

Thoughts on Utility, Grid, on demand, Cloud computing and Appliances
NCOIC Conference, St Petersburg, FL

Mark Cathcart, IBM Distinguished Engineer,
m_cathcart@us.ibm.com

<http://ibmcorner.com>

Dec. 6th, 2007 - (c) IBM Corporation 2007

INNOVATION THAT MATTERS

IBM

Agenda



- › **Looking back**
- › **Evolution of Grid and on demand**
- › **Cloud Computing**
- › **Where Next?**
- › **The Appliance conundrum**

IBM Evolution towards Internet scale data centers



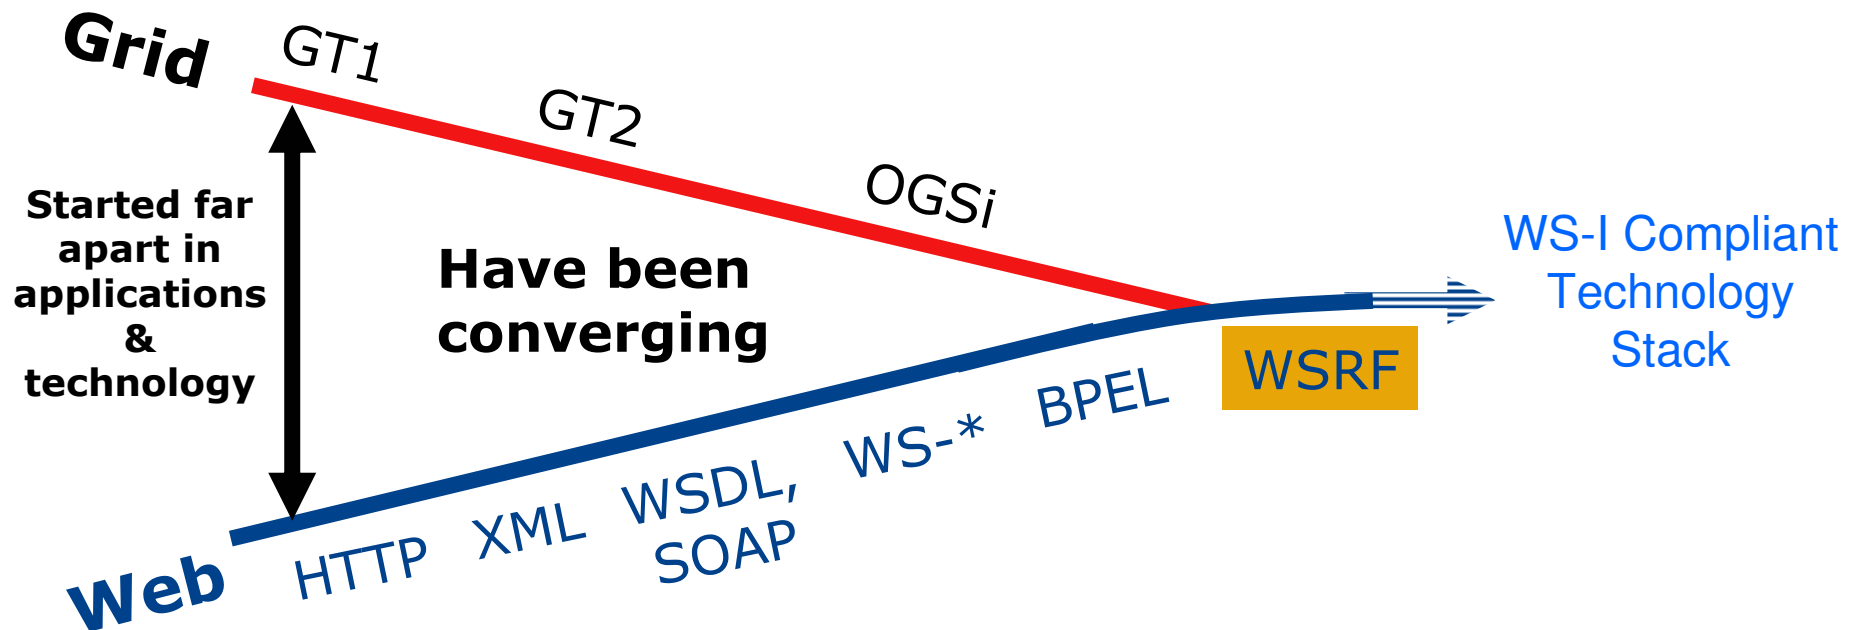
- › **Utility computing initiative started in 1999**

- › **Investment in Grid computing in 2001 onwards**
 - Through Globus, Univa – Globus Toolkit
 - Contributed to Web Services
 - Open Grid Services Interface(OGSI) and Architecture(OGSA)

- › **Web Services evolution**
 - From web technology to secure Enterprise interoperability

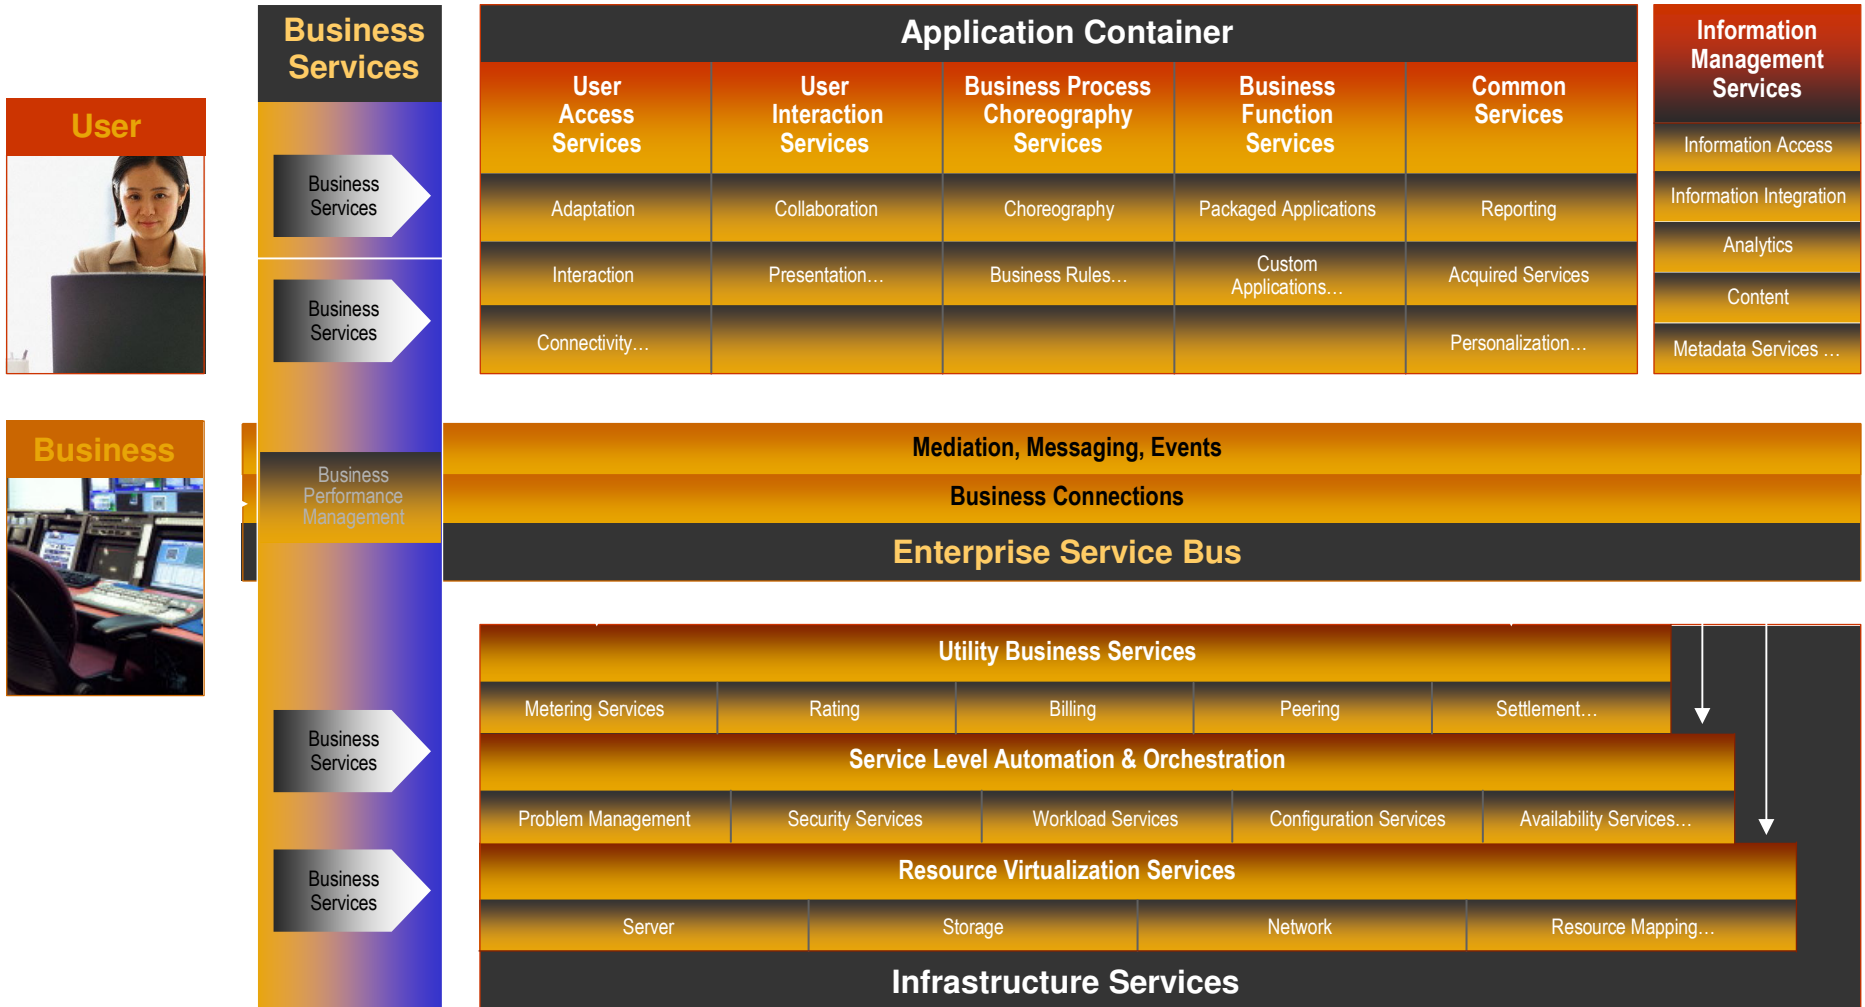
- › **Open Standards for Managenments and Interoperability**
 - Web Services Distributed Management(WSDM), CIM, WSRF, et al

Grid and Web Services Standards



Convergence of Core Technology Standards allows Common base for Business and Technology Services

The Operating Environment Architecture

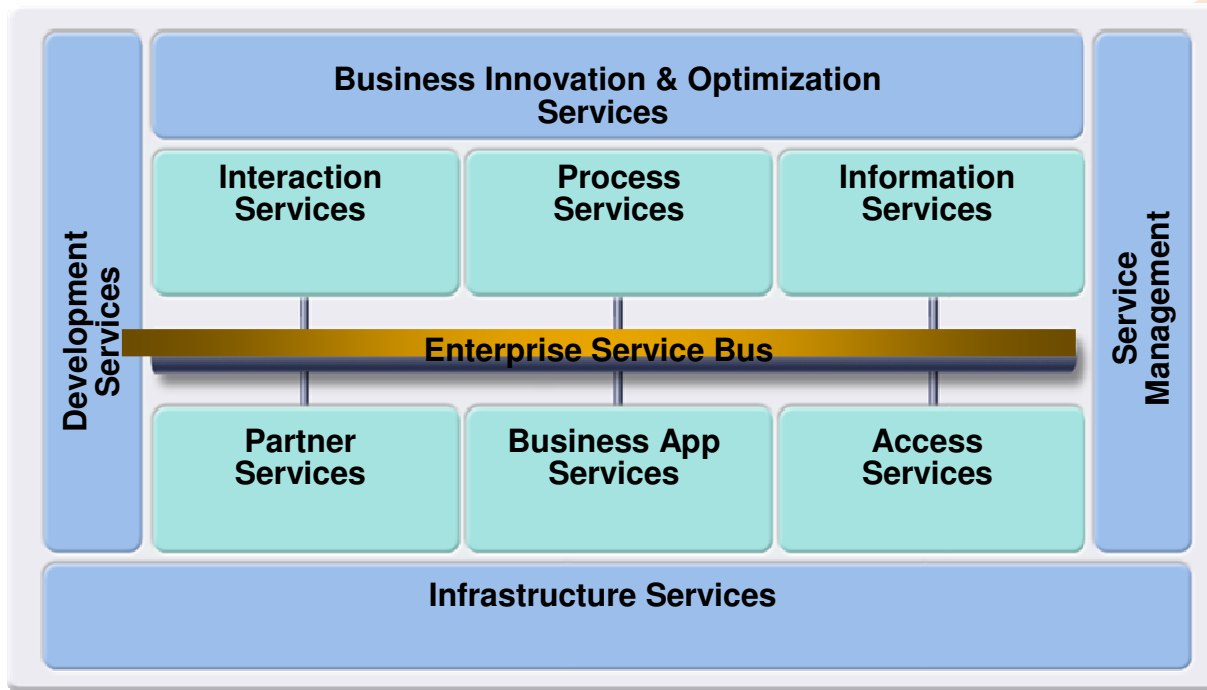


SOA Reference Architecture

The Enabling Foundation for an SOA Product Strategy



IBM SOA Reference Architecture



- 1 Provide a common, open-standards-based run-time...
- 2 ...on top of which we build added-value targeted at specific domains...
- 3 ...providing a platform for clients and ISVs / SIs to build solutions...
- 4 ...around which we apply a common set of quality of service extensions

Built on a common framework for: Tooling, Runtime, Messaging, and Infrastructure

IBM Unveiled Plans “Blue Cloud” on 11/15



- › Family of ready-to-use cloud computing offerings
 - Based on open standards and open source software together with IBM software, systems technology and services
 - First offering to support Power and x86 processors
 - Plans to support System z and highly dense rack clusters
- › Reduces IT management complexity and increases business responsiveness
- › Supports both existing and emerging, data-intensive workloads
- › Vietnamese Ministry of Science and Technology plans to leverage Almaden Research Center’s cloud environment to run country innovation portal

Initial Blue Cloud offering targeted for Spring, 2008 availability

Forces Driving Cloud Computing





Explosion of data intensive applications on Internet



Fast growth of mobile commerce



Increased network capacity and availability



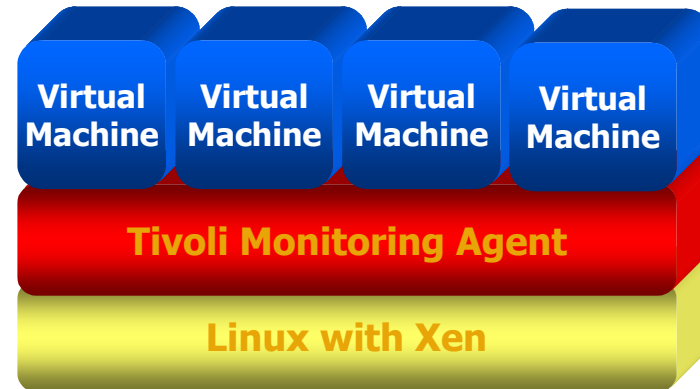
Advances in computer architecture and price/performance

Requires massively scalable cloud infrastructures to serve billions of heterogeneous browser-based clients

Blue Cloud Initial Offering




Delivers a massively scalable and flexible compute platform for hosting both existing and emerging data-intensive workloads.



Virtualized Infrastructure Based on Linux & Xen

- Based on open standards and open source software
- Includes IBM software, systems technology and services
- Supports both Power and x86 processors
- Web 2.0 resource reservation system



Provisioning Management Stack



IBM and Massive Scale Computing



- › IBM is uniquely qualified to lead in cloud computing
- › Blue Cloud builds off of IBM's decades of experience developing and leading massive-scale computing:
 - Parallel Sysplex
 - IBM's Deep Blue SP Cluster
 - Blue Gene
 - Grid Computing

Business Benefits of Blue Cloud



Cost efficient model for creating and acquiring information services

Reduces IT management complexity

Increases business responsiveness with real-time capacity reallocation as demand for compute power grows

Powers both existing and emerging data-intensive workloads

Vietnam Ministry Leverages Cloud to Run Innovation Program



Students

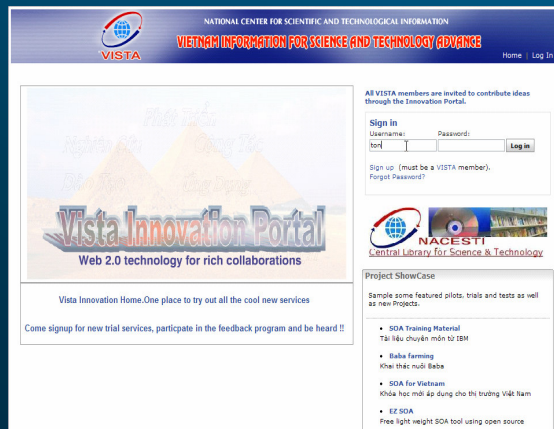


Teachers




Researchers

VISTA Innovation Portal (VIP)



VIP pilot hosted on IBM's Blue Cloud computing infrastructure at Almaden



Blogs

Wikis

Forums

Profiles

Social Tagging

Information Discovery

IBM Innovation Factory

VIP, powered by IBM Innovation Factory, provides a platform to foster collaborative innovation among major universities and research institutes.



First Principles For Simplification



- › **Decompose large/hard problems into multiple smaller/easier ones**
 - Form complex systems from modular subsystems (e.g., using services model)
 - Employ layered/hierarchical system structures
- › **Avoid interlocking dependencies**
 - Decouple components (both hardware and software) using virtualization, standard interfaces, and service abstractions
 - Provide compatibility per component type/family, and across versions over time
- › **Use standard interfaces, formats, and functions**
 - Use standard components and configurations; reduce the number of unique building blocks
- › **Use building blocks that hide their internal implementation**
- › **Leverage homogeneity**
 - Aggregate multiple building blocks of like type into pools that are managed as a single entity
- › **Limit change to where it really matters**
 - Exploit mature products and processes to minimize problems/incidents and leverage existing skills
- › **Unify management infrastructure and centralize points of control.**
 - Avoid use of multiple, interdependent management infrastructures, functions, interfaces
- › **Reduce the need for user/operator skills and effort**
 - Simplify and automate tasks that are complex or costly
 - Provide intelligent default behaviors

IBM Virtualization Strategy



Service Management: Tivoli

1. Tivoli roadmap to fully support and leverage virtualization on IBM and industry platforms
2. Common development with Systems Director
3. Advantage IBM platforms through additional exploitation and solution orientation

Process Management

Change Mgmt, Configuration, Problem Mgmt, Incident Mgmt, Availability Mgmt

Business System Management

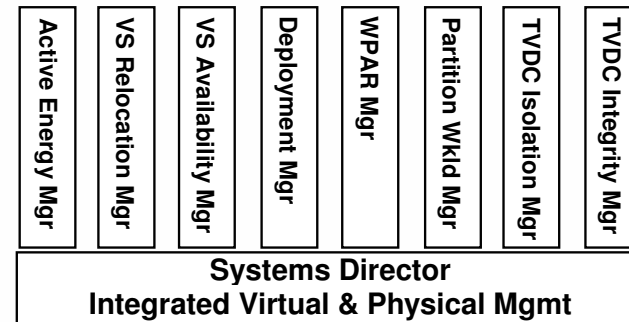
Business Continuity QOS (& Capacity Management) Resource Pool Mgmt.

System Lifecycle Management

Create, Deploy, Discover, Monitor, Update, Audit /Compliance/Secure, License


Platform Management: Systems Director

1. All IBM platforms and multiple x86 hypervisors
2. Advantage System z & p - common user interfaces with high volume servers for skills leverage
3. Support multiple x86 hypervisors – including VMware – to alleviate VMware domination
4. Compelling roadmap, infuse selective acquisitions



Virtualization Technologies


1. Sustain leadership in System z & p and storage: customer value from superior scaling and RAS
2. Foster open industry alternatives to VMware in x86 space to alleviate their market domination
3. Drive adoption of IBM storage virtualization



Classic Complex Data Center




Business Processes as Services



Topologies of federated services are mapped onto large numbers of diverse virtual and physical resources from competing vendors

Sea of Heterogeneous Servers, Storage, Networks and Their Virtualization




- › Data center complexity has reached crisis levels and is continuing to increase
- › Businesses spend a large fraction of their IT budgets on data center resource management rather than on valuable applications and business processes

Data Center “implementation” Approach To IT Simplification




Business Processes as Services



Topologies of federated services are mapped onto smaller numbers of systems. implementations replace ad hoc groups of systems.

Sea of Heterogeneous Servers, Storage, Networks and Their Virtualization




› An **implementation** is a new kind of parallel system designed to simplify data centers

- It consists of a pool of compatible networked systems with integrated virtualization and management software that manages and leverages virtualization
- It can scale from a few to many thousands of servers, while having management complexity / cost essentially independent of its size and like that of a single system
- Parallel Sysplex, Blue Gene, and Google are examples of implementations in many ways
- Virtual resource mobility is an important capability within most implementations

Data Center “implementation” Approach To IT Simplification

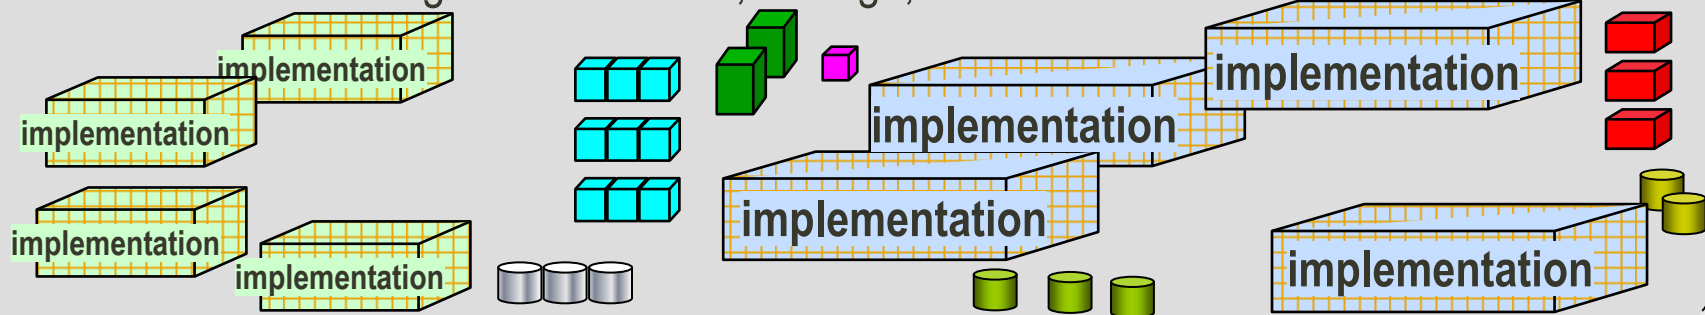


Business Processes as Services



Topologies of federated services are mapped onto smaller numbers of systems. implementations replace ad hoc groups of systems.

Sea of Heterogeneous Servers, Storage, Networks and Their Virtualization



› An implementation is a new kind of parallel system designed to simplify data centers

- It consists of a pool of compatible networked systems with integrated virtualization and management software that manages and leverages virtualization
- It can scale from a few to many thousands of servers, while having management complexity / cost essentially independent of its size and like that of a single system
- Parallel Sysplex, Blue Gene, and Google are examples of implementations in many ways
- Virtual resource mobility is an important capability within most implementations

Blue Cloud implementation Illustration



IT Cloud:

- Heterogeneous with many traditional networked systems (not shown)
- Immense scale (may even span datacenters)
- Network service (data center fabric)
- Security service (federated identity)
- Service Registry

“implementation”

Autonomic Resource Pool: (1 of N per cloud)

- Homogeneous systems with virtual resource mobility functions (z/VM, PHYP, Xen)
- Locally networked servers and storage / SVC
- Integrated software stacks
- Libraries of virtual resources (e.g., images)
- Scales from a drawer to many racks
- Integrated service management:
 - Workload optimization
 - Availability
 - Recovery
 - Change/release mgmt
 - Security

Blue Cloud implementation Examples



IT Cloud:

- Heterogeneous with many traditional networked systems (not shown)
- Immense scale (may even span datacenters)
- Network service (data center fabric)
- Security service (federated identity)
- Service Registry

Consolidation service

Virtual Client service

SAP Application service

z/VM, LPAR service

PHYP service

Xen service

Hybrid service

DB service

File system service

Storage backup, archive, service

Implementation Value Proposition



- › **implementations are highly scalable, virtualized systems that provide a single system management image – they look like single systems with respect to management.**
 - implementations dramatically reduce the effort customers must expend in managing their physical and virtual hardware resources.
- › **implementations apply the first principles of simplification:**
 - Insulate IT software stacks from changing hardware
 - Make hardware/firmware maintenance non-disruptive to virtual resources and their software stacks/applications
 - Avoid the need to upgrade software stacks in order to move them to newer hardware
 - Aggregate like resources into pools that can be managed as a single entity independent of their aggregate scale
 - Integrate IT components so they work together “out of the box”
 - Support software appliances for a variety of benefits
 - Separate application management from IT resource management, allowing customers to focus more on their applications and less on the IT plumbing
 -



Services Map To Images Which Deploy To implementations



Processes

#1: Tools & Standards for image creation


image



image



image



#2: GRID-SOA Interfaces

GRID-SOA Abstraction ('place images in implementations')

#3: Deploy Images from libraries

IBM implementations

#4: Optimize implementation for performance, availability, and efficiency



System z



System p



System x

Moving toward Internet Scale Data Center



Simplified

- Consolidate from many to fewer data centers
- Server/Storage consolidation
- Server / Storage Virtualization
- Improve business resilience
- Reduce Systems Management Complexity
- Green
- Converged Network

Shared Infrastructure

- Highly Virtualized pools of resources (homogenous clouds)
- Automated Service management
- Green design
- Optimized Appliances and Hybrid Systems
- Advanced provisioning and workload management
- High performance network

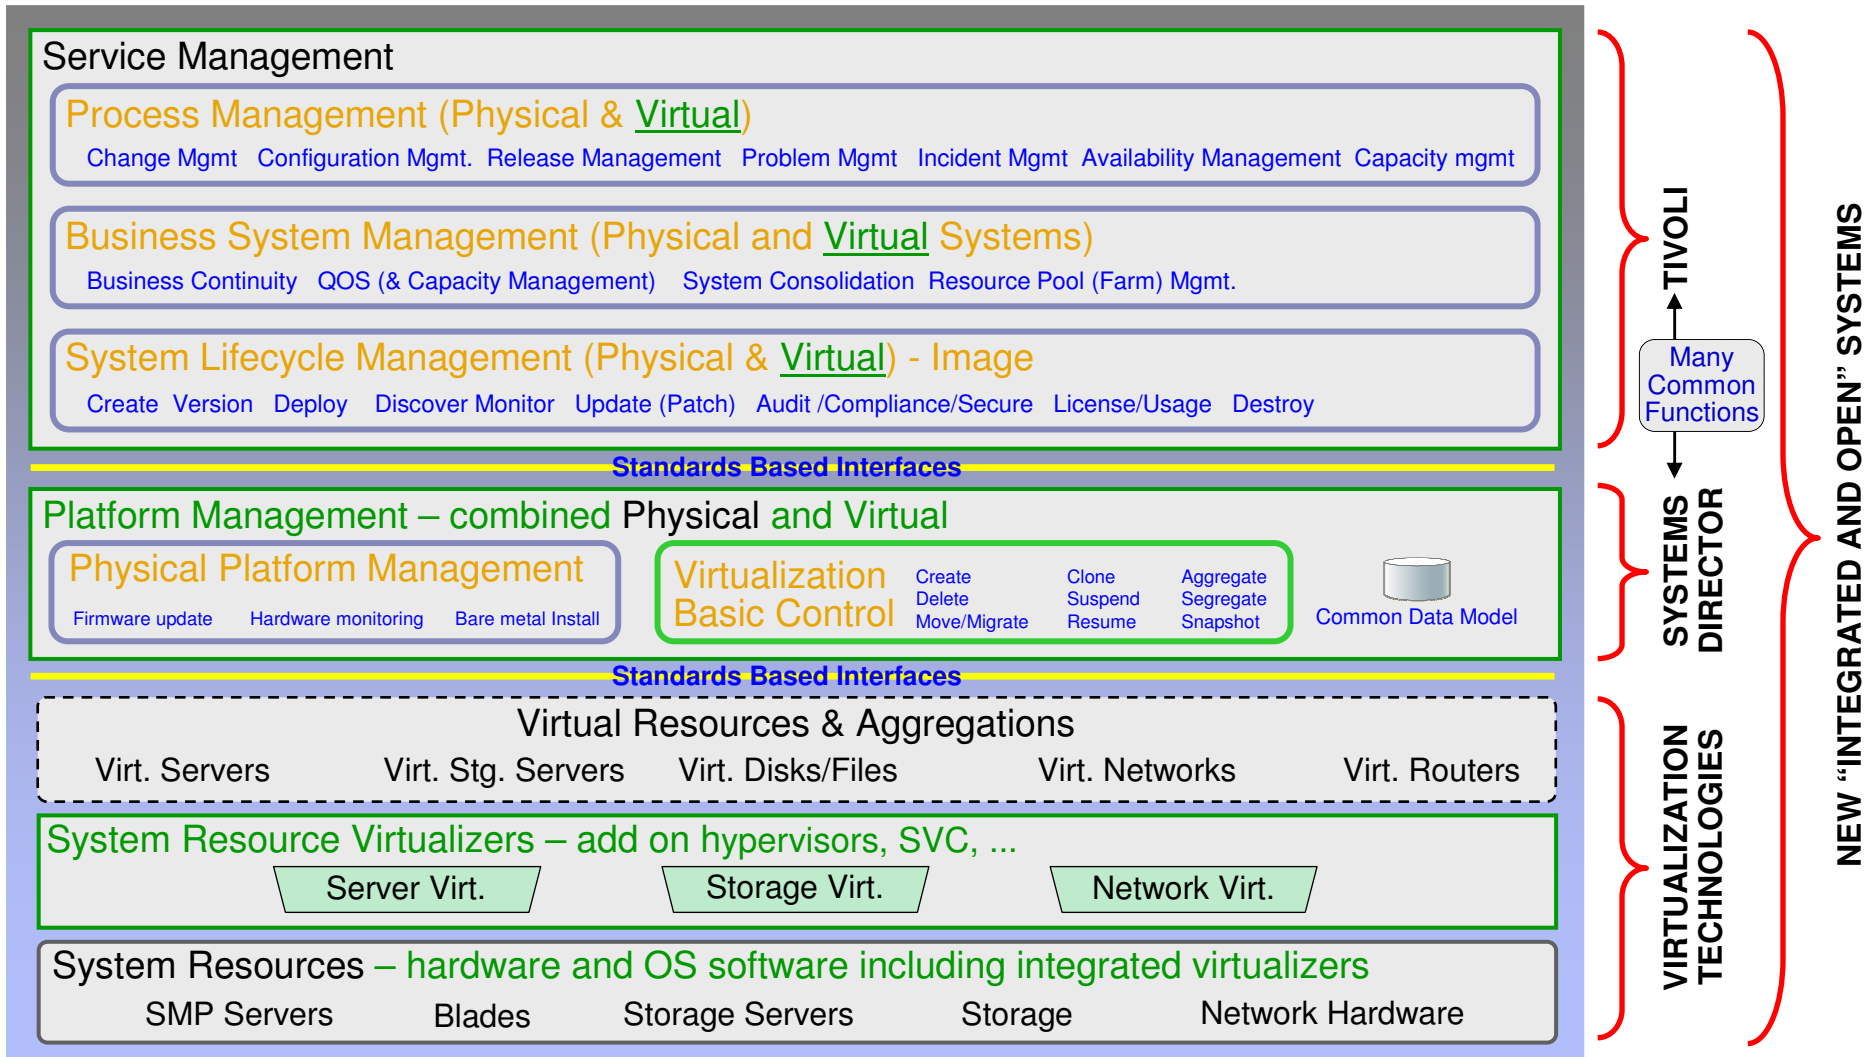
Service Orientation

- Internet scale
- Embedded super computers
- Cloud computing
- Web + Enterprise platform
- Converged network

The Stack: Four Key Infrastructure areas



IBM Plays



Virtual Software Resource Pools?



Composite Configurations of virtual implementation resources

- Networked services made of virtual resources (atoms and molecules) provisioned from virtual resource libraries
- SOA-based

Virtualized Software Resource Pool

- Middleware and Application Pools, hosted on virtualized hardware resource pool
- Mobility of software resources within pool and with like pools
- Initial middleware pools: WAS, db, transaction, file system, LAMP, windows consolidation
- Initial application pools: collaboration, desktop PS, SAP, etc.

Virtualized Hardware Resource Pool

- Pool of compatible servers with integrated hypervisors
- Mobility of all virtual resources within pool and with like pools
- Three main types: System z with z/VM, System p with PHYP, System x with Xen
- Scales from a drawer to many racks

implementation application mgmt. software

- Service Lifecycle management
- Service Registry
- Service Management (monitoring, optimization, security, continuity, process automation)

implementation resource mgmt. software

- Service and application mgmt of Pools
- Automated software lifecycle mgmt
- Automated workload optimization, HA, backup and recovery
- Virtual software image libraries

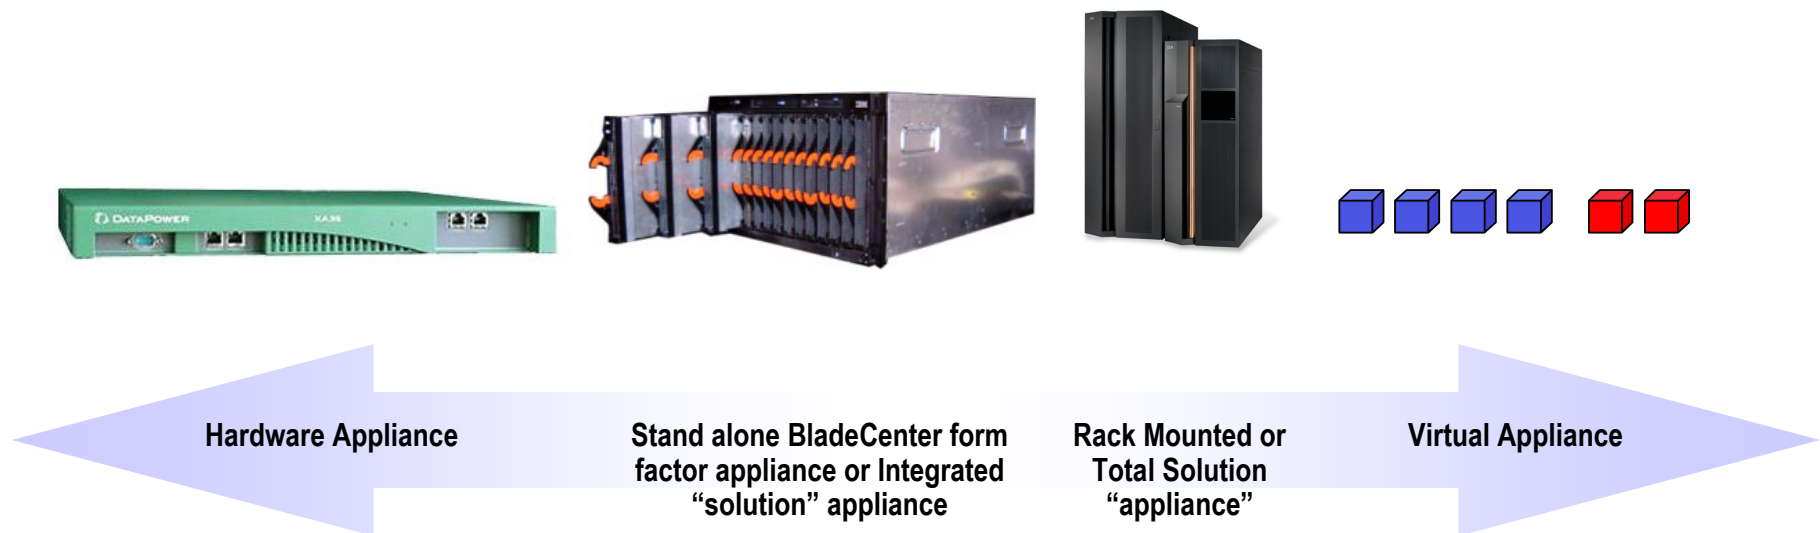
implementation resource mgmt. software

- Service and platform mgmt of physical and virtual resources
- Same APIs/GUIs as for SMP server
- Automated QoS/HA/Power optimization
- Virtual resource libraries

The Appliance spectrum



› Appliances are a spectrum, not just stand-alone device



- › Special purpose hardware, optional in all cases
- › Power processor advantage – low power options, p5+ to p7, PRISM
- › Solution “appliances”
 - Total solution package incorporating H/W, S/W and possibly services

The Power appliance advantage



› Built on the Power platform



How can you “build” a solution today?



“A la carte”

“Integrated”




VMWorld 2007 S288511 Virtual Appliances and the New Datacenter - Massimo Re Ferre', Senior IT/Architect, IBM

How can you “build” a solution today? The true story...



“A la carte”

- Very inefficient
- Lots of hw / os / sw dependencies
- Server proliferation
- Supportability matrixes is more than a nightmare....



“Integrated”

- Solves some issues
 - but doesn't scale.....
- Still too many servers
- Still running at 5 / 20 %
- 100 solutions = 100 different servers/appliances
- HA issues ?
- ISV just can't handle this model
... they deal with SW ... not HW.
- etc. Etc.




VMWorld 2007 S288511 Virtual Appliances and the New Datacenter - Massimo Re Ferre', Senior IT/Architect, IBM

In medio stat virtus

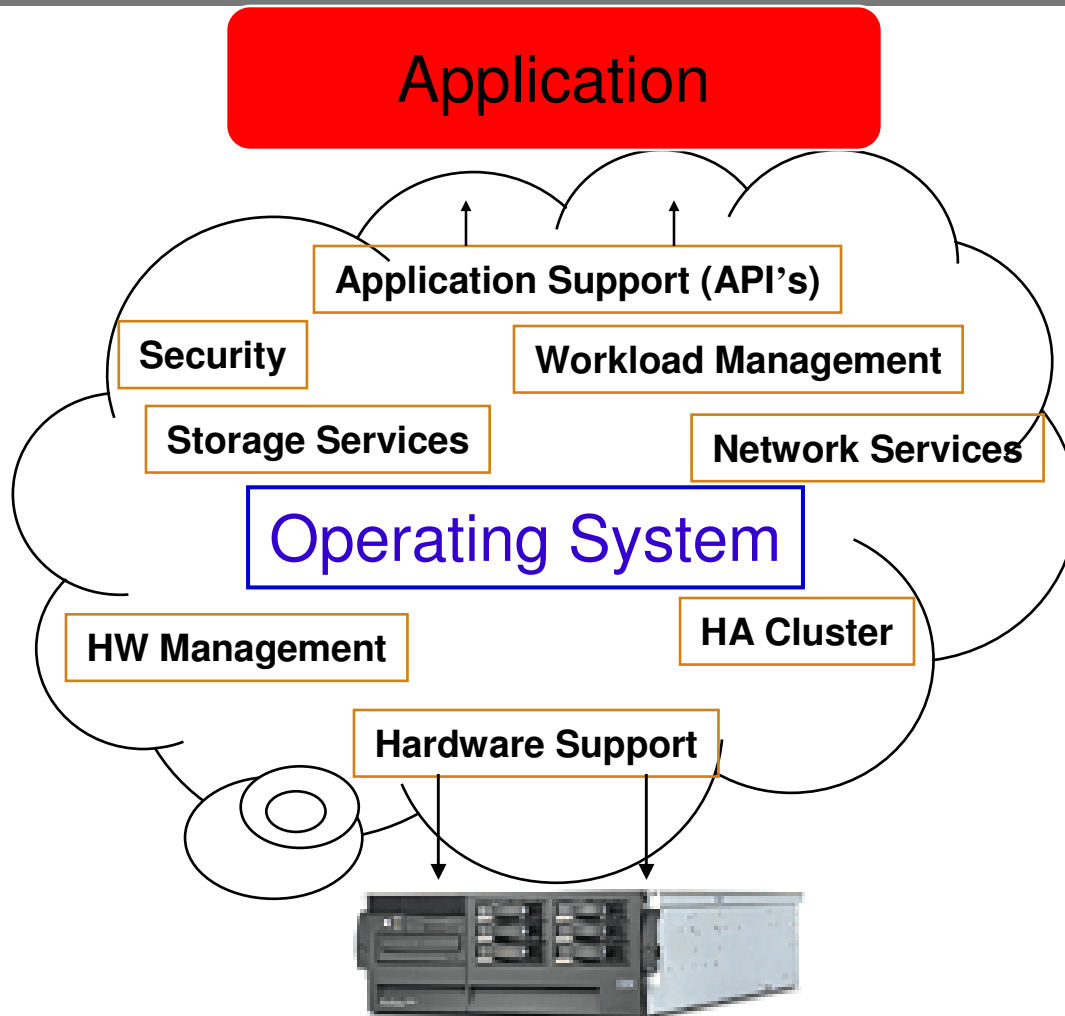
Virtue stands in the middle. Virtue is in the moderate, not the extreme position. (Horace)



- Why don't we take the best of both worlds ?
- Why don't we take a "hardware appliance" and virtualize it ?

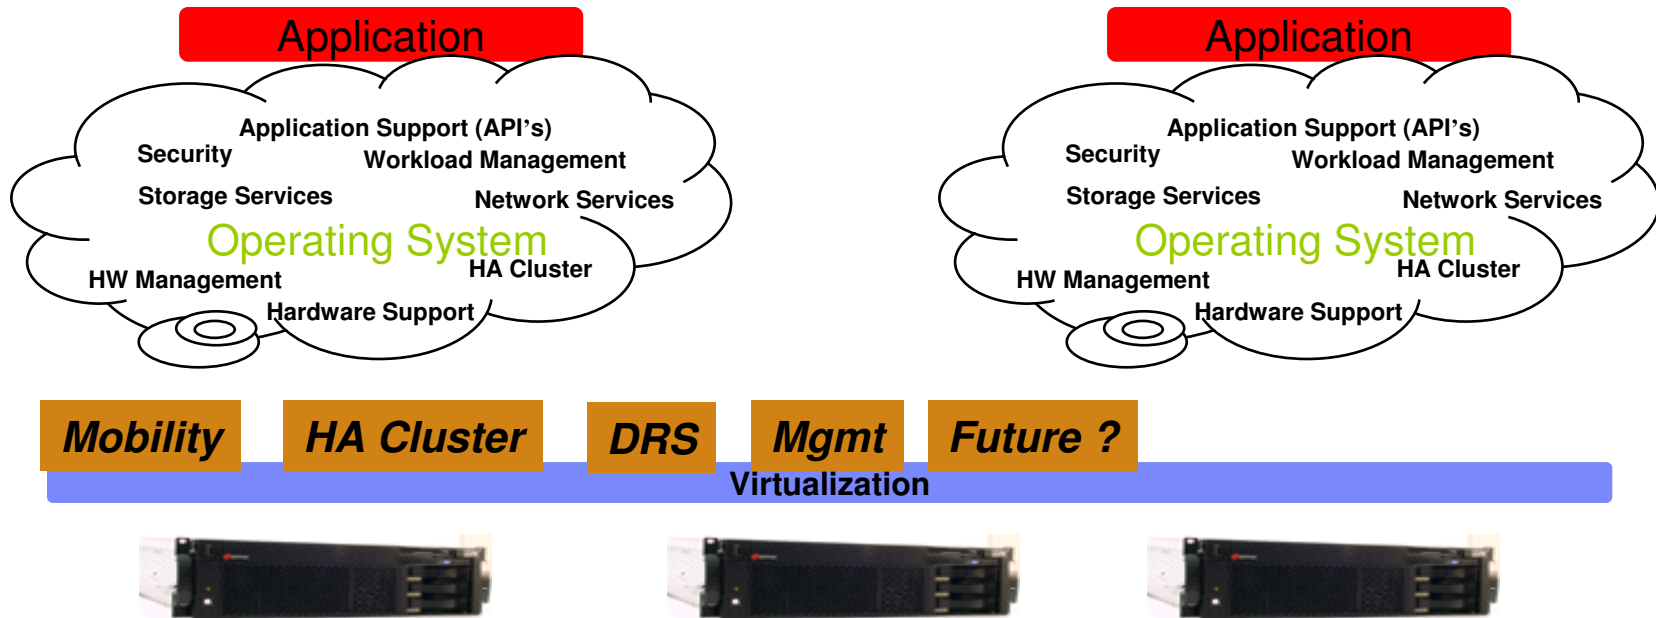


The current Open Software Stack



- › A default Win2003 install is roughly 2GB
- › Windows 2003 has roughly 50/60 components/services
- › Most of this code is not even used at run time
- › Lots of dependencies among application, hardware and OS services
- › Most of the time the application is bound to and meant to “possess” the whole stack
- › Causes server sprawl
- › Application migration problems

But have you noticed what's happening?



- Many typical “OS services” are draining into the virtual infrastructure
 - Fast becoming a “composite OS”
- this is diminishing, or at least changing, the role of the OS (as we know it today at least)

Virtual Appliance continuum(1)



› Virtual appliance

– Bi-directional delivery channel

- Entry level, restricted function version
 - i.e. *EMC Avamar virtual edition*
- Up-sell to full external appliance
 - i.e. *EMC Avamar datacenter edition*
- Use same code-base, packaging

– Consolidation option, one or more servers running multiple virtual appliances

- Deployed into common OS virtualization – WPARs
- Deployed using own Server packaging, ie custom Linux
- i.e. *Cisco vFrame* – network, security, XML virtual appliances in network “server”

Virtual Appliance continuum(2)



› Virtual Appliances

- Easy install, for SMB on existing server
 - Does not require new rack, BladeCenter
 - If Software Appliance, does not update, disrupt existing OS/Apps
- Efficient use of existing infrastructure
- Exploit Capacity upgrade on-demand
 - Exploit available capacity
- Use accelerator or co-processor Firmware “masked” operation
 - Flexible charging option

Summary



- › **Utility computing infrastructure can be implemented now!**
 - Build from components
 - Use industry standard, services based management

 - Deploy solution offerings
 - IBM Bluecloud
 - Use Grid providers

 - Focus on implementations
 - Need to understand application needs
 - Not just infrastructure

- › **Traditional OS functions**
 - Implemented in infrastructure fabric
 - Composite OS
 - OS constructed from distributed software and fabric services

Service Oriented Infrastructure



› IBM Virtualization and the On Demand Business – Redbook

- Book from 2004, a little dated but Chapter-3 sets out Service Oriented Infrastructure
- <http://www.redbooks.ibm.com/abstracts/redp9115.html>
- Or Google for “Virtualization and the on demand business”
- Author Chapters 1-4

