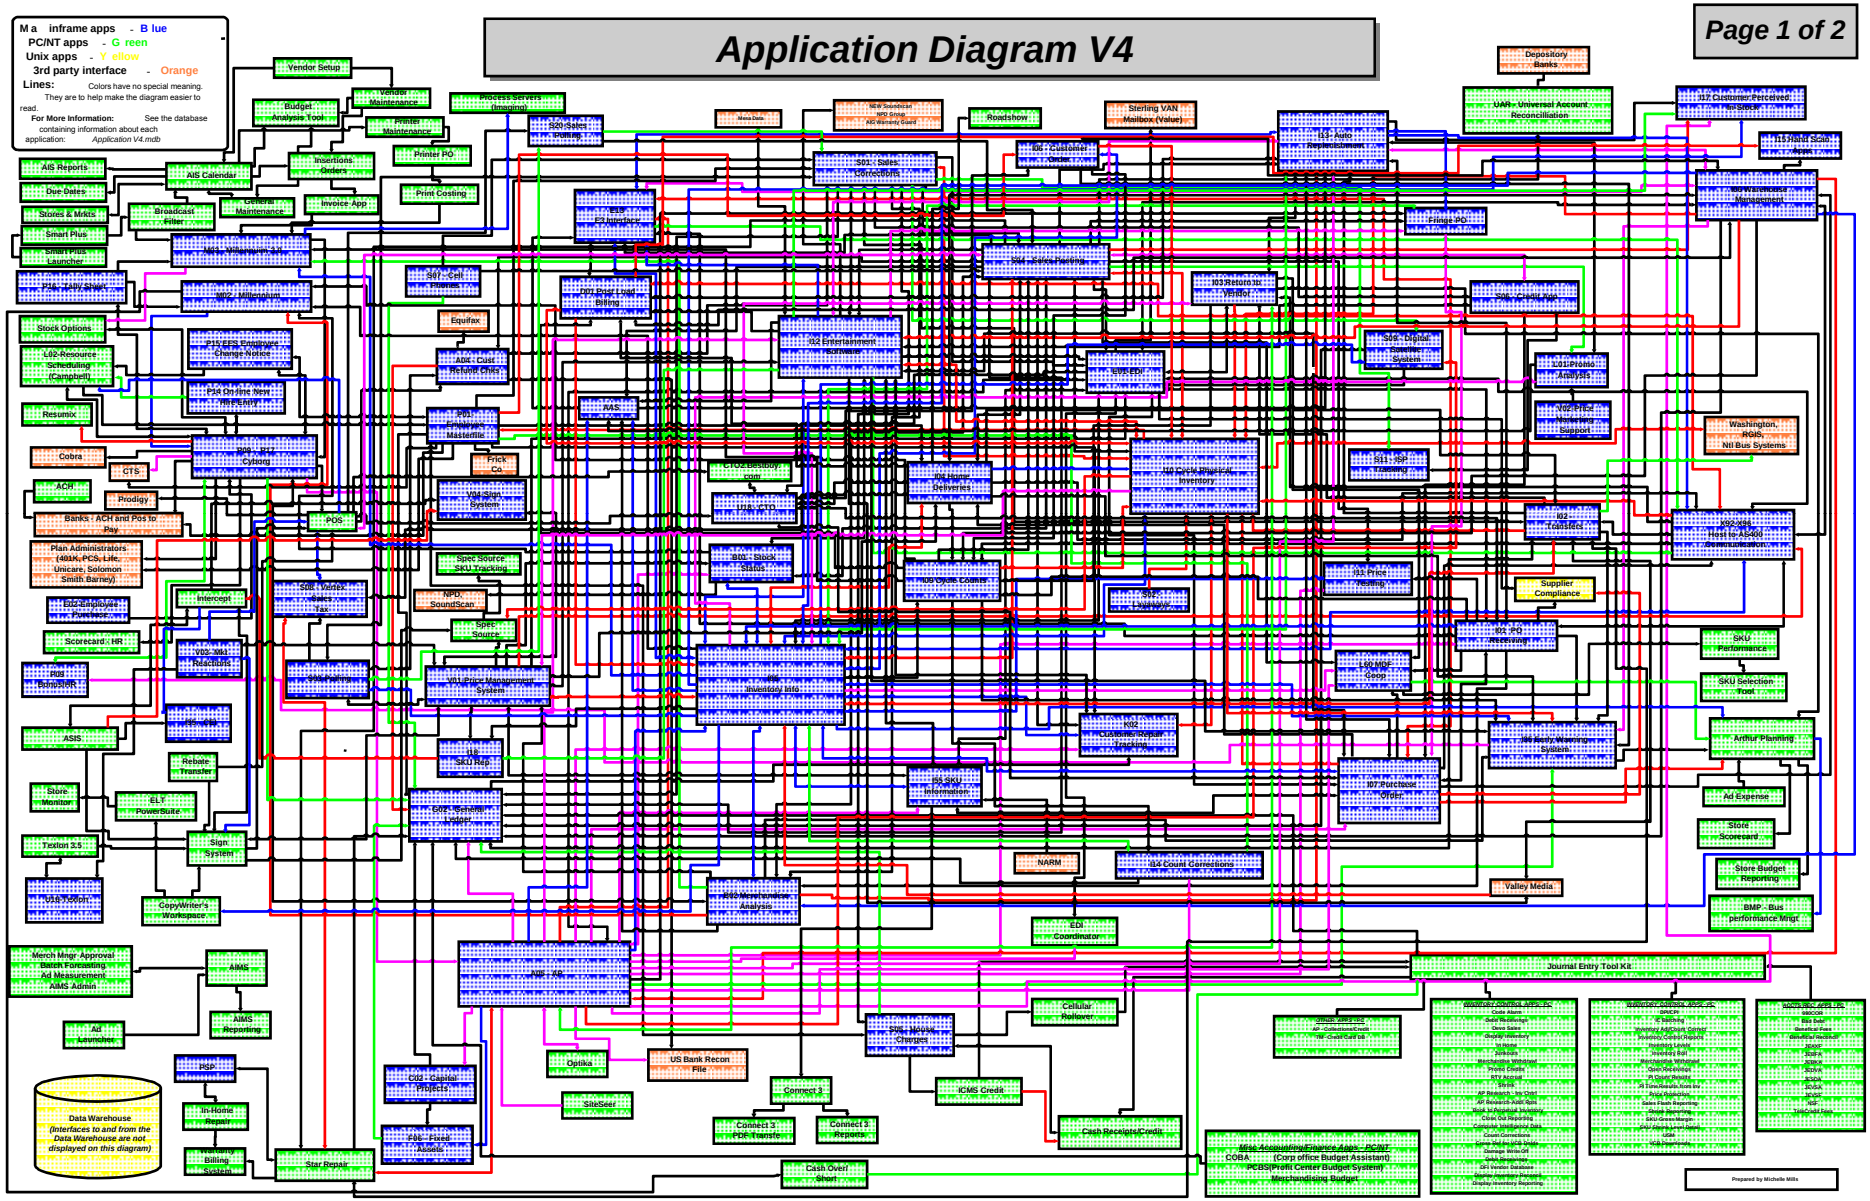

Source: IBM client engagement experience

Not shown: physical and virtual resources used

Basic Forces Will Drive Increases In IT Complexity

- **The numbers of systems deployed will continue to grow rapidly, driven largely by:**
  - New applications (Web-based apps, surveillance, operational asset mgmt., etc.)
  - Improving hardware price/performance
- **The diversity of IT products will increase as competing suppliers continue to introduce new applications, systems, and management software products**
  - Today's innovations become tomorrow's legacy
- **The coupling of IT components is extensive and increasing, driven by application tiering, advances in high-performance standard networks, …**
- **The virtualization of resources will affect existing IT processes and can lead to virtual server sprawl**
  - Introducing virtualization can have significant "hidden costs", and requires considerable skills, planning, and discipline

Virtualization – its current and expanding IT role

Virtualization Functions and Benefits

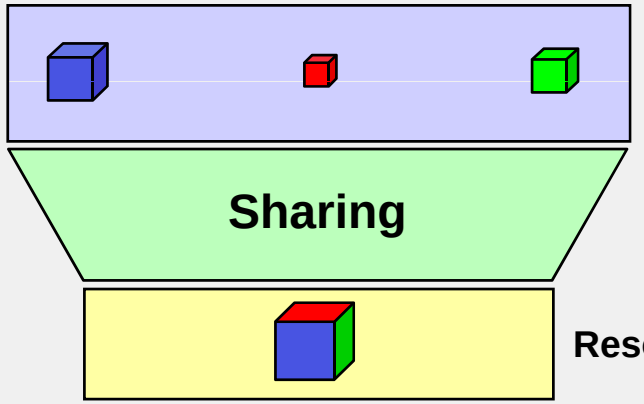

Examples: LPARs, VMs, virtual disks, VLANs
Benefits: Resource utilization, workload manageability, flflexibility, isolation

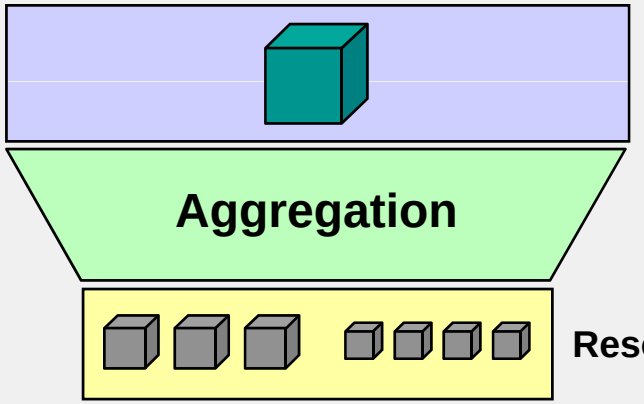

Examples: Virtual disks, IP routing to clones
Benefits: Management simplification, investment protection, scalability

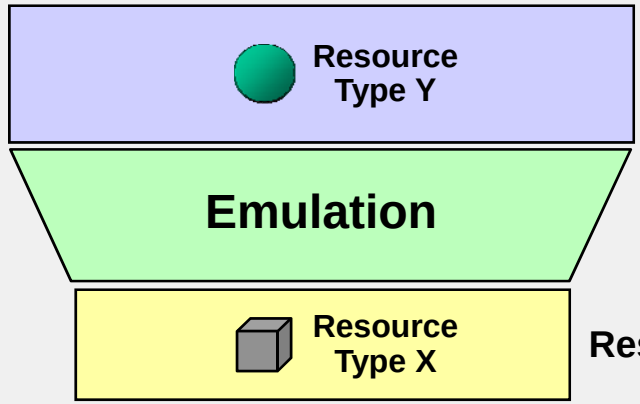

Examples: Arch. emulators, iSCSI, virtual tape
Benefits: Compatibility, software investment protection, interoperability, flexibility

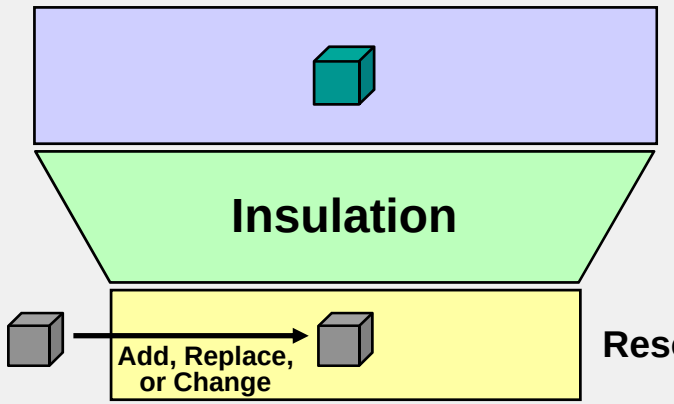

Examples: Spare CPU subst., CUoD, SAN-VC
Benefits: Continuous availability, flexibility, software investment protection

Vision: Virtualization Will Transform IT

IT Without Virtualization

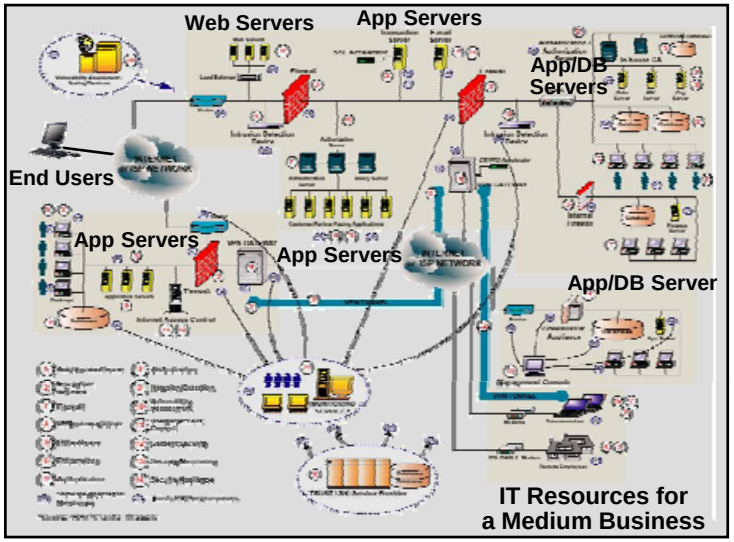

- Rigid configurations
- Fixed resources per server
- Low server utilization
- Wasted energy and floor space
- HW changes impact SW assets
- Servers managed individually

Future Virtualized IT

Virtual Environment

- Virtual resources are easier to deploy, grow, move, …
- Virtual resources, configurations, and workloads are decoupled and insulated from physical environment

Virtual Servers | Virtual Storage | Virtual Networks | Virtual Clients

Virtualization

Decouples Virtual and Physical Environments

Physical Environment

- Physical resource changes can be made without impact to running IT workloads
- Improved HW utilization and energy efficiency

Blades | SMP Servers | Storage Servers and Storage | Network Hardware

- **Virtualization technologies and emerging management software will significantly improve the efficiency of IT data centers.**

Server Virtualization Approaches

Hardware Partitioning

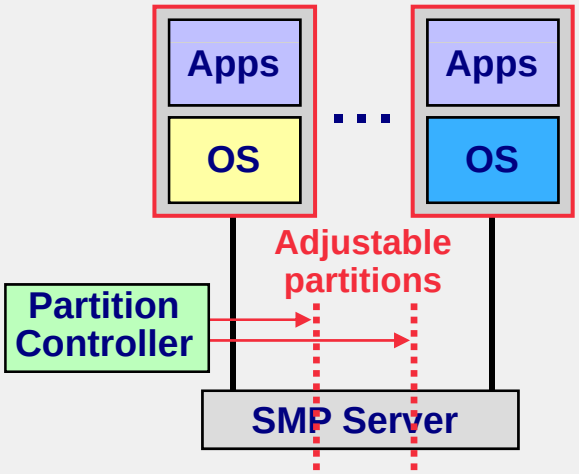

Server is subdivided into fractions each of which can run an OS

– Granularity coarse but improving
– Number of virtual machines is limited but may be sufficient
– Hardware fractions become dedicated to partitions
+ Board-level partitioning may provide electrical isolation

• Dying out as an approach

Bare Metal Hypervisor

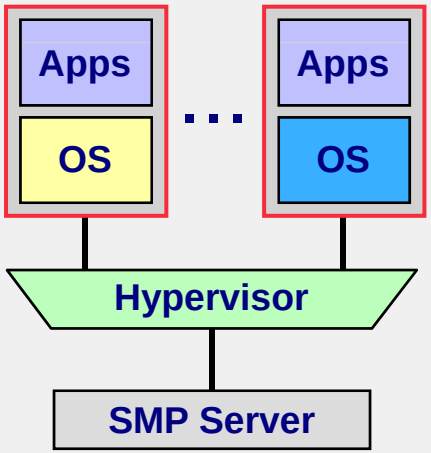

Hypervisor software/firmware runs directly on server

+ Granularity may be very fine
+ Sum of virtual resources may far exceed physical resources
+ Dynamic timesharing of all physical resources
+ Virtual machines may have dedicated hardware resources

• Will be pervasive on servers and integrated as firmware
• May play role on client devices

Hosted Hypervisor

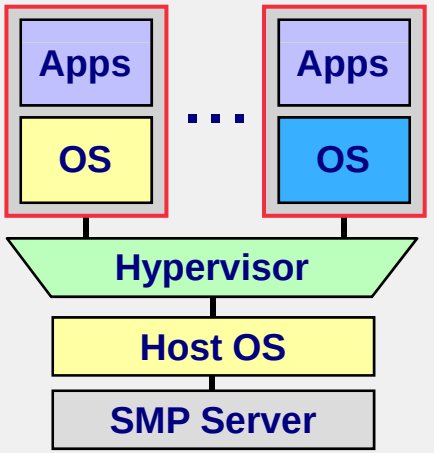

Hypervisor software runs on a host operating system

+ Granularity may be very fine
+ Sum of virtual resources may far exceed physical resources
+ Hypervisor can exploit OS services and device drivers
– Host OS is large single point of failure and overhead factor

• Will continue to be valuable as client-OS-based function

Server Virtualization Hierarchy

- Most special purpose
- Most lightweight
- Most HW independent

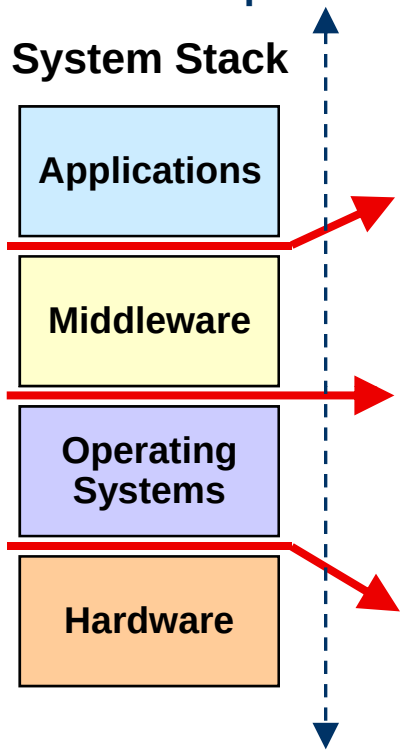

- Most general purpose
- Most heavyweight
- Most HW sensitive

Virtual Resources	Implementation Methods
• **Virtual runtimes** (application containers)	• **Middleware provides JVM, J2EE, or CLR application containers** • **Multiple middleware instances act as one**
• **Virtual operating systems** (application containers)	• **OS creates virtual OS environment per app.** • **Each container has its own name space, files, root, …**
• **Virtual servers** (virtual machines / LPARs) • **Virtual I/O** • **Virtual networks** • **Virtual devices** (CPUs, memory, I/O adapters, …)	• **Hypervisors** • **HW partitioning** • **Virtual I/O Servers** • **Self-virtualizing I/O adapters** • **In-memory VLANs**

- **There are three levels in the system stack at which virtualization is often done: (1) hardware virtualization, (2) OS virtualization, and (3) middleware virtualization**
- **Hardware virtualization is the most general purpose, but higher levels can be more efficient, more hardware independent, and even OS independent**

Storage Virtualization Approaches

Server Based

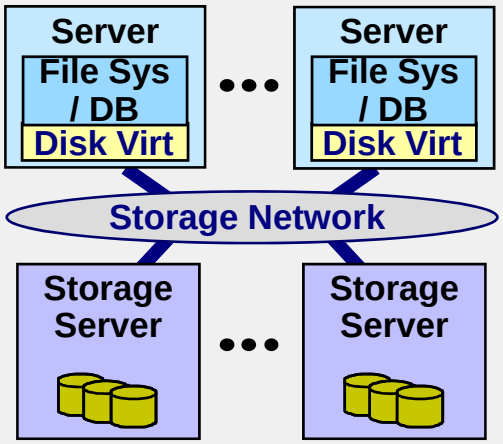

- Virtualization is done in servers by OS, hypervisor, or MW
- Supports heterogeneous storage servers

Storage Server (Array) Based

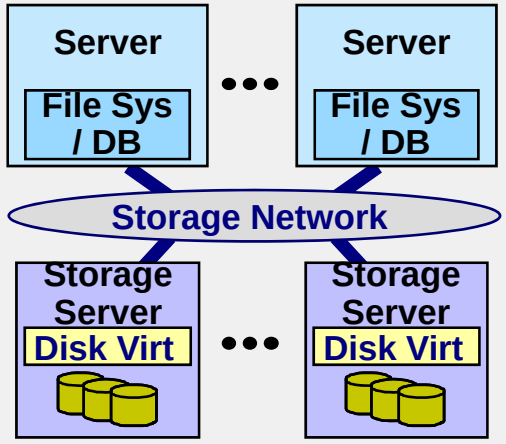

- Virtualization is done in storage servers
- Supports hetero-geneous servers
- Virtualization pools typically limited to single stg. server

Network Based via In-Band Switch

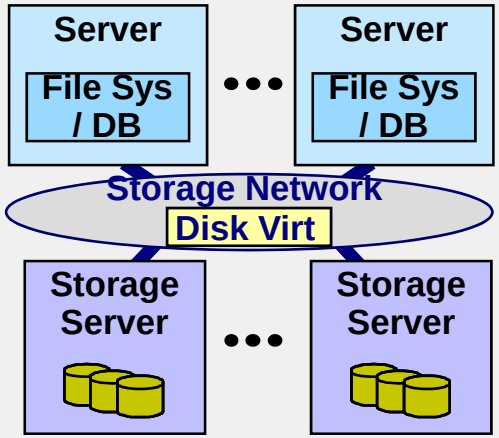

- Virtualization is done in switches
- Mapping tables for direct accesses

Network Based via Intermediary Server

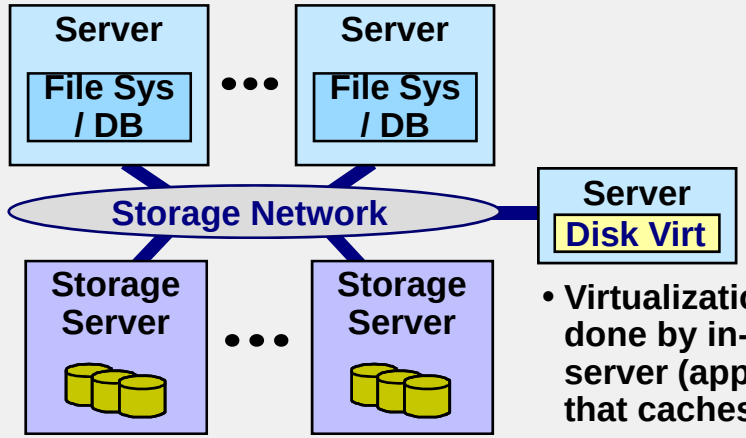

- Virtualization is done by in-band server (appliance) that caches data

- Server-based, storage-server-based, and appliance methods will continue to grow in use
- Service management software will need to support the multiple implementations
- The in-band switch approach will continue to die out

Network and I/O Virtualization

- **Virtualization of networks and I/O resources will increase**
  - Virtual networks and virtual I/O adapters are becoming virtual resource objects that can be provisioned, moved while active, … like their server and storage counterparts
- **Networks and I/O adapters will provide their own self-virtualization capabilities**
  - I/O adapters are implementing NPIV, SR-IOV, … functionality, providing controlled sharing similar to what has been done on mainframes since the 1980's and on Power Systems since mid-2007.
  - Ethernet will support dynamic VLANs and virtual Fibre Channel networks (VSANs)
  - Security and access control capabilities will be built-in
- **Servers will continue to provide internal virtual networks which may exploit server memory and/or optimized cluster-local networks**
  - Server virtual networks will continue to offer significant performance and cost benefits
  - Server virtual networks are implementing management, control and dynamic plane functions, to enable automation and easier integration with external network managers.
- **Hierarchical service management software will coordinate the virtualization of networks and I/O resources, throughout and across entire data centers**
  - Integrated solutions will be used in more homogeneous configurations
  - Adoption of rip-and-replace and break-new-ground solutions will be limited by economic realities

Expanding IT Role Of Virtualization

Virtualization with new management software will provide major new IT benefits

- The industry is investing to make this happen and it will unfold over the next 3-5 years
- The resulting high value to customers will render non-virtualized IT obsolescent

Future: Virtual IT Configurations

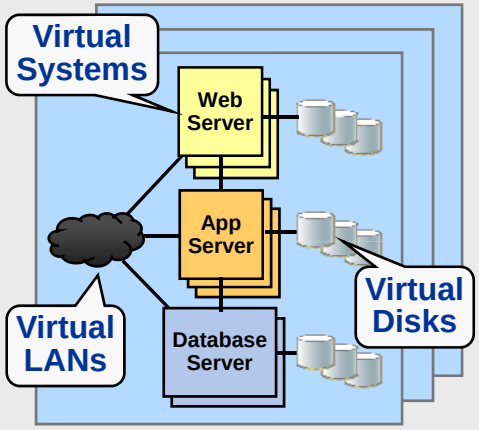

Virtual resources are fully self-describing objects

- **Creatable / adjustable** dynamically
- **Movable** while active
- **Durable** over HW generations
- **Storable / versionable** in libraries
- **Distributable** as appliances

Pools of physical systems are manageable as single systems but more scalable, resilient, …

Emerging Benefits:

- **Reduced IT management costs**
- **Better SW investment protection**
- **Dynamic energy optimization**
- **Simplified HA\DR/PD solutions**
- **Improved resource optimization**
- **Greater IT agility**
- **Ready-to-run packaged software**
- **Improved security foundation**

Emerging: Multi-System Virtualization Mgmt. Software

Provides virtual resource mgmt. and virtualization-based solutions spanning physical resource pools

Maturing: Physical Systems with Local Virtualization

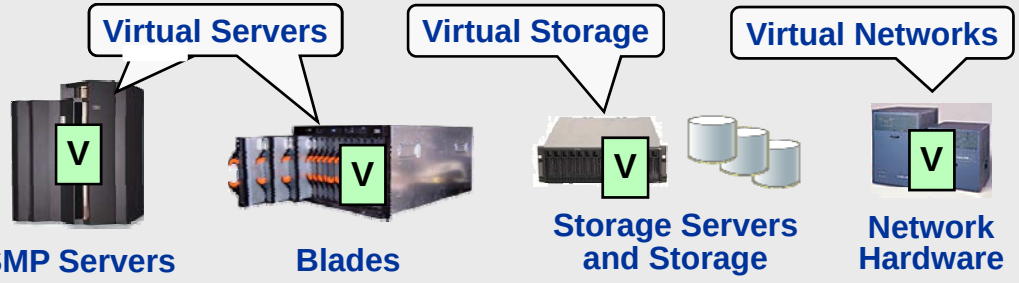

Today's Benefits:

- **Better hardware utilization**
- **Improved IT agility**
- **Lower power consumption**

Stages of IT Infrastructure Transformation

Strategic Imperatives
- **Evolve data centers from ad-hoc to well-architected**
- **Better manage the IT complexity that will continue to exist**
- **Reduce the IT complexity that needs to be managed**

Integration for Simplification

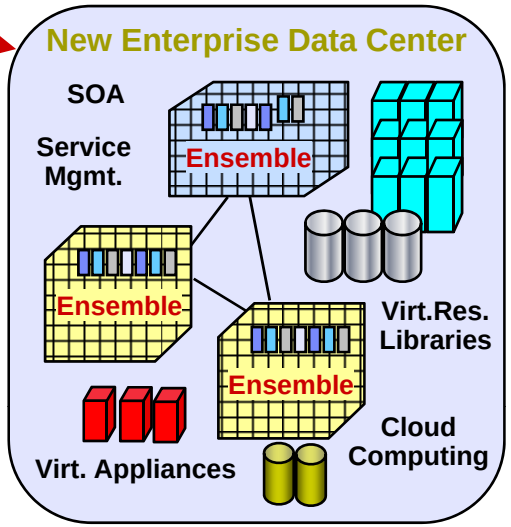

Virtualization has major benefits but introduces more complexity

Objects and Pools

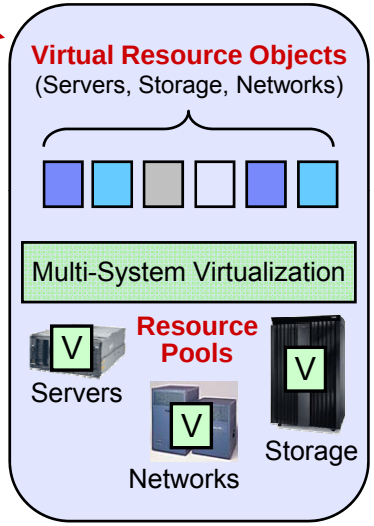

Physical Consolidation

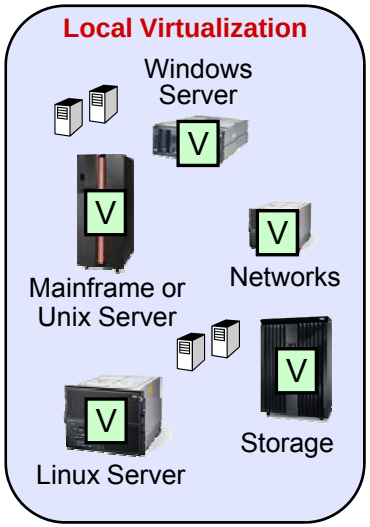

- Reduced scale-out complexity
- Integrated autonomic mgmt
- Dynamic energy optimization
- Business resilience / security

- Better SW investment protection
- Simplified HA solutions
- Improved resource optimization
- Ready-to-run packaged software

Continuing Advances

- Better hardware utilization
- Improved IT agility
- Lower power consumption

Continuing Advances (I/O, network, resilience, performance, …)

Applying Simplification First Principles to IT

- **Virtualize as many resources as possible**
  - Allows the elimination of dependencies between virtual and physical resources
- **Provide compatibility across families and generations of resources**
  - Enable virtual resources to function unchanged despite changes to underlying real resources
- **Leverage greater IT infrastructure reuse and standardization**
  - Make better use existing assets and skills to reduce overall heterogeneity
  - Promote and adopt open standards wherever possible
- **Modularize the data center**
  - Group physical resources into pools of like types
    - Provide virtual resource mobility within these pools
  - Make groups of resources look like single resources to users and resource managers
    - Hide their internal design complexities and their parallelism
    - Provide a “single system management image” per pool
- **Apply engineering first principles to the management software itself**
  - Don’t over-manage – “perfect is the enemy of good”
  - Use hierarchical, modular management software (self-managing server/ensemble tier)
  - Provide independent role-based management of virtual and physical IT environments
  - Provide consoles and tools that can be used to manage both virtual and real resources
  - Decouple data mgmt. from physical storage mgmt.
  - etc.

Key Technologies for IT Simplification

1

- **Service oriented architecture** technologies frame business processes as services, facilitating deployment, composition, reuse, modular change, agility, efficiency, …

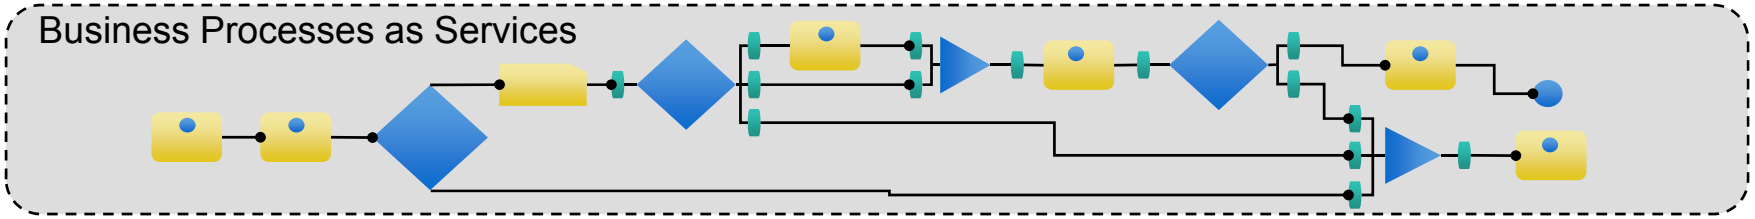

2

- **Service management** software spans the diversity of heterogeneous physical and virtual resources, providing unified cross-platform management in support of SOA
  - Services registry, service life cycle mgmt., image libraries, multi-tier composition, provisioning, autonomic optimization, network services, security services, …

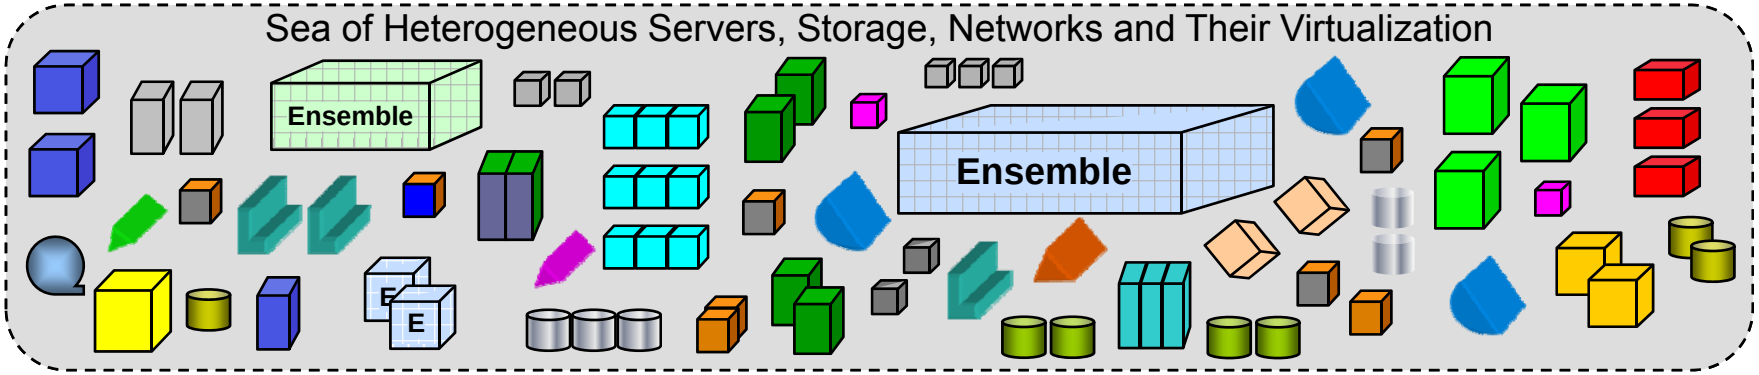

3

- **Ensembles**, scalable pools of like systems that are manageable as single systems, and traditional **scale-up servers** will replace multitudes of individual servers and reduce the labor required for physical systems management
- **Virtual Resource Objects, Virtual Appliances, Libraries,** and **Tooling** will reduce the labor required for software stack development and management

Ensemble Components

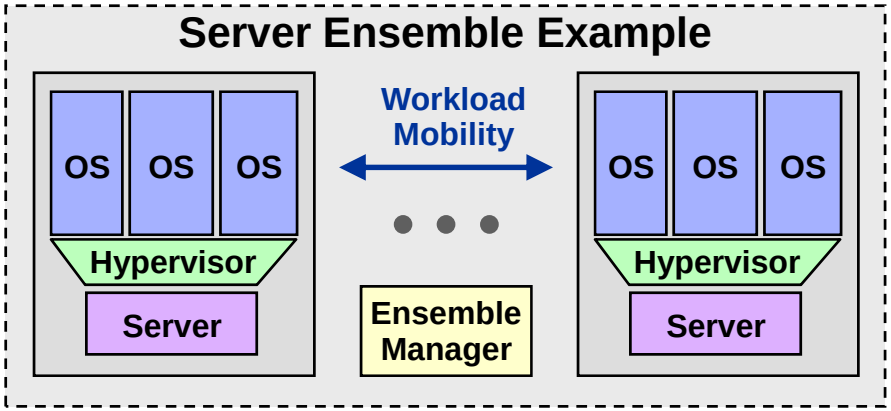

Ensemble definition:
a pool of like systems that is manageable as a single system

An Ensemble generally consists of the following components:

- A pool of compatible system nodes (e.g., N standard servers; typically homogeneous)
- Virtual resource mobility within an ensemble and with compatible ensembles
- The networks which interconnect the ensemble nodes (may be local / optimized)
- Resource virtualizers (hypervisors, I/O virtualizers, storage virtualizers, …)
- An ensemble manager appliance that provides platform management for the ensemble virtual and physical resources
- Tools for planning, ensemble creation, P2V migration, image mgmt. & composition, …
- Ensemble-local automated optimization software of performance, availability, energy usage, security, … with intelligent defaults
- Multi-system services (locking, caching, message queuing, …) may be integrated with some ensembles

Groups of Servers vs. a Server Ensemble

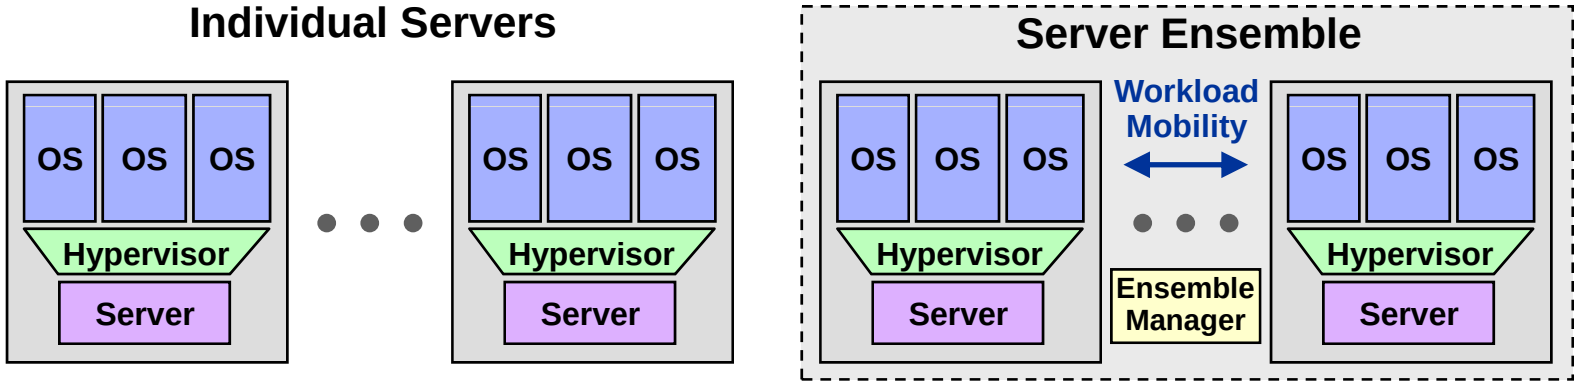

# of Things to Manage	N virtual servers; M physical servers	N virtual servers; one physical ensemble
Create, Test, and Maintain	Do it yourself; few assemblies are alike	Standard "off the shelf" assemblies
Management Automation	Add-on software, custom scripts, …	Built-in optimizations; intelligent defaults
Management Interfaces	Many individual knobs and variables	Menus of selectable standard behaviors
# of Consoles	Separate consoles for physical & virtual	Single console; in-context functionality
Data Center Mgmt. Arch.	Monolithic; spans d. c. heterogeneity	Hierarchical; pool-level modularity

Ensembles will significantly reduce IT resource management complexity and cost, and improve IT functionality, in consumable data center increments

End-to-End Customer Usage Scenario

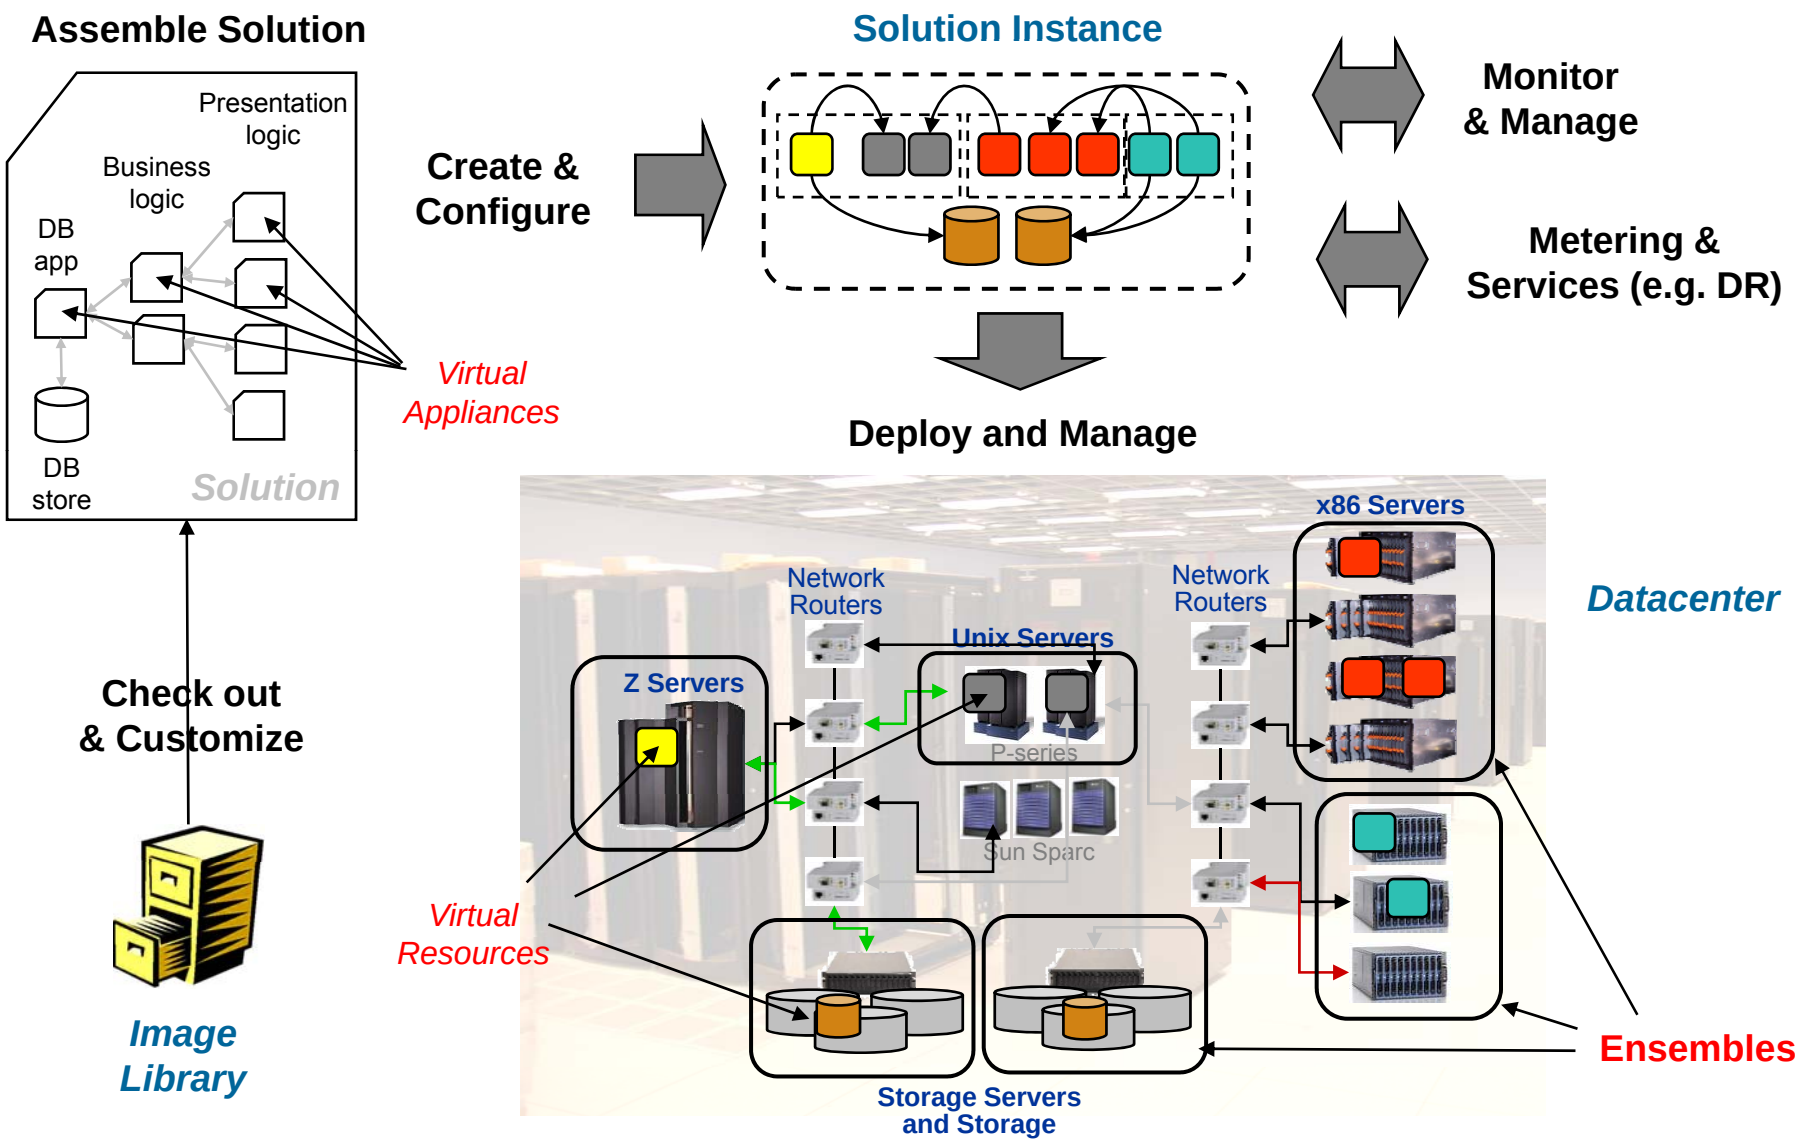

Business Resilience – advances through virtualization

Resilience: **the ability to be ready – to take advantage of good changes and bounce back from bad changes**

Business resilience

Application resilience

Infrastructure resilience

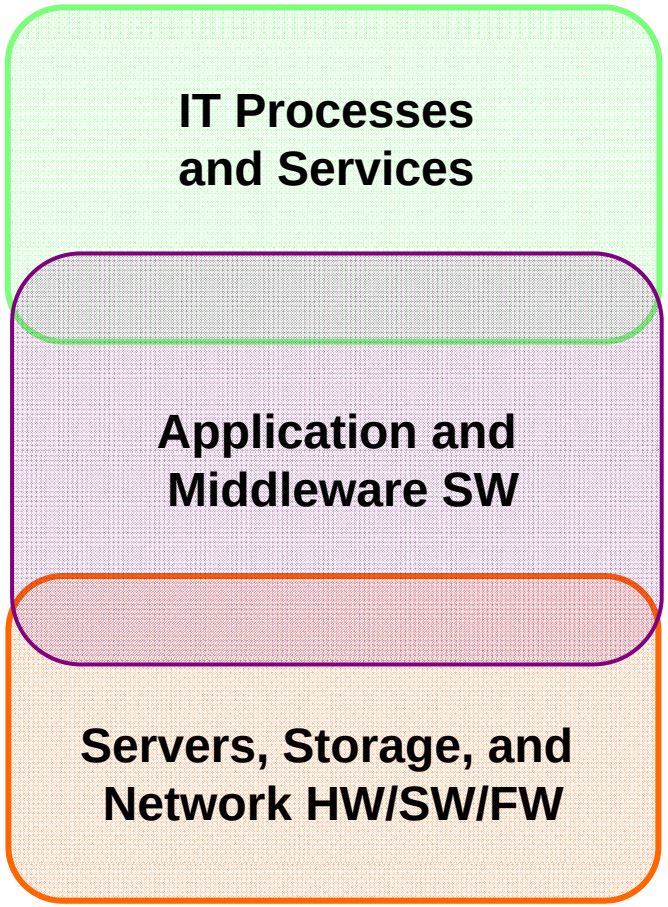

"Virtual Resources" Approach To Server Resilience

Traditional HA software

- Gets added to software stacks
- Can integrate with OS / MW / app
- User provides scripts that specify actions to be taken per event

\+ Can provide high HA functionality

– Costly to deploy and manage

– Brittle to change

"Virtual Resources" approach

- HA is provided outside SW stack
- Automatic HA event handling
- User selects behavior from menu

\+ Easy to deploy and manage

\+ Consistent for diverse SW stacks

– Less HA functionality (e.g., can't recover applications)

Both approaches are needed

Availability Event	LPAR (VM) Action Taken	WPAR Action Taken
HW/FW upgrade or maintenance	LPAR active evacuation	WPAR active evacuation
Predicted hardware failure	LPAR active evacuation	WPAR active evacuation
Unpredicted hardware Failure	LPAR remote reboot	WPAR remote restart
LPAR OS upgrade	LPAR orderly shutdown	WPAR active evac. to local or remote LPAR
Predicted LPAR failure	LPAR orderly shutdown	WPAR active evac. to local or remote LPAR
Unpredicted LPAR failure	LPAR local reboot	WPAR restart on local or remote LPAR

Actions depend on whether LPARs and/or WPARs are used

Use of LPAR Mobility for Server Downtime Avoidance

Event notification:
server will power down for
- **Hardware maintenance**
- **Firmware upgrade**
- **Predicted HW failure (PFA)**
- **Energy conservation**

Server restored
- **Testing run if needed**
- **All LPARs moved back**
- **Server set to "online"**

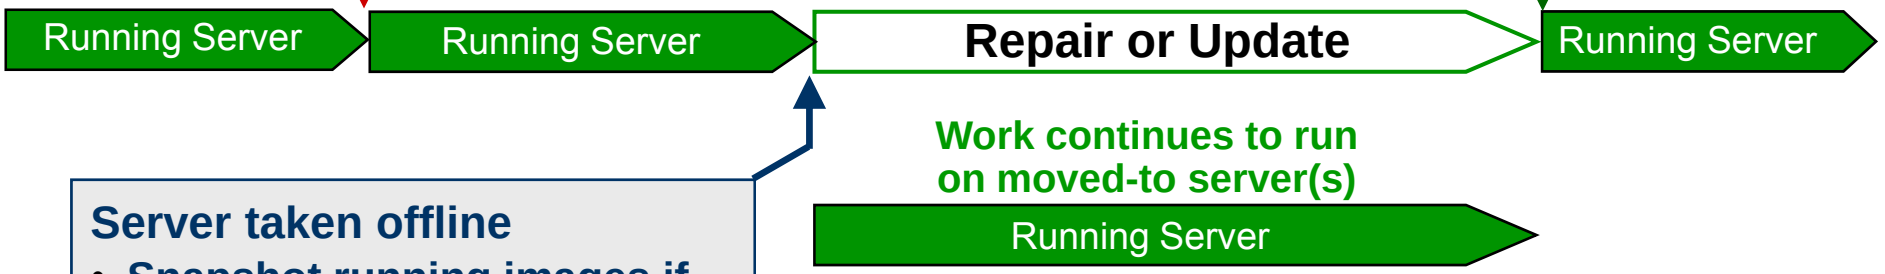

Work continues to run on moved-to server(s)

Server taken offline
- **Snapshot running images if needed for possible backout**
- **Select migration destinations for all LPARs**
- **All LPARs moved while active to other servers**
- **Server set to "standby"**

No downtime
No manual intervention

Tri-Active Disaster Transparency – IBM GTS Vision

"Three is the Magic Number for Five Nines"

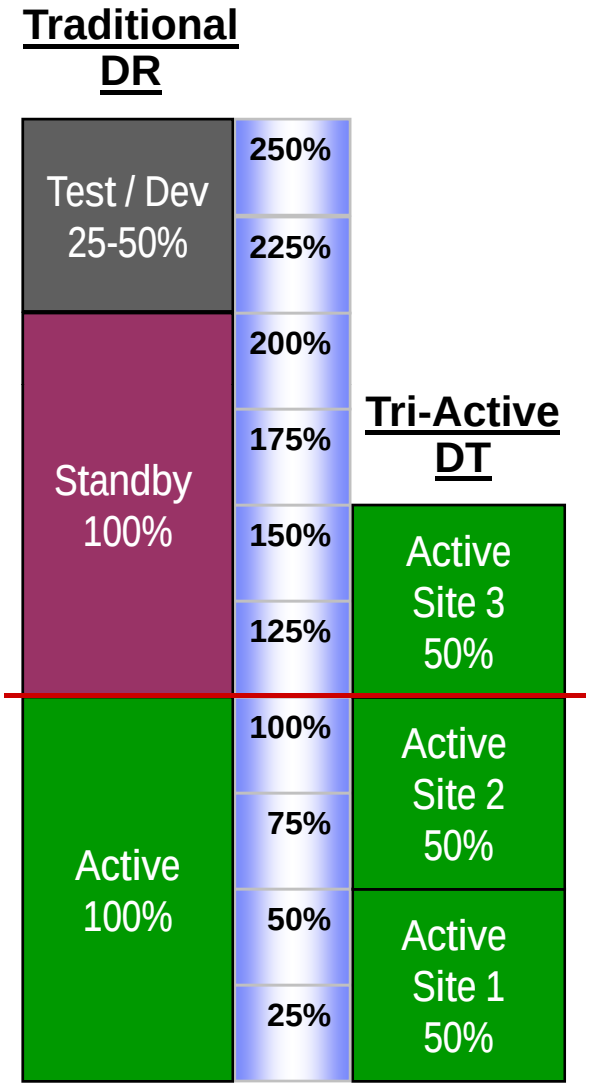

Tri-Active DT Attributes

- Active-Standby pairs are distributed across three sites using virtualization
- Continuous application delivery is decoupled from physical infrastructure
- Disaster Transparency replaces Disaster Recovery
- Production Capacity: 150% of required level
- Failure Impact: 33% (down to 100% of required capacity)
- Test / Dev needs are met using Standby capacity – and preempted quickly when production needs that capacity
- Recovery Time: Seconds
- Maintenance windows are measured in minutes or seconds, not hours
- Service Restoration is a concept nobody remembers

Virtualization will enable dynamic sharing and repurposing across data centers

Security – securing the virtual data center

Security and Virtualization

- **The growing use of virtualization introduces new IT security risks and challenges**
  - Sharing of physical resources creates new exposures for unauthorized interactions
  - Virtualization software itself may be compromised
- **A solid foundation of virtualization technology is critical to achieving IT security**
  - Virtualization technologies need to be hardware-integrated and protected
  - Security services need to be provided as virtual appliances
- **Security management of the virtualization technologies is required to reduce the risk of security exposures and to enable security policy enforcement**
  - Isolation Management will enforce restrictions on administration and resource sharing
  - Integrity Management will maintain a resource inventory and act as an early warning system for anomalies (detect and report)

Security Role of Hypervisors

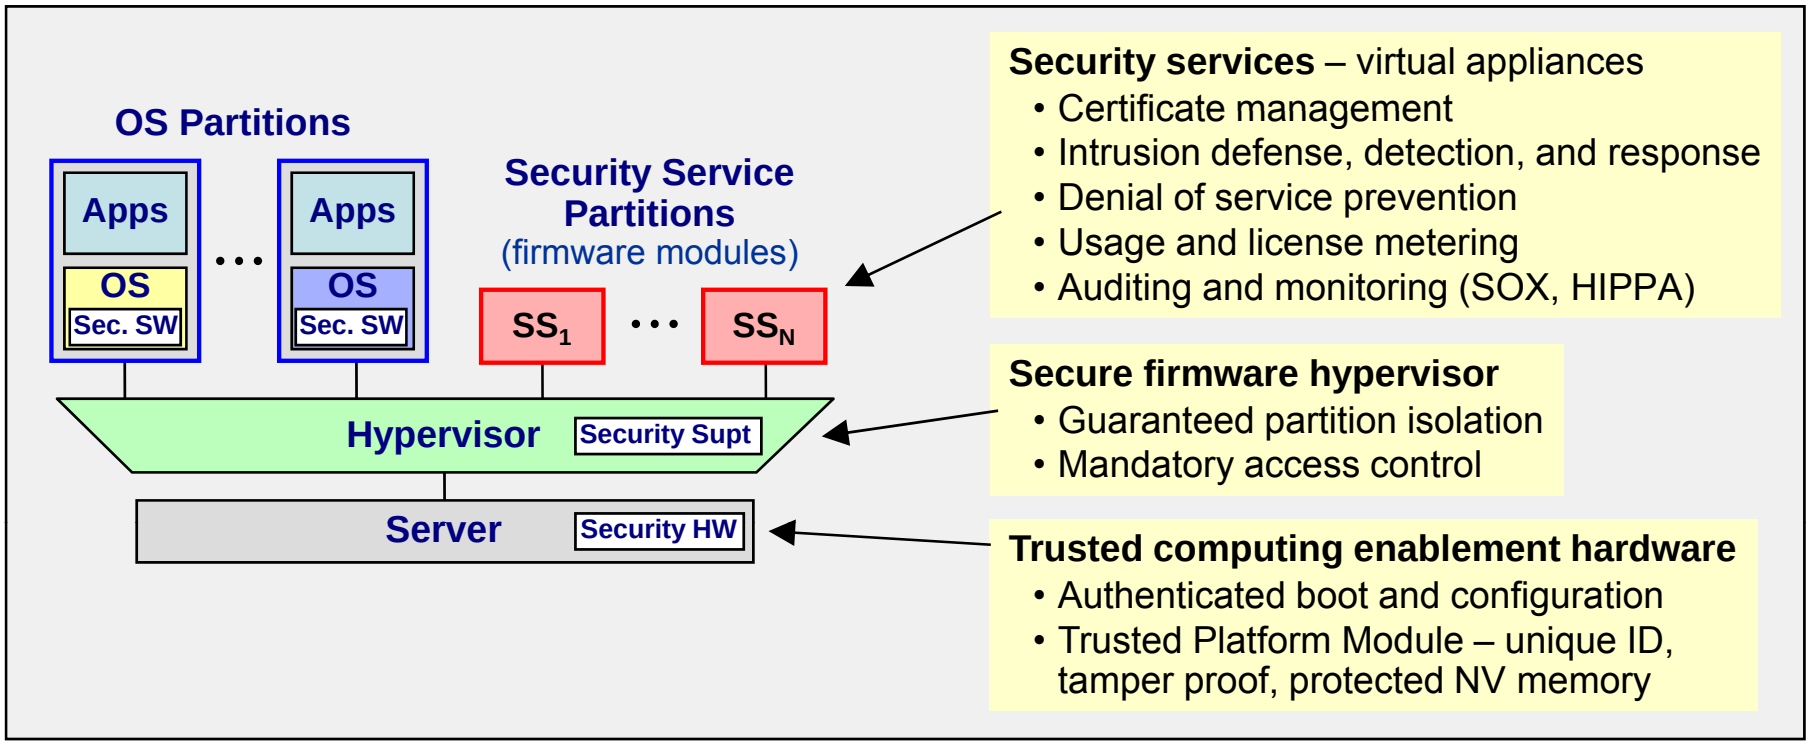

- **General purpose operating systems are a weak foundation for secure computing**
  - Large and complex with many latent bugs
  - Constantly changing – security certifications are quickly made irrelevant
  - TCP/IP stacks make them vulnerable to attack by viruses, worms, and hackers
- **Hypervisors will be used to establish a solid foundation for secure computing**
  - Small enough to be fully inspected and certified
  - Limited functionality – may rarely need to be changed – suitable as a BIOS extension.

Security Is Limited By The Weakest Link

Trusted Virtual Data Center (TVDc)

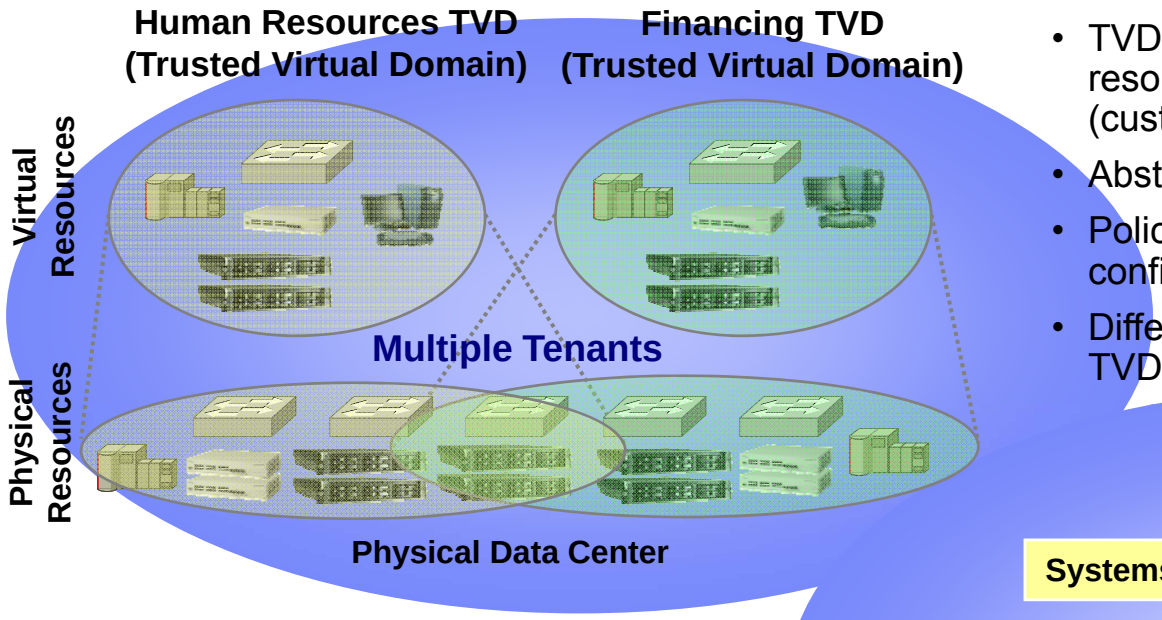

Centralized IT Security Management

- TVD: Grouping of VMs, network and storage resources that support common objective (customers, workloads, etc.)
- Abstracts the physical infrastructure
- Policy-driven: consistent security configuration and management)
- Differentiated views: data center admin. vs. TVD admin

Distributed Enforcement

- Very strong, coarse-grain security guarantees – cannot be bypassed by VMs
- Single data center security policy across different platforms and hypervisors
- Containment (malware, insider attacks) & Trust

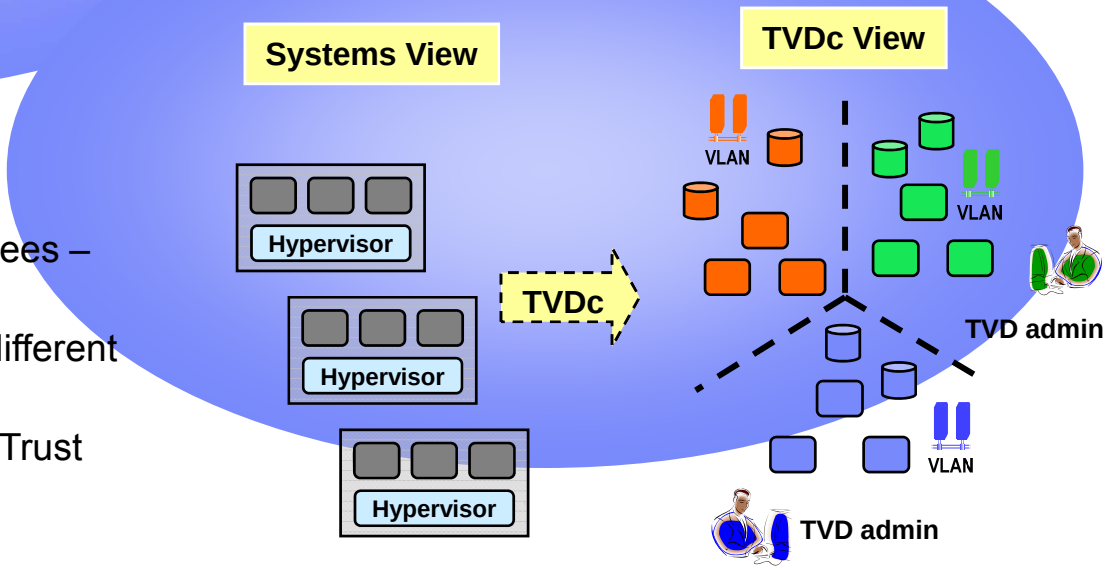

Cloud Computing – the next phase of Internet-based sharing

What Is Cloud Computing? A Simple View

Cloud Computing is an emerging style of computing in which applications, data, platforms, and resources are provided as services to users over the Web.

- ⇨ The services provided may be available globally, always on, low in cost (economies of scale), "on demand", massively scalable, "pay as you grow", …
- ⇨ Consumers of the services need only care about what the service does for them, not how it is implemented

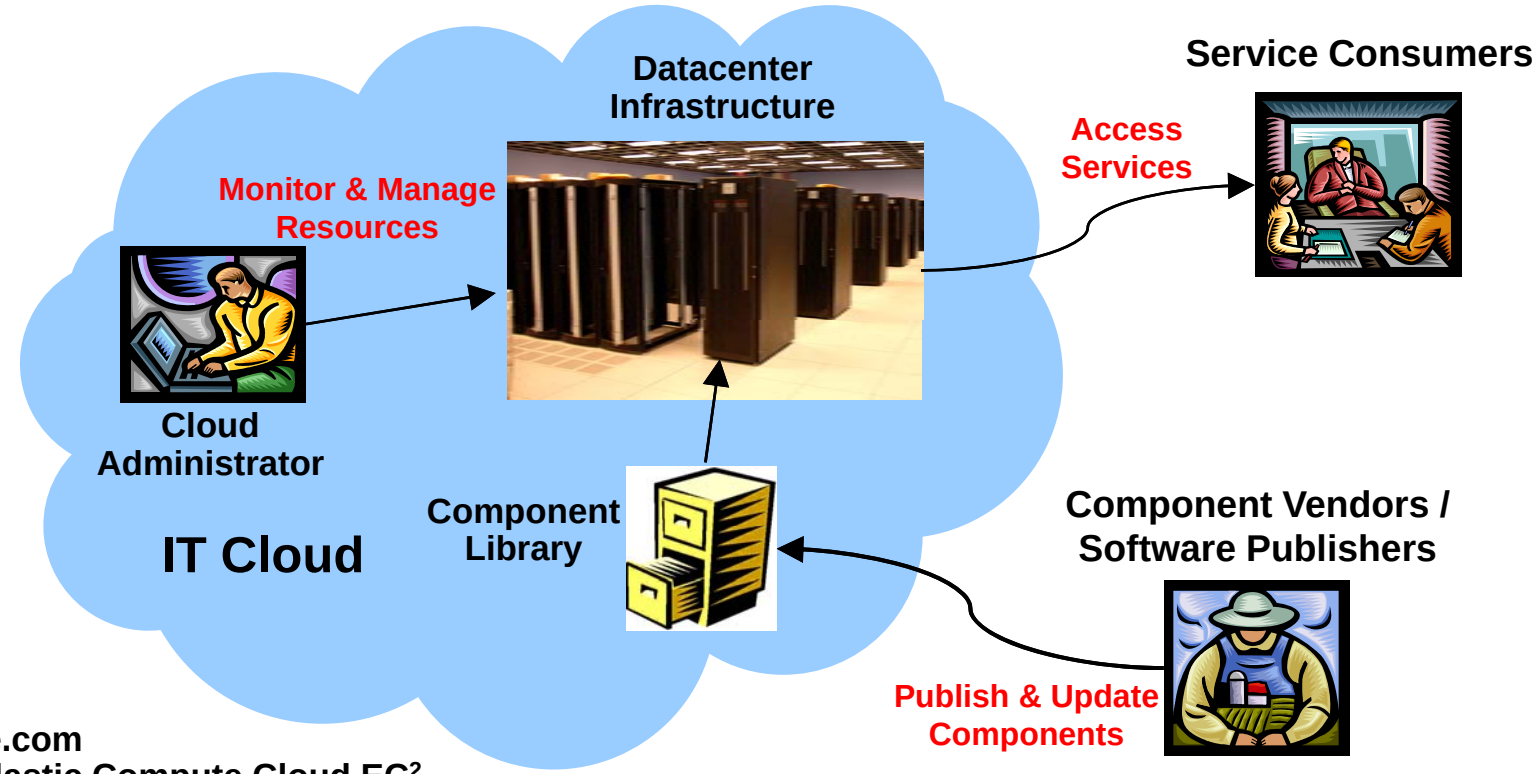

Examples:
- **Salesforce.com**
- **Amazon Elastic Compute Cloud EC²**
- **IBM Research Compute Cloud RC²**

Evolution of Sharing on the Internet

Services Sharing — **Cloud Computing**

- Everything as a service over the Web: SaaS, utility computing, IT services, …
- Ubiquitous, always available, scalable, …

Resource Sharing — **Grid Computing**

- Standards and software for sharing of remote resources and collaboration
- Mainly used for highly scalable HPC jobs

Information Sharing — **The World Wide Web**

- HTML page format, HTTP protocol, and Mosaic browser for document exchange
- Initially in universities; worldwide adoption

Network Sharing — **Inter-Networking and the Internet**

- Inter-Networking of regional networks with TCP/IP
- Began to replace regional alternatives
- Worldwide adoption

Networking — **Networks**

- Multiple regional networks linking computers
- Initially at universities and national labs

Diversity of Cloud Services

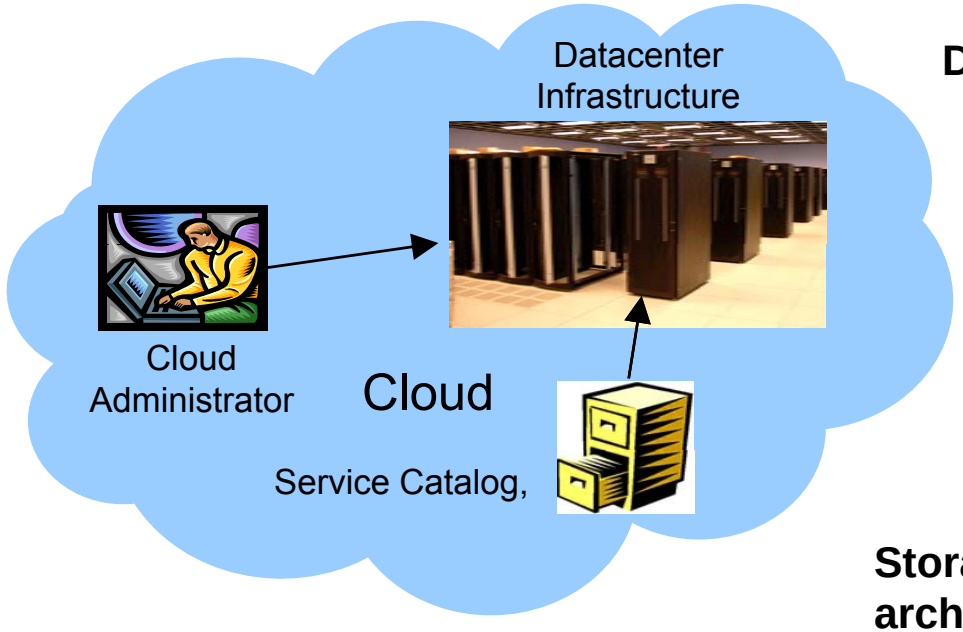

Web Application Services

Collaboration Services

Database Services

Compute Services

Storage Services

Job Scheduling Services

Storage backup, archive… service

Virtual Client Services

Content Classification

Examples of Cloud Services

1. **Software as a Service (SaaS)**
   - Applications are delivered through the browser to multitudes of customers.
   - Customers avoid upfront investment in servers and software licensing.
   - Providers spread application hosting costs over large customer base.
2. **Managed IT Services as a Service**
   - Services such as a virus scanning service for e-mail or an application monitoring service are provided as managed services to IT groups rather than to end-users
3. **Web Services as a Service**
   - Web service providers offer APIs that enable developers to exploit functionality over the Internet, rather than delivering full-blown applications.
4. **Development Environments as a Service**
   - Vendors provide development environments as a service.
   - Clients build applications that run on the provider's infrastructure and are delivered to end users via the Web from the provider's servers.
5. **IT Infrastructure as a Service (Utility Computing)**
   - Data center infrastructure elements (storage, virtual servers, and even virtualized data center partitions) are provided on demand over the Web.
   - Initial adoption is mainly for supplemental, non-mission-critical needs, but broader use may follow.

Types of Clouds and Cloud Services

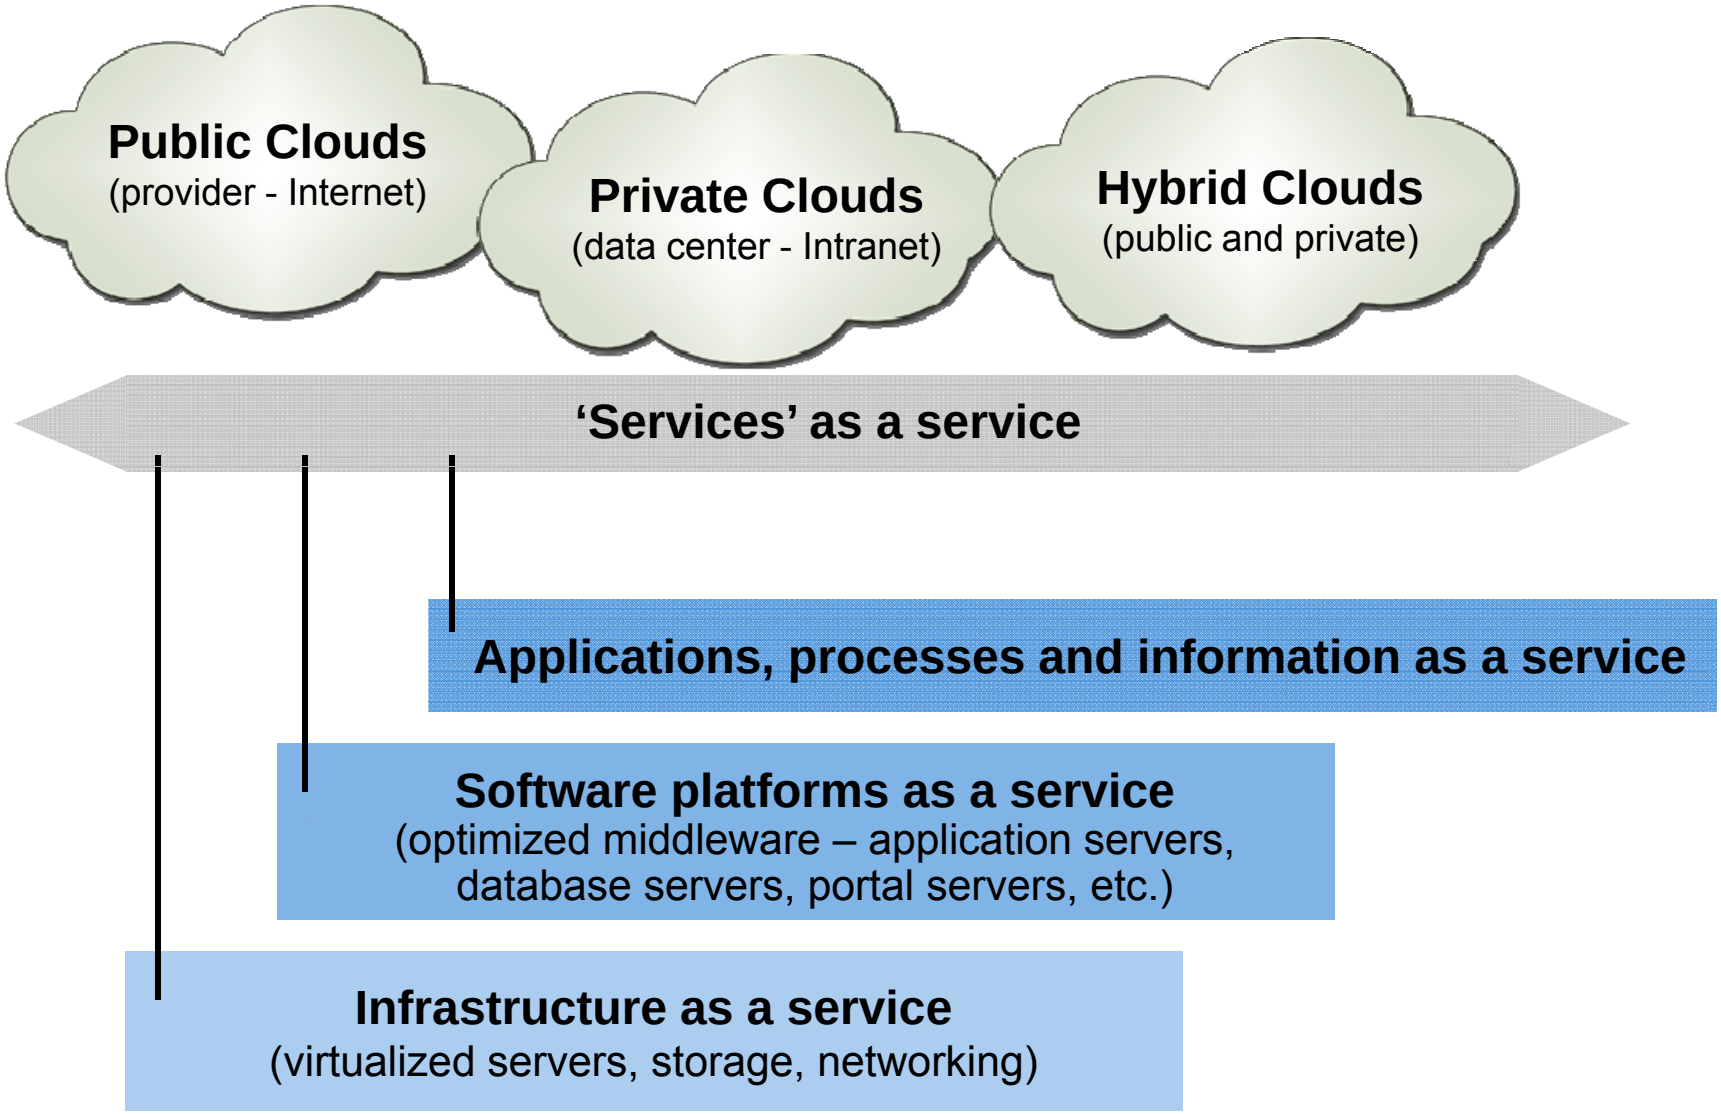

Notable Cloud Computing Players

- **Amazon**
  - Web-based retail sales business
  - Elastic Compute Cloud (EC^2): virtual servers as a service (Xen/SLES based)
  - Simple Storage Service (S3): online storage as a service
- **Google**
  - Leading "search" application (using proprietary MapReduce; Linux based)
  - Innovative advertising-based business model
  - New suite of free word-processing and spreadsheet software via browser
- **IBM**
  - Worldwide cloud deployments (e.g., using Hadoop open source version of MapReduce)
  - Research Compute Cloud (RC^2): Self-service, Web-based cloud for IBM Research
  - NC State University Virtual Compute Lab: Xen/Linux cloud for teaching and research
- **Microsoft Azure**
  - Windows Live software suite (includes electronic mail, photo-sharing, …)
  - Direction: suite of services and applications for personal and community use across the PC, the Web, and the phone; including data management
- **VMware vCloud**
  - New initiative to deliver cloud computing by federating compute capacity on demand between virtual datacenters and cloud service providers

Summary

Summary

- **Consolidations provide compelling cost savings today by improving hardware utilization, energy efficiency, and ease of software stack deployment**
- **Introducing virtualization where it has not be used before increases complexity and cost, offsetting virtualization benefits**
- **Integration of standard hardware, virtualization technologies, management software, platform software … will be used to hide many virtualization complexities**
- **Emerging service management software will use virtualization to achieve a breakthrough in IT simplification and effectiveness**
- **Cloud computing will be employed widely, leveraging virtualization to implement a diversity of valuable public and private services**

Closing Thoughts

- **Virtualization will be vital, but it is only one of the required technologies**
- **IBM will provide a full range of the advanced systems, software, and services required for dynamic infrastructures and cloud computing**

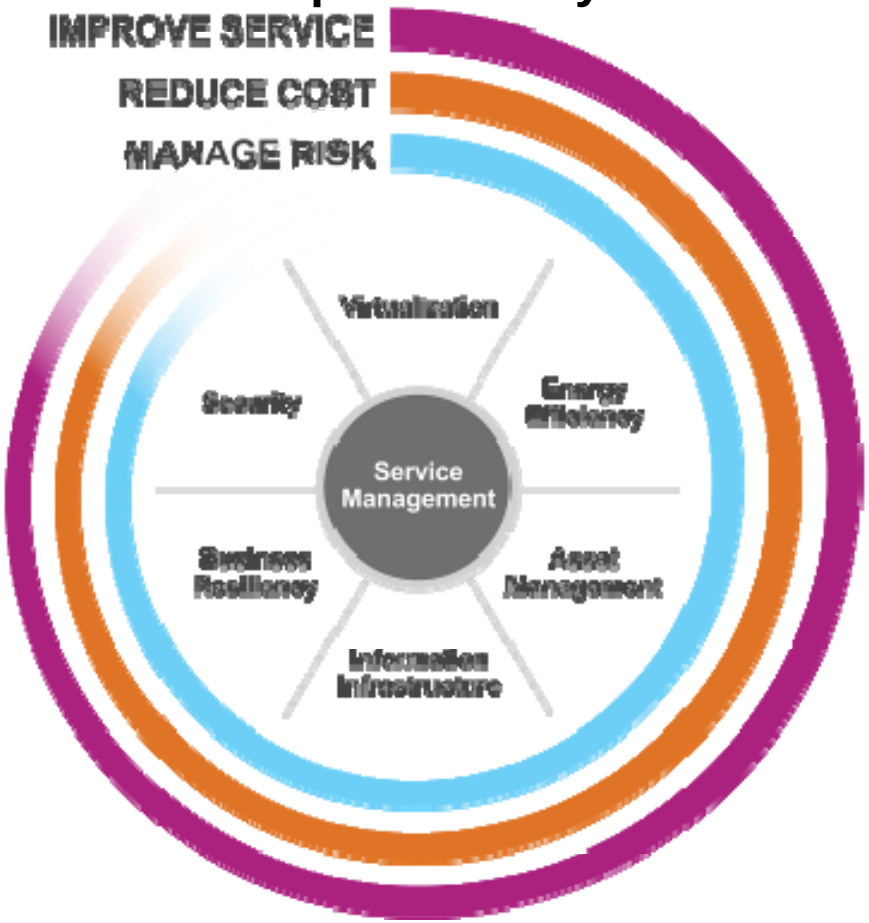

- **Critical customer requirements:**
  - ✓ **Improve Service** (Quality, Speed, Agility)
  - ✓ **Reduce Cost**
  - ✓ **Manage Risk**
- **Required enabling technologies:**
  - ✓ **Service Management**
  - ✓ **Asset Management**
  - ✓ **Virtualization**
  - ✓ **Energy Efficiency**
  - ✓ **Business Resiliency**
  - ✓ **Security**
  - ✓ **Information Infrastructure**

The end.

Thank you!

Any questions or thoughts?

Trademarks and Disclaimers