Welcome to NC State University and IBM– Cloud Computing Summit

Janis Landry-Lane, IBM World-Wide Deep Computing

Andy Rindos, IBM Center for Advanced Studies

Agenda– Thursday, March 26th

8:30 am - 8:45 am	**Welcome**
8:45 am - 10:25 am	**VCL and HPC at NC State (overview/demo) - Mladen Vouk et al (NC State)**
10:25 am - 10:45 am	**VCL at NCCU & the TTP - Cameron Seay (NCCU)**
10:45 am - 11:00 am	**BREAK**
11:00 am - 11:15 am	**VCL at ECU - Wendy Creasey (ECU)**
11:15 am - 11:40 am	**Cloud Computing for K-8 Pedagogy (VA VCL) - Mark Gardner (VA Tech)**
11:40 am - noon	**SURAgrid: A Regional Community Cyberinfrastructure - Gary Crane (SURA)**
noon - 1:00 pm	**Working Lunch**
1:00 pm - 1:30 pm	**IBM university relations in cloud computing - Andy Rindos (IBM, RTP CAS)**
1:30 pm - 2:30 pm	**Blue Cloud strategy and Tivoli solutions - Pratik Gupta (IBM Tivoli)**
2:30 pm - 3:15 pm	**WebSphere Clouds & the Apache project - Matt Hogstrom (IBM WebSphere)**
3:15 pm - 3:30 pm	**Break**
3:30 pm - 4:15 pm	**Blue Cloud and University Collaborations - Dave Doria (IBM STG)**
4:15 pm - 5:00 pm	**IBM datacenter/cloud computing services -- Craig Nygard (IBM GTS)**
5:00 pm - 5:30 pm	**First day wrap-up**

Agenda– Friday, March 27th

8:00 am - 8:45 am	**Arrival and Discussion**
8:45 am - 9:30 am	**Cloud computing support of K-12 in NC - Phil Emer (NC State Friday Institute)**
9:30 am - 9:45 am	**BREAK**
9:45 am - 10:30 am	**IBM Research and Cloud Computing - Vas Bala (IBM Research)**
10:30 am - 11:00 am	**Wrap-up**
11:00 am	**Depart for IBM RTP**

Agenda– Friday afternoon, March 27th (optional)
Please sign up!

12:00 pm – 12:15 pm	**WELCOME, Amy Freeman (working lunch provided by QLOGIC)**
12:15 pm - 1:15 pm	**iDATAPLEX DEEP DIVE, and A COMPARISON of iDATAPLEX to OTHER OFFERINGS, Karl Dittus**
1:15 pm -- 1:45 pm	**Tour of the iDATAPLEX facility, with Karl Dittus**
1:45 pm – 2:00 pm	**BREAK**
2:00 pm – 2:45 pm	**Data Center Assessment and Planning for Best Practices, Brett Lehman**
2:45 pm – 3:45 pm	**Optimizing the Storage Subsystem for iDATAPLEX, Andy McNeil**
3:45 pm -- 4:00 pm	**Pricing exercise, next steps for NSF Submission, Jay Bonanno**

The Cloud™ Promise……

"Deliver Cloud Computing as a utility, by making software even more attractive as a service…… and deliver this over the Internet. With Cloud, you have the illusion of infinite resources, pay as you go, and rapid deployment."

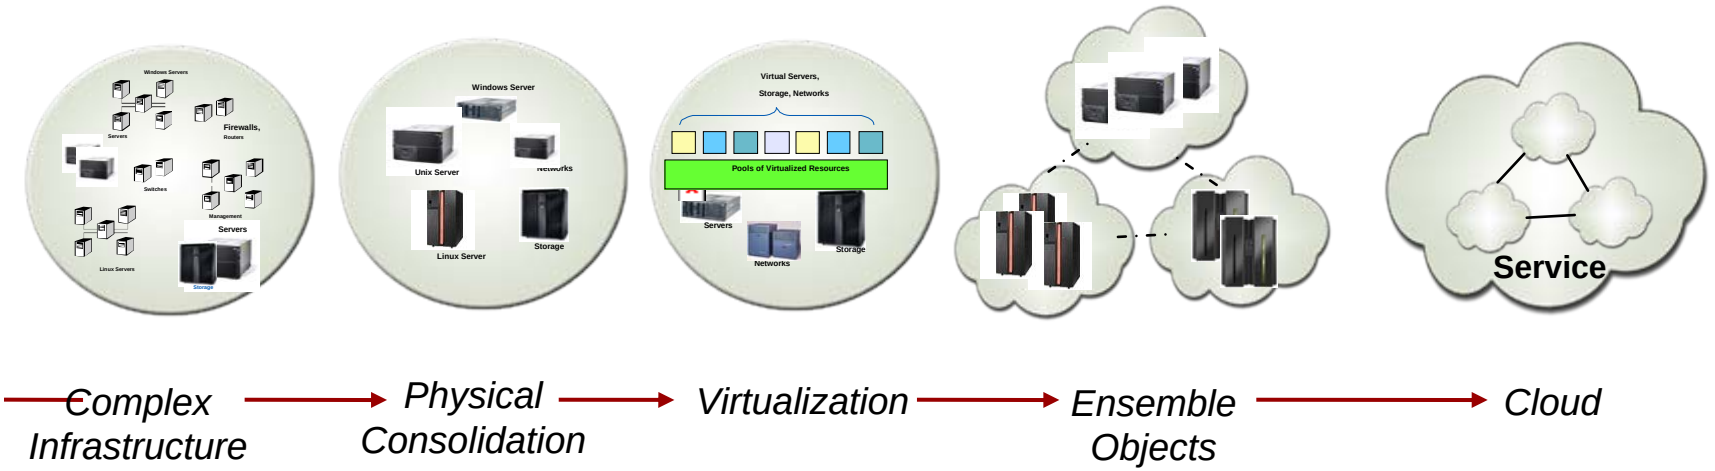

Some of the Best Practices Come from our CUSTOMERS

- **How do we effectively use all of the HPC cycles available?**
- **How do we service all our HPC needs (large scale batch) and also service our LPC needs (computing labs/running small models)?**
- **How do we maximize our investment in IT?**
  - Economies of scale (reduce costs)
  - Utilization (of existing investment)
  - Services to build the next generation of computational skills

IBM's Research Cloud Computing Initiative

- **Architecture and management of**
  - Server ensembles
  - Storage ensembles
  - Network ensembles
- **Virtualization**
  - Image catalog
  - Security/trusted virtual data center
  - Solution composition
  - Solution deployment
- **Large scale processing**
  - Hadoop/MapReduce
  - Distributed storage/Unstructured data
  - Analytics/System S
- **Federation of clouds**
  - Across geographies/administration
  - Dynamic adjusts to workloads
- **Collaboration services and infrastructure**
  - Collaborative applications
  - Business mash-up platform
- **Enabling cloud computing services**
  - Service creation, migration, deployment
  - Global delivery, virtualized data centers
- **Living lab**
  - In support of our own researchers
  - Experimenting with new technology
  - Reference center for customers

The Evolution of Cloud Technology

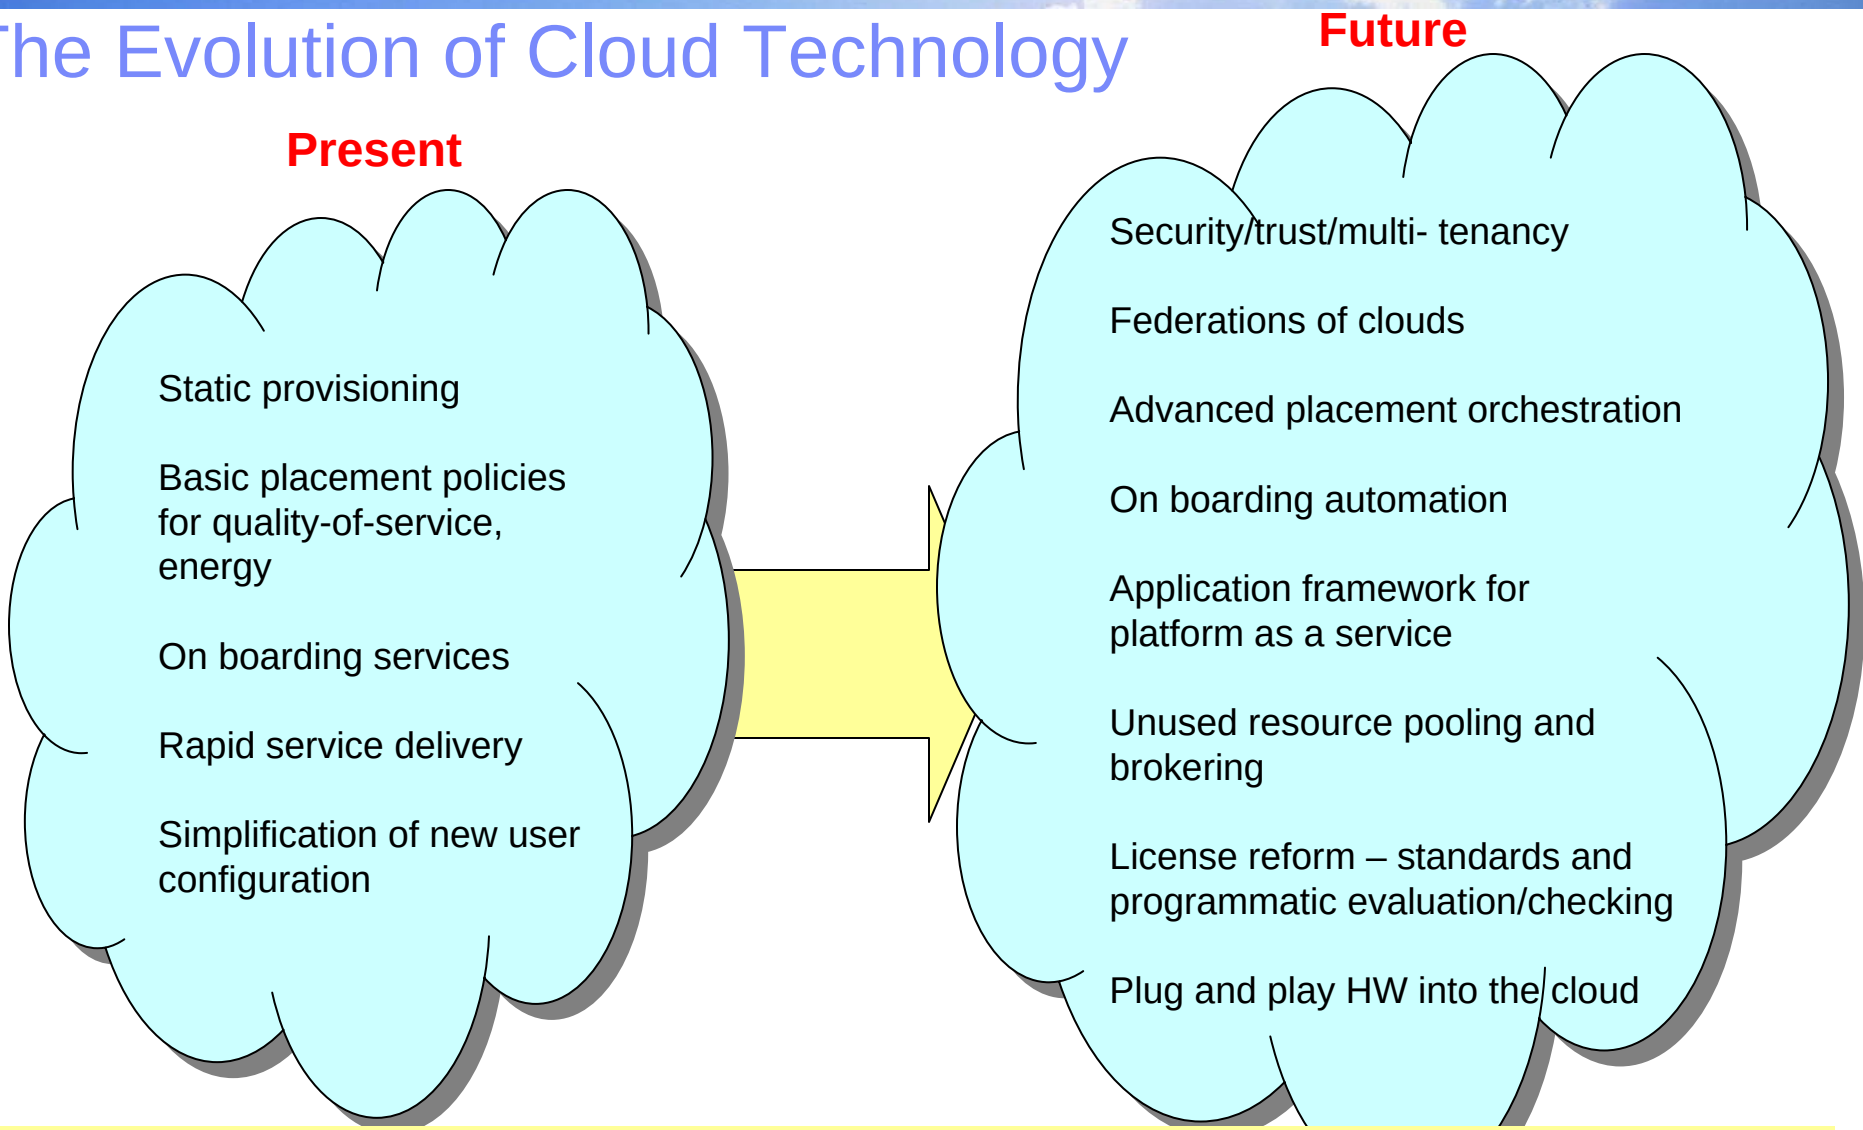

- *Economies of scale will be achieved through increases in dynamic provisioning, hibernation, over provisioning*
- *IBM products and services for automatic and seamlessly transition to cloud resources*
- *Hardware optimized and configured for cloud deployments*

The IBM Cloud Computing Centers serve clients around the world

Source: HiPODS

Some say, “CLOUD COMPUTING IS JUST GRID COMPUTING WITH LIPSTICK”, now we will find out!

Thank You!

NC STATE UNIVERSITY

IBM

Using VCL to Power "Clouds"

Sam Averitt, *Director of the Center of Excellence for Cloud Computing*

Andy Kurth, *VCL Development Team*

Aaron Peeler, *VCL Manager*

Henry Schaffer, *Professor of Genetics, and member of the VCL Special Projects Team*

Eric Sills, *Assistant Vice-Provost for Research Computing*

Sarah Stein, *Professor of Communications and member of the VCL Special Projects Team*

Josh Thompson, *VCL Development Team*

Mladen A.Vouk, *Professor and Department Head of Computer Science, and Associate Vice-Provost for Information Technology*

North Carolina State University, NC 27695
Raleigh, NC, USA

Outline

- Overview
  - VCL, Clouds and Service-Oriented Computing
  - Economics of Cloud Computing
  - Architecture
- Demo
- High-Performance Computing and Some Architectural and Implementation Details

Authors would like to thank all our VCL collaborators and colleagues and especially Michael Bugaev, Dennis Norris, Shawn VanHulst, Wake Tech's Darryl McGraw, IBM's Dr. Andy Rindos, and RENCI's Patrick Dreher.

http://vcl.ncsu.edu

VCL and Cloud Services

Virtual Computing Laboratory is an award winning* open source innovation.

http://incubator.apache.org/projects/vcl.html

(*) 2007 Computerworld Honors Program Laureate Medal for technical innovation

Why Cloud Computing?

- **Interest** in cloud computing has grown significantly over the past few years both in the commercial and non-profit sectors.
- In the **commercial** sector, various companies have advanced **economic arguments** for the installation of cloud computing systems to service their clients' needs.
- **Non-profit educational** institutions have (as always) a number of **distinct needs** that need to be taken into account if the clouds are to succeed in this environment.
- Cloud Computing means many (and often different things) to different people – and that is OK.
- Bottom line is highly reliable delivery of complex on-demand and scheduled **SERVICES** (over the network)

Service Composition

"One-button" or invisible/seamless services and very user-friendly interface.

- *Reservation groups*
- *More complex images*
- *Image groups (Clouds)*
- Workflow construction
- Base-line images (e.g., VMware, Xen, XP, Linux, IBM System z, …)

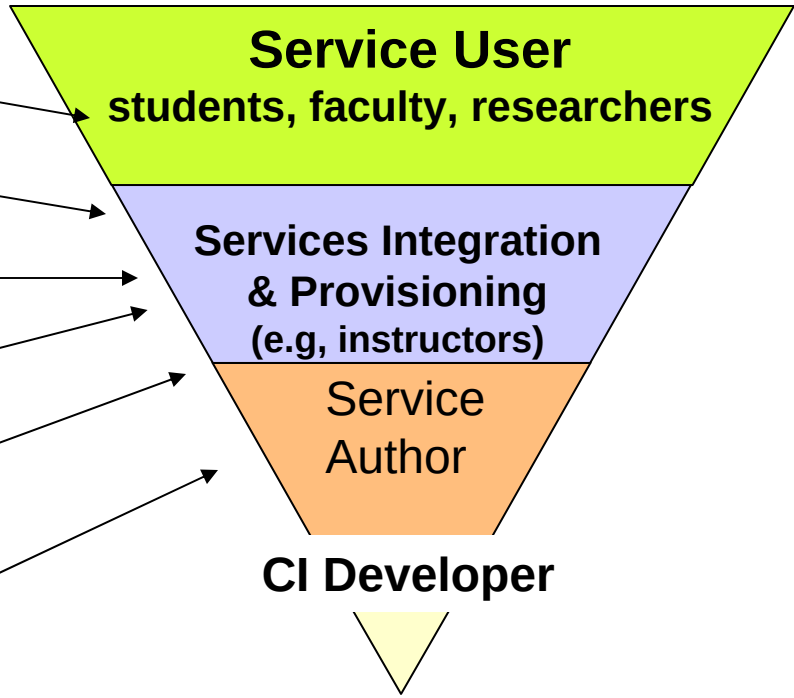

Infrastructure as a service (IaaS), Cloud as a Service (CaaS)?, Platform as a Service (PaaS), Application as a Service (AaaS),

VCL "Seat" Services

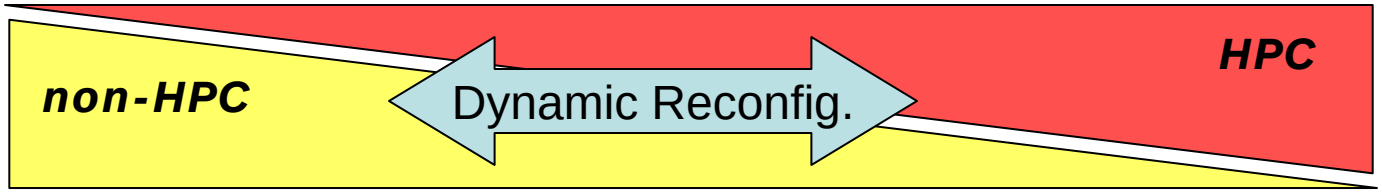

Undifferentiated Resources

Single Seat (VCL-Desktop)	Multiple Synced Seats (VCL-Class)	Servers (VCL-Server)	Environments (VCL-subCloud)	HPC (VCL-HPC)

Differentiated Resources

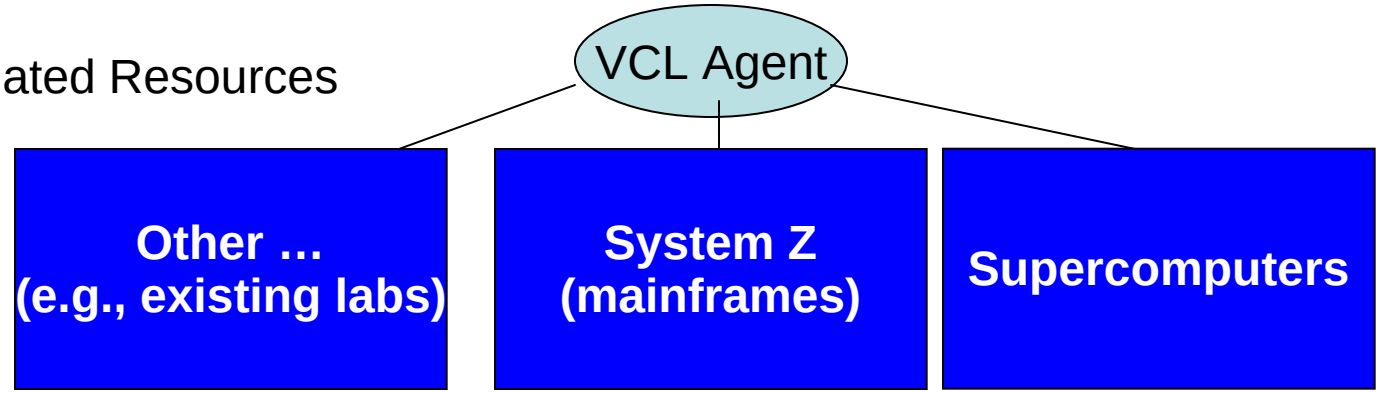

Storage

HPC and Cloud Services
(also, using Workflow technology to integrate heterogeneous resources)

Private, Public,

Homogeneous and Heterogeneous Clouds

Typical HPC Use of VCL

Internet

Login Node

HPC Scheduler

Job

HPC Storage

Compute Nodes

Typical "Cloud" Use of VCL

On-demand construction and reservation of clusters of homogenous or non-homogenous resources, operating systems and apps.

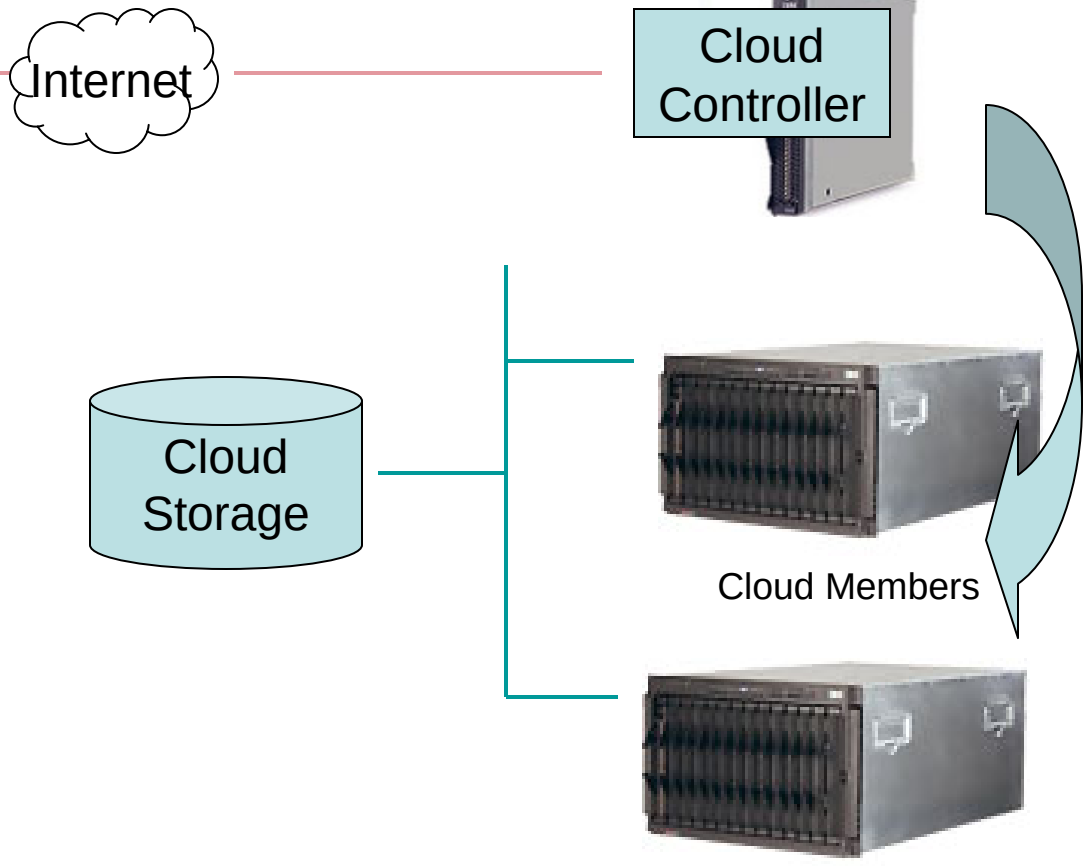

Workflows and Integration of Heterogeneous Resources

Private, Public,
Homogeneous and Heterogeneous Clouds

The preliminary analysis of the VCL data suggests that the model it is based on is very economical and cost efficient.

Business Model

Current VCL (only NC State University):

1. cca 2,000 blades (cca 1,300 to 1,500 in production)
2. open to 30,000+ students and faculty
3. cca 500 to 600 in non-HPC mode, the rest in HPC mode

VCL Usage 2004-2008

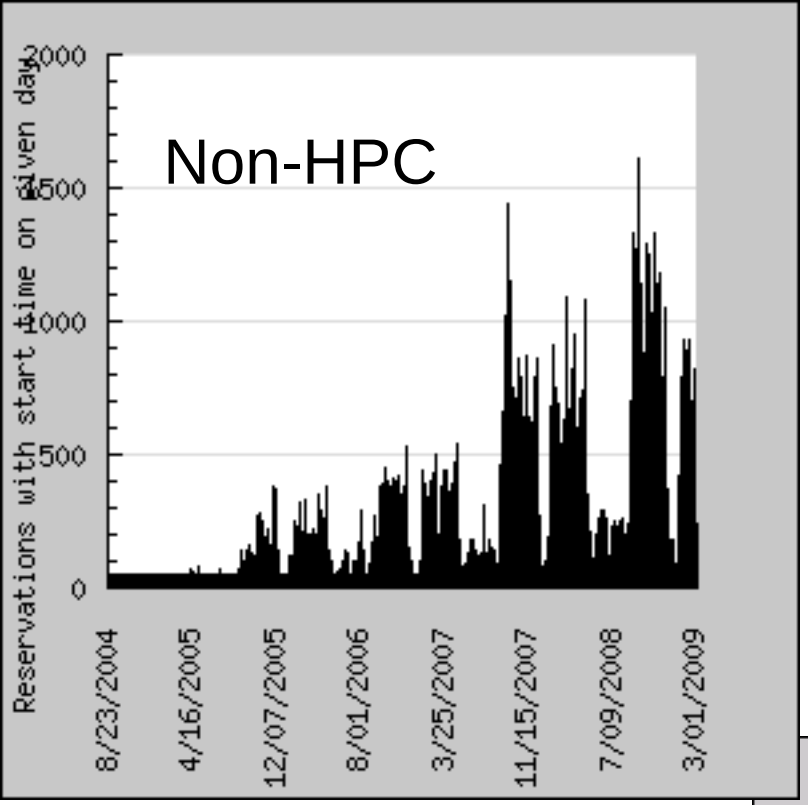

Non-HPC:

Total Reservations:	352,488
"Now" Reservations:	338,245
"Later" Reservations:	24,876
Unavailable or failed:	10,633
Failed:	5,080
Reliability:	*0.969 – 0.985*

Non-HPC Reservations:

0 - 30 Min:	132,052
30 Min - 1 Hour:	77,023
1 Hour - 2 Hours:	75,809
2 Hours - 4 Hours:	54,922
> 4 Hours:	23,315

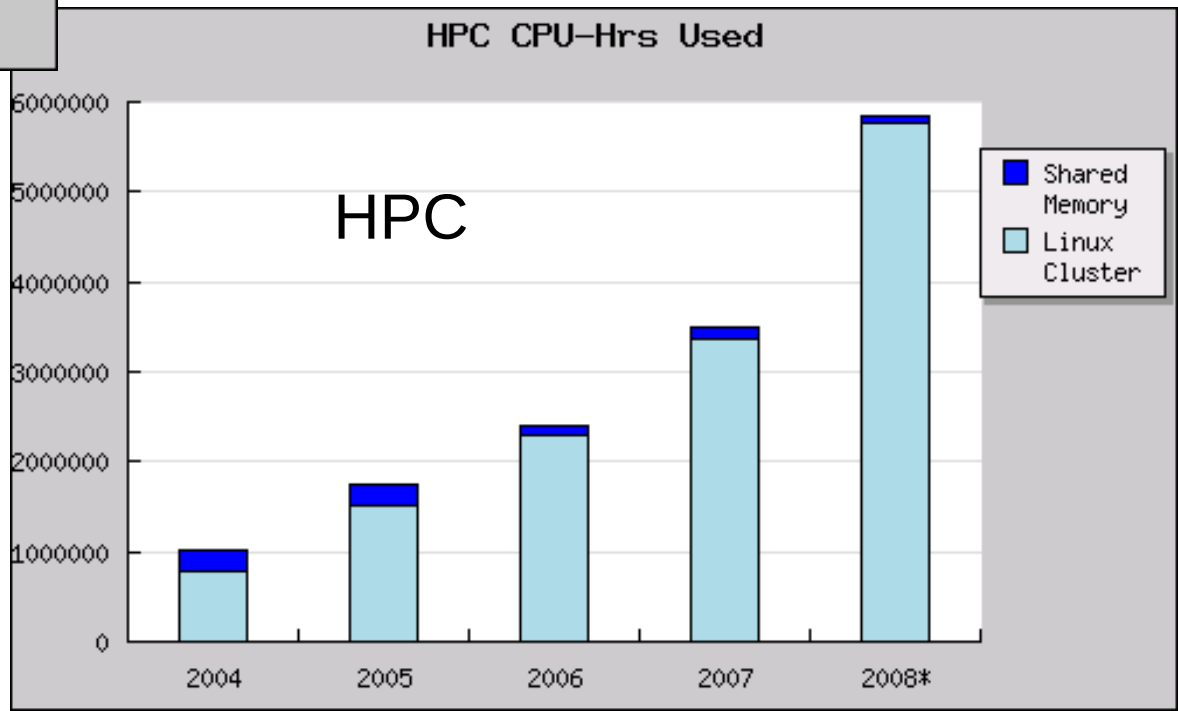

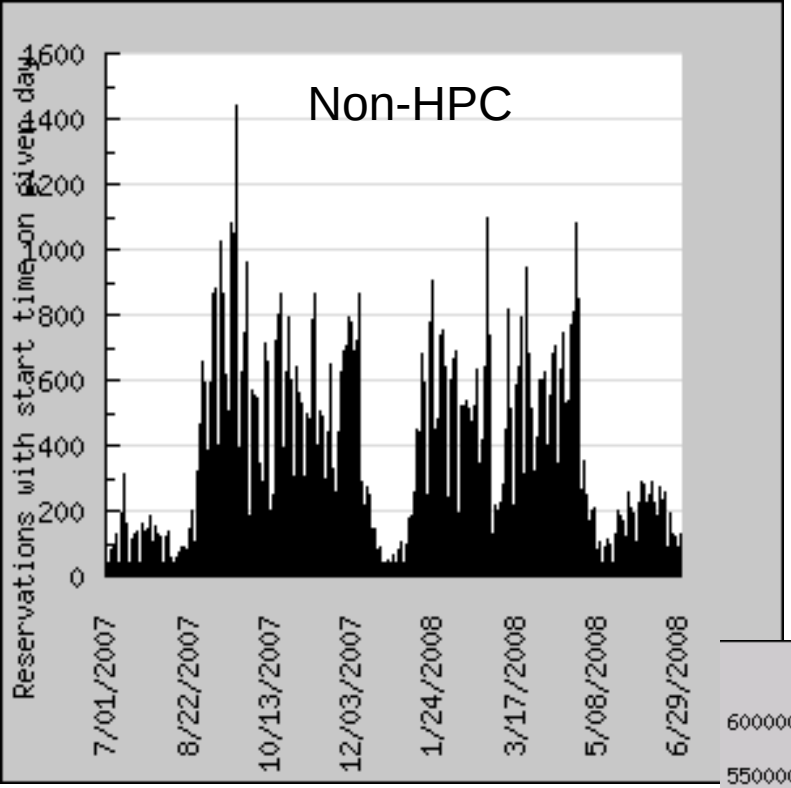

Non-HPC:

Total Reservations:	130,800
Total Hours Used:	198,583
"Now" Reservations:	125,278
"Later" Reservations:	11,436
Unavailable + Failed:	5,914
Failed:	1,611
Reliability:	*0.955 – 0.988*
Load times < 2 minutes:	109,223
Load times >= 2 minutes:	21,577

VCL Usage

1-Jul-07 to 30-Jun-08

0 - 30 Min:	48,614
30 Min - 1 Hour:	31,014
1 Hour - 2 Hours:	27,421
2 Hours - 4 Hours:	22,222
> 4 Hours:	7,443

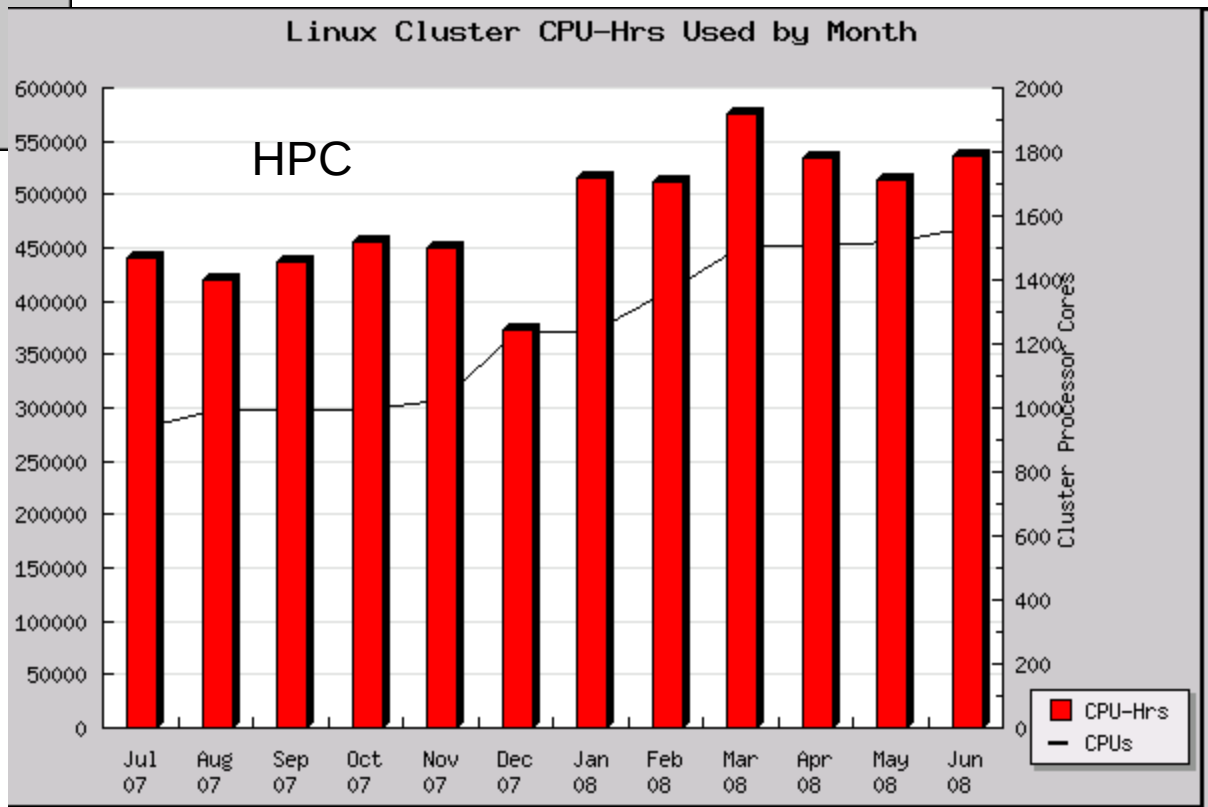

November 2008 Non-HPC

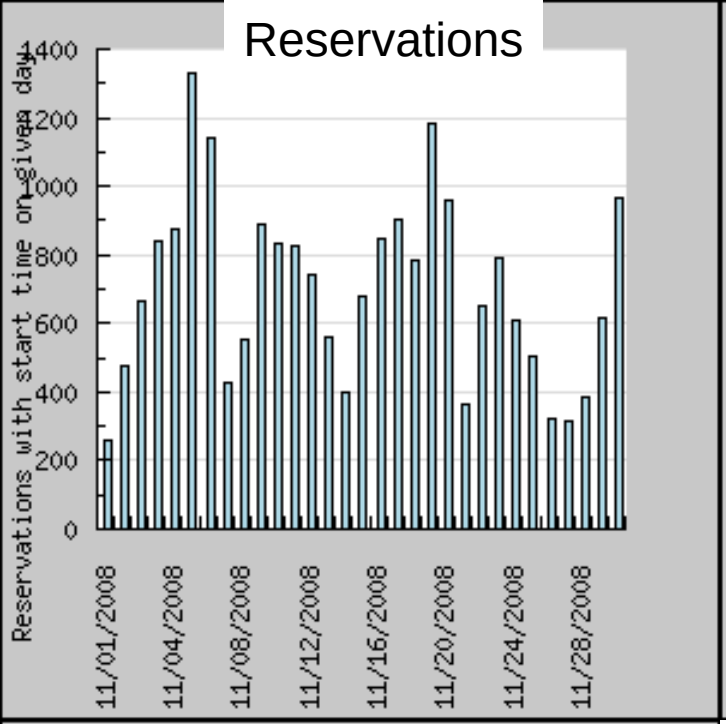

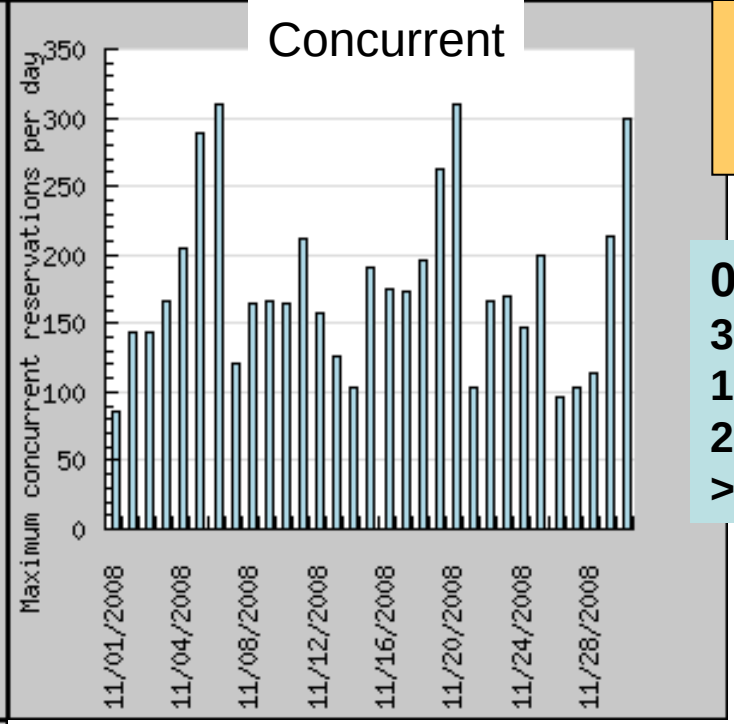

0 - 30 Min:	5959
30 Min - 1 Hour:	5069
1 Hour - 2 Hours:	5604
2 Hours - 4 Hours:	3224
> 4 Hours:	1847

cca 500 blades

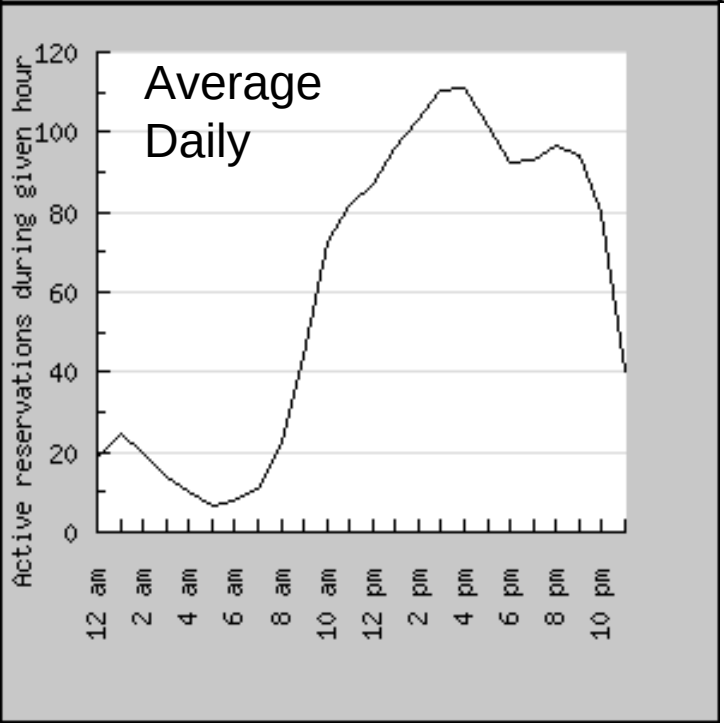

Total Reservations:	20,686
Total Hours Used:	31,853
"Now" Reservations:	19,770
"Later" Reservations:	1,933
Unavailable + Failed:	1,017
Failed:	429
Reliability:	*0.950 - 0.979*
Load times < 2 minutes:	17,013
Load times >= 2 minutes:	3,673
Total Unique Users:	4,095

Case-Study: Wake Tech Community College Resource Consolidation

- 60,000 students
- Pilot project with cca 800 students
  - Some introductory class laboratories.
  - Using VCL with about 60 blades, no bare-metal loads (virtualization using VMware)
- Lab cost savings: cca 50%

Cost Factors

- **Utilization** (70-80%)
  - Virtualization
  - Timesharing, e.g., Lab spaces (25:1)
  - Load-smoothing – e.g., in 2008/09 cca 160,000 non-HPC reservations, cca 7 million HPC CPU hrs
- Resource **lifetime** - Refresh cycle (yearly), resource lifetime (cca 5 years) – yearly down-migration of resources
- Architectural **consolidation** savings (e.g., NCCCS)
- Reduced **administration** and maintenance costs (2-3 FTEs for about 2,000 blades)
- **Efficiency**
  - E.g., One stop shopping (augmentation)
  - E.g., Distributed burden of image creation (600+ images)
- **"Green"** - Power savings (Blades)
- Other …

The *Economics* of Cloud Computing

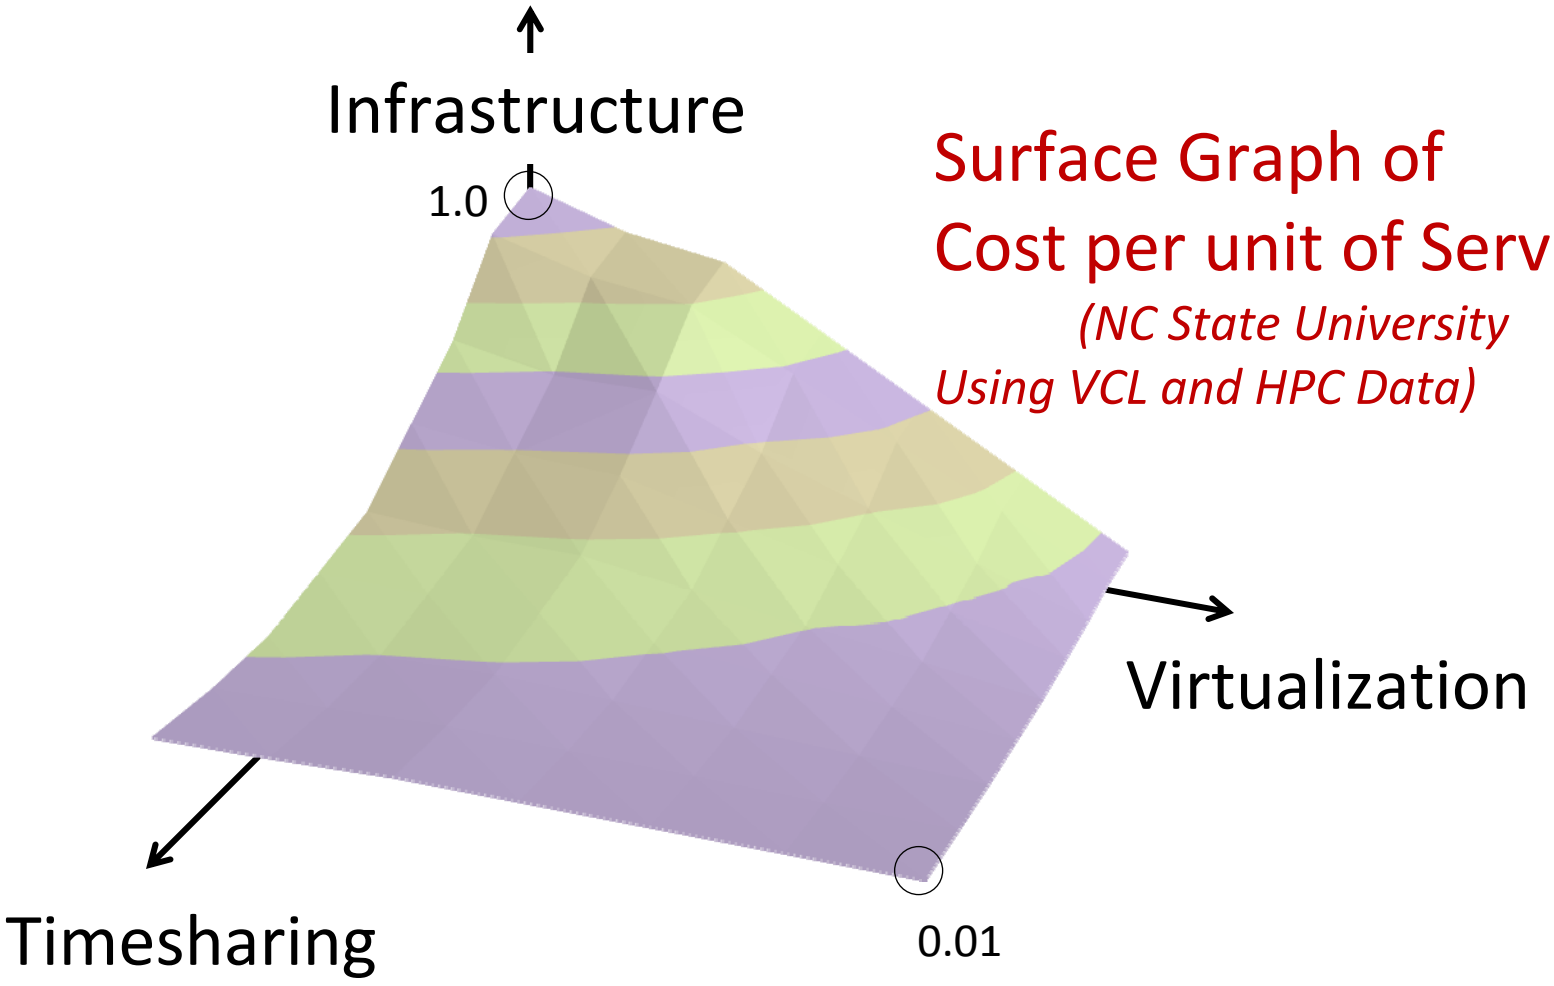

Surface Graph of
Cost per unit of Service
(NC State University
Using VCL and HPC Data)

Economics

- In 2008, about 7,300,000 CPU hours (about 7 million on HPC and about 300,000 on non-HPC) on about 1,500 blades (cca 3,000 processors) – about 1000 in HPC mode.
- High utilization on the average, but in reality low on non-HPC side (over provisioned to handle peak loads), very high on the HPC side.
- About $2 million annually (refresh, management and maintenance, improvements, personnel, …).
- About 27 cents per CPU hour (when desktop services are combined with “filler” batch or asynchronous services).
- This can come down to 10 to 15 cents per CPU hour with scale-up, large-scale virtualization, and new hardware (moving to quad-core processors).

Shades of Things to Come

Plans

- Virtualization variety (VMware, XEN, KVM, …)
- Pro-active and speculative scheduling
- Automated image construction
- Government and military-level security options (Secure Open Systems Initiative - SOSI)
- UNC build-out
- Community Colleges and K-12
- Increased performance
- Seamless resource sharing
- Modularization
- Other ...

http://vcl.ncsu.edu

Demo

Virtual Computing Laboratory is Open Source

http://incubator.apache.org/projects/vcl.html

http://vcl.ncsu.edu

Architecture

Virtual Computing Laboratory is Open Source

http://incubator.apache.org/projects/vcl.html

VCL
Manager
& Scheduler

Authentication
Service

Internet

NC State Computational "Cloud" is powered by VCL

VCL Database

Node Manager #1

Node Manager #2

Node Manager #n

Storage

Image Repository

Image Repository

Image Repository

University Labs

Tera-Grid

z-Series

Storage

Virtual or Real Differentiated
Resources

Virtual or Real
Undifferentiated Resources

Differentiator: Bare-Metal + Virtual, from Desktop to HPC on-demand, Open Source

Users

"Images"

"Service" Image Stack

Images
&
Environments
IaaS, PaaS, AaaS

VCL Manager

Resources

(Virtual) Hardware
Blades, servers, desktops, storage …

Undifferentiated Local or distributed

xCAT	VCL code	IBM TM
WebServer	DataBase	Etc.

Reliability, Component-Based, Scalability, Economy

Typical Student Computing, Desktop Augmentation, Use of VCL

Internet

Web Interface

Scheduler

VCL DB

Management Node

Server

Service Environment (Image) Library

VCL Components

- Web Interface/Scheduler
- Database LAMP (Linux/Apache/MySQL/php/perl) server
  VCL scheduler code and DB schema
- Management node xCAT & VCL management node code
- Servers Servers - physical and/or virtual to be managed by VCL

Small VCL Configuration

- 1 BladeCenter E chassis
  - 2 Ethernet Switch Modules (BNT Layer 2/3 copper)
  - Power supplies 3&4 (for 7 or more blades)
  - Chassis network module to connect management node to storage
    - Fiber Channel - Optical pass through
    - iSCSI - Copper pass through
- 2-14 HSxy Blades
  - At least on blade configured to attach to external storage for Image Library (FC, iSCSI, …)
  - Server for scheduler, database, and management node
  - Server(s) to deliver VCL services
- Storage for Images
  - FC or iSCSI storage array (few TB)

Small VCL Configuration

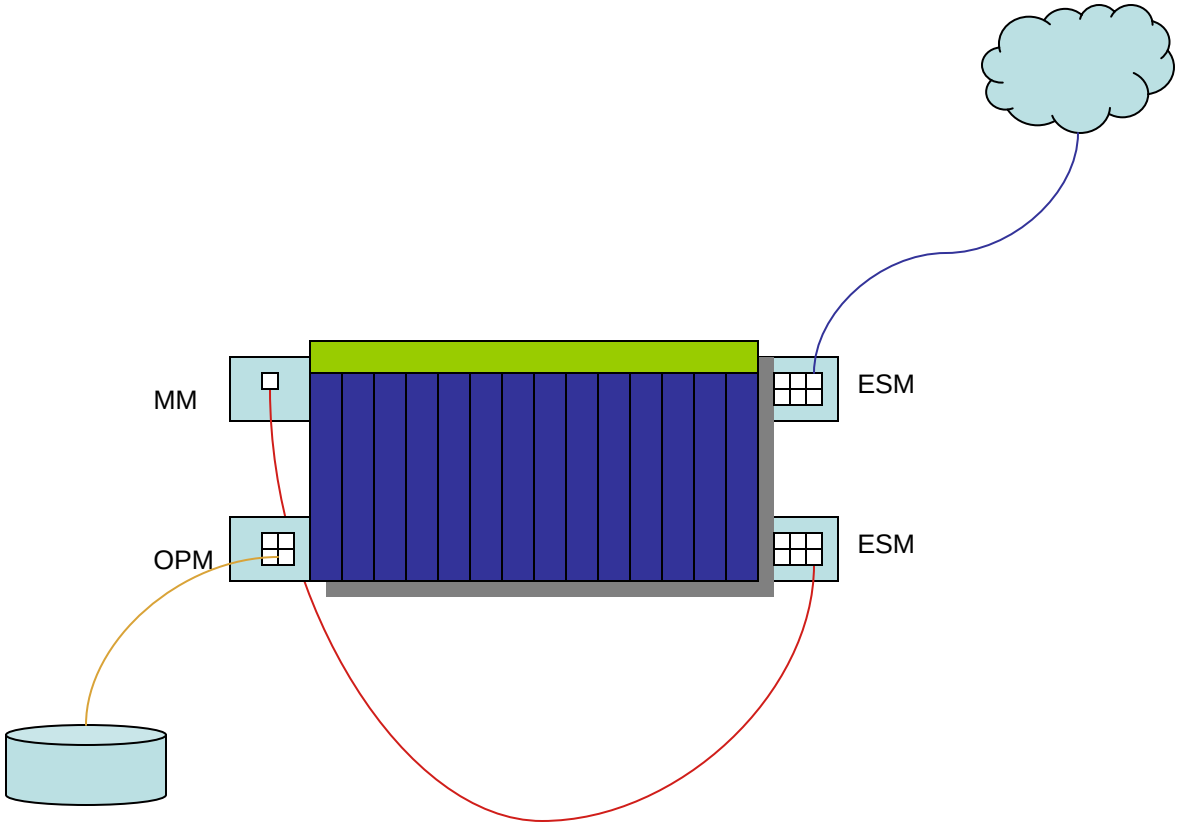

Scaling BladeCenter VCL Configuration

- Network switch
  - Cisco 6509e (or equivalent in your favorite network vendor flavor)
  - 3 separate networks (at least)
    - Network connected to Internet for user access
    - Private Network connected to VCL management node (for loading and managing images)
    - Private Management network (connecting BladeCenter Management Modules and VCL management node - controls power on/off, reboot, …)
- VCL Management nodes
  - One management node for every ~100 blades
  - Physical connection to storage array - shared file system (GFS, GPFS) for multiple management nodes at one site

Scaling VCL

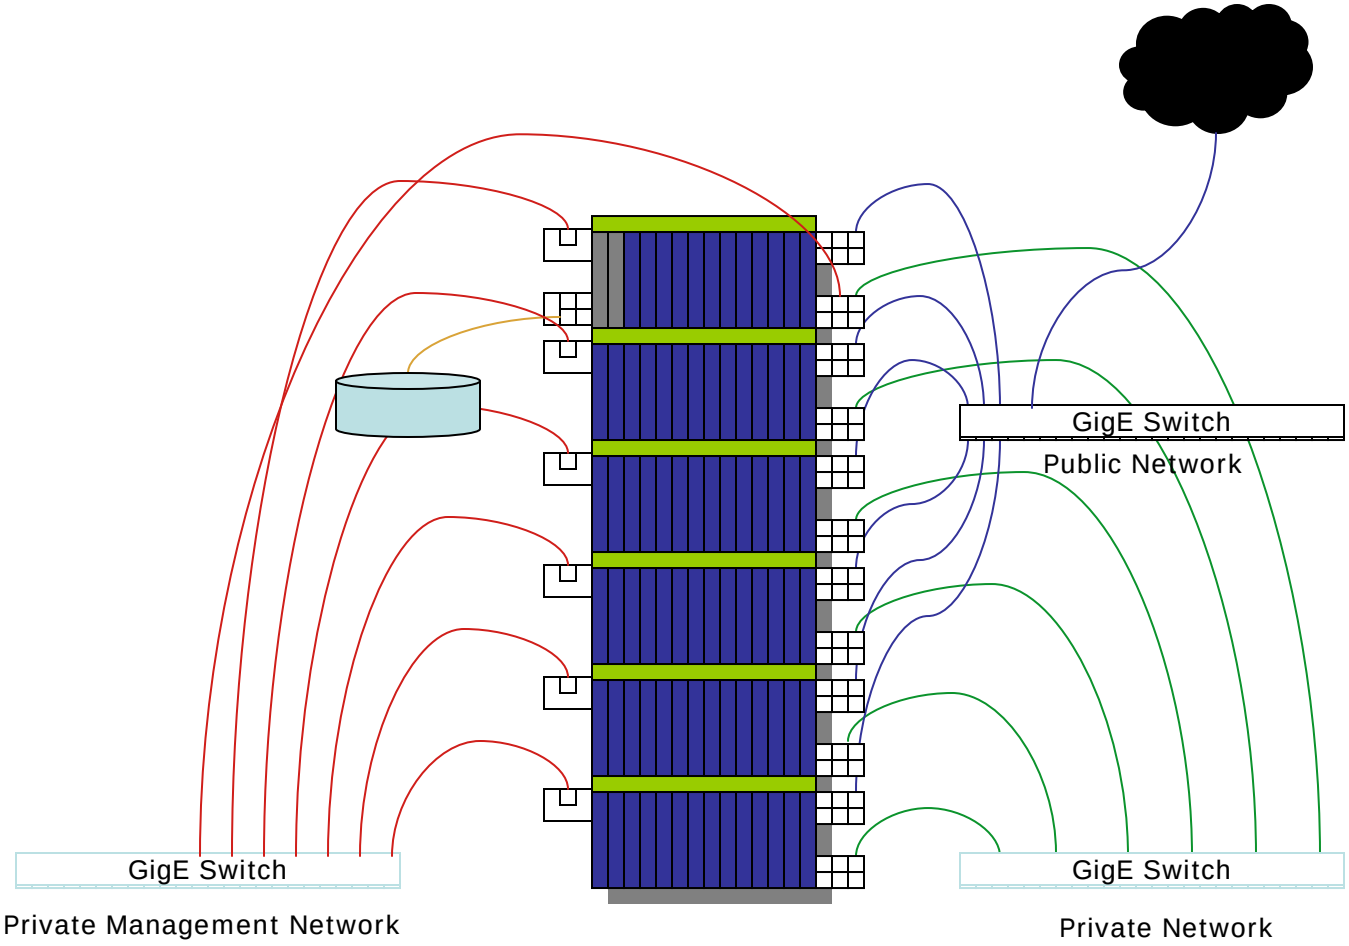
GigE Switch
Public Network
GigE Switch
Private Management Network
GigE Switch
Private Network

HPC Cluster in VCL

- Network switch
  - Add another private network for message passing traffic - use NIC that would be used for Public network user access
- BladeCenter Chassis
  - Configure two VLANs in one chassis switch module.. one for public Internet access and one for private message passing interface
- VCL management node
  - configures blade VLAN based on image metadata

HPC Cluster in VCL

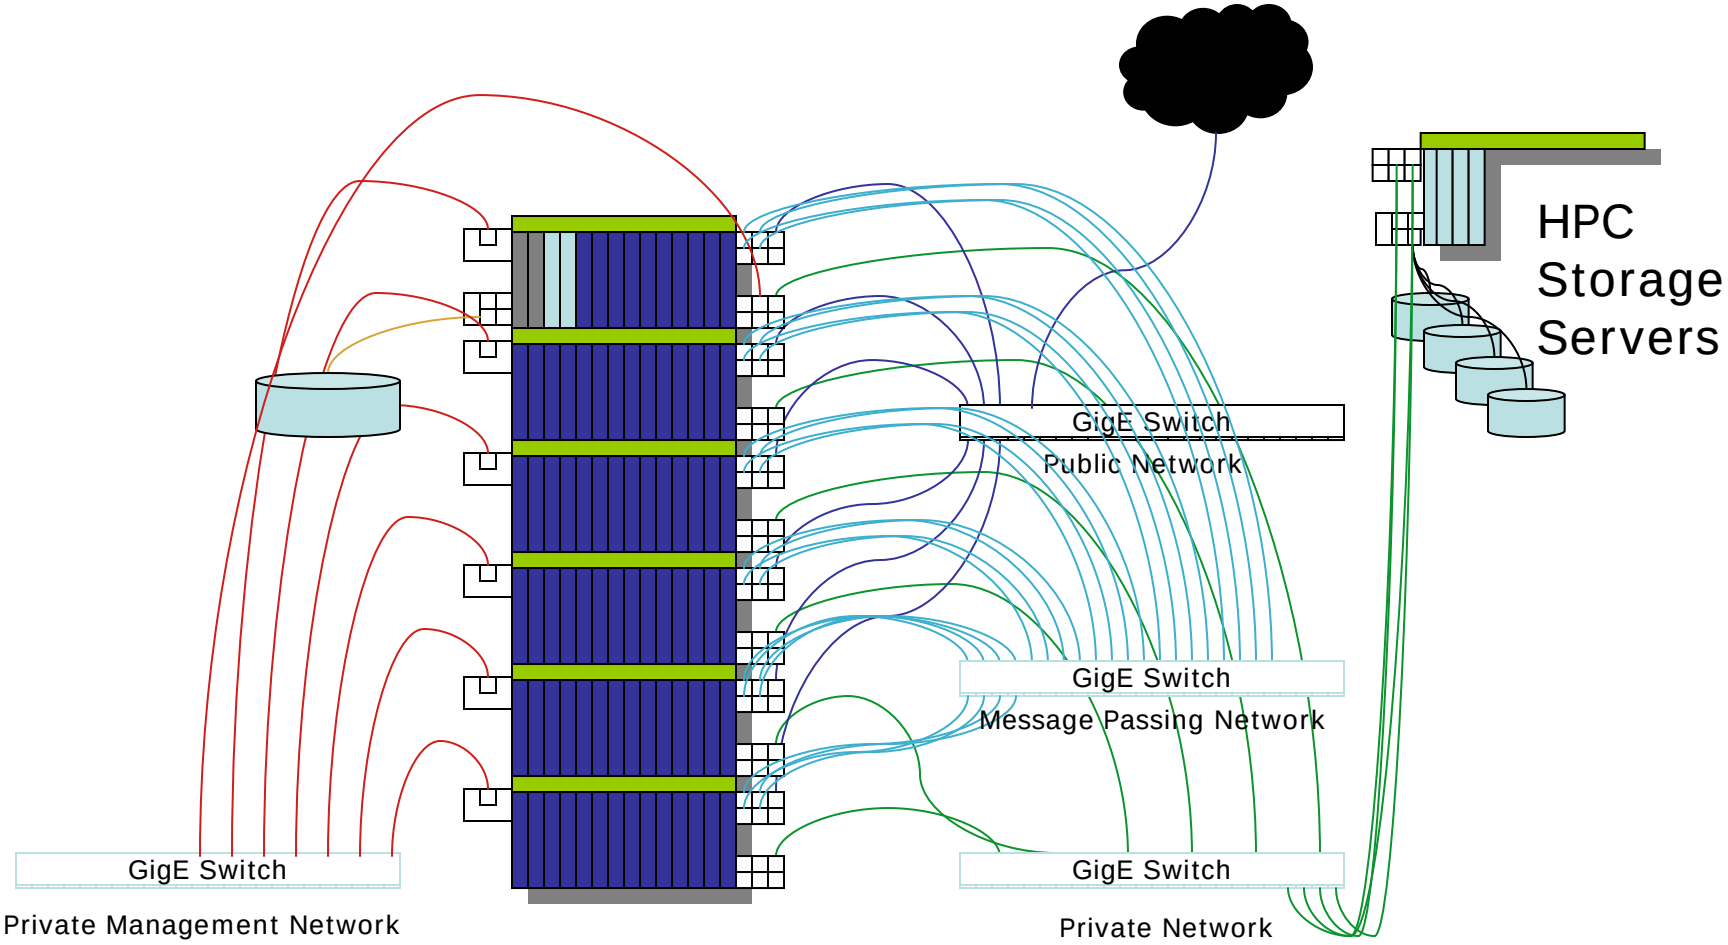

Adding Low Latency Interconnect for HPC workload

- BladeCenter chassis (not chassis housing management nodes)
  - Chassis network module for low-latency interconnect
    - Optical pass through (Myrinet, InfiniBand)
    - IB Switch
- Blade servers
  - Daughtercard for low-latency interconnect (Myrinet, InfiniBand)

Large Scale VCL Deployment

- iDataPlex - ~84 physical servers/rack
- LAMP & Management node servers
- Network switch
  - 1 less network - no separate management network port (combined with one of two GbE ports)
  - Server switches in iDataPlex rack
- Storage

VCL at NCCU and the Technology Transfer Project

Cameron Seay

School of Business

North Carolina Central University

ABOUT NCCU

Founded in 1910, North Carolina Central University is a Historically Black College (HBC) that has provided academic opportunity for the citizens of North Carolina and students from around the world

Currently, NCCU is a liberal arts university with a science and technology focus

Our current enrollment is around 7,500

I am with the Computer Information Systems program, which is housed in the School of Business (CIS has around 100 students, the School of Business has around 1,000)

ABOUT TTP

The Technology Transfer Project (TTP) was begun in 1996 to assist HBCUs in their technology initiatives

The TTP is contained within the Executive Leadership Council (http://www.elcinfo.com), a group of Executives that work with HBCUs

The schools currently involved with the TTP include:

Alabama A&M University,
Hampton University,
Florida A&M University,
Howard University,
Morehouse College,
Morgan State University,
Norfolk State University,
North Carolina A&T University,
North Carolina Central University,
Southern University, Baton Rouge,
Tennessee State, and
Tuskegee University.

VCL at NCCU:

In the Fall of 2005, NCCU began to work with NC State and IBM to develop a presence of the Virtual Computer Lab (VCL) at NCCU

In 2005 and 2006, NCCU was award successive Shared University Research (SUR) Grants by IBM to facilitate this goal

The result is a small VCL infrastructure tied directly to the VCL management logic maintained by NC State

NCCU Blade Center at Present

Housed at MCNC, RTP, NC

NC State
Management Module`

MGTModule

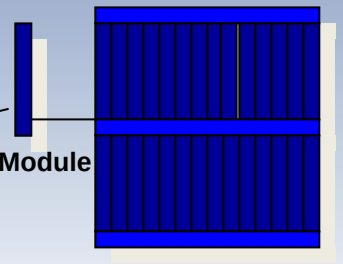

14 Blades

Chassis 2 (Virualization HPC)

MGTModule

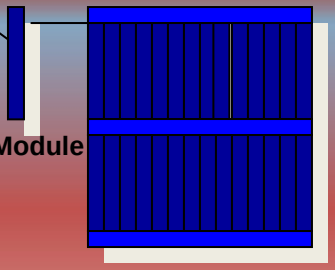

10 Blades

Chassis 1 (SandBox)

NC State Infrastructure

VCL at NCCU

Over 500 unique users since Fall 2007

45 separate images in use

Most popular applications: MS Office, SAS, SPSS, 3270, WAMP

Drs. Alisha Malloy and Donna Grant of NCCU have a project at Hillside New Tech High School involving Alice and MS Office

VCL and the TTP

Deanna Roquemore of Southern University will install a VCL site at her School

Gerald Whitaker, Hudson Defoe and Clifton Wood have created a VCL instance at Morgan State

Installations are planned for NC A & T and Howard University

The entire TTP and other HBCUs can collaborate to form an HBCU "Cloud"

The TTP Cloud can extend beyond the TTP schools to smaller HBCUs and the K-12 environment

Future of VCL/Cloud for TTP Schools

VCL and Cloud penetration will deepen and widen among TTP schools

Cloud functionality will become more seamless among the schools

System z will play a much bigger role in TTP based Cloud, especially in server and desktop virtualization

More schools will become involved, including high schools and middle schools in the communities near the TTP universities

Smaller HBCUs will be encouraged to participate

Future of VCL/Cloud for NCCU

More exposure of students/faculty to Linux and Open Source tools

Greater utilization of VCL among the academic units

Greater extension of the NCCU VCL instance to the external community (high schools, community centers, etc).

Inclusion of System z technology into the NCCU VCL infrastructure

Expansion of NCCU VCL Hardware (Sun Grant, etc)

TTP Cloud

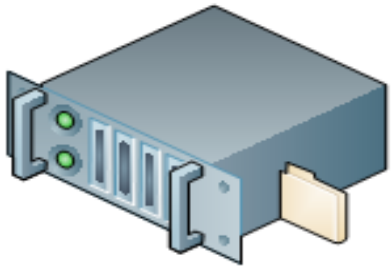

Howard
(proposed)

z10

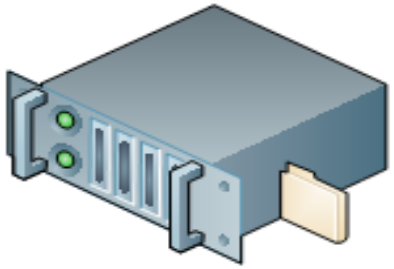

A & T
(proposed)

Shaw U.
(proposed)

Bennett Col.
(proposed)

H.S.

Middle School

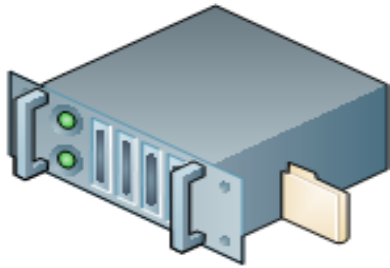

Morgan

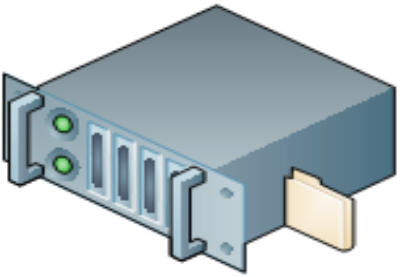

Southern

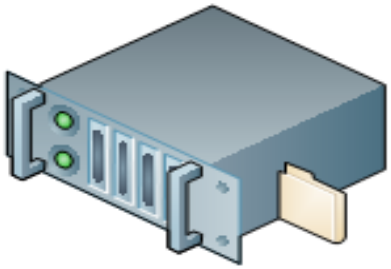

NCCU
(MCNC)

THANKS!

cseay@nccu.edu

virtual computing experiences

east carolina university

wendy creasey

partnership 2006

great support from ncsu

doubled capacity

moving from mcnc to ecu

84 blades

since august . 8511 reservations
. 13,619 hours used

Fall 2008	20 images
	28 courses
	20 faculty
Spring 2009	15 images
	23 faculty
	29 courses

interesting applications

.cisco labs – using dynamips cisco emulation – simulate cisco hardware

.big image, virtual network adaptor

.linux image using two vm's – students had to set up firewall rules and test – faculty used vnc viewer to log in and provide live assistance

.two vmware machines- simulate one machine attacking another machine

more interesting applications

.ni circuit design software

.several applications using license servers

.special education software bundle to expose de students to an array of tools.

.sql server

.linux programming course using different compliers

.rational rose

lots of standard software

.software includes adobe products, autocad, office, project, visio, sas, spss, mathematica, matlab…

future

.need fast persistent space

.learning curve when hardware is at ecu vs. mcnc

.interest in adding more hardware devices

.kvm over ip

“the vcl has allowed me to give basically all the … students access to software that otherwise was date restricted… and or impossible to install …”

faculty

“vcl is terrific. it is practically indispensible for teaching courses dependent upon software at a distance.”

faculty

“I believe vcl should be a key component to how we teach.”

faculty

“it was easy to use…it worked great. plenty of bandwidth and plenty of cpu.”

“I just have to say why didn’t this come out sooner?...”

“I was surprised how fast everything downloaded. Setup was very easy!”

students

the best part has been the support and partnership with ncsu

Cloud Computing at VT

Wu Feng and **Mark K. Gardner**

CS/ECE *Office of IT*

SyNeRG

IBM Cloud Computing Seminar
March 26-27 2009

VirginiaTech
Invent the Future

The Vision

HPC

Enterprise

Academic

VCL

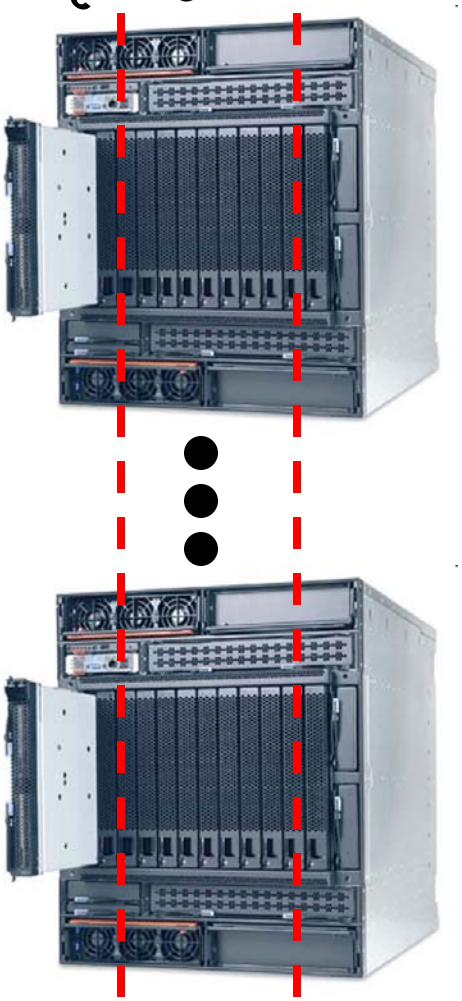

Cloud

The Roadmap

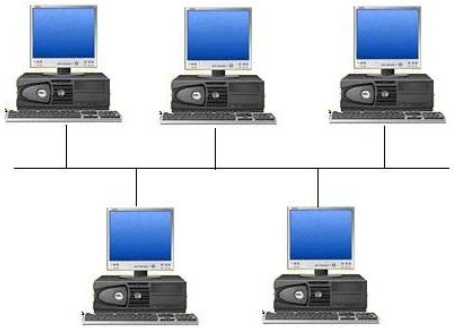

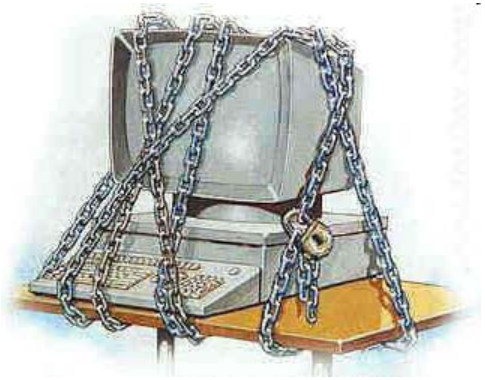

- Research
  - Network Issues
  - Security Issues
  - Rural Economic & Education Stimulus

- Pilot Project
  - Virtualization for K-8 Education
  - Remote Campus and Colleges
- Production
  - Enterprise
  - Academic
  - HPC

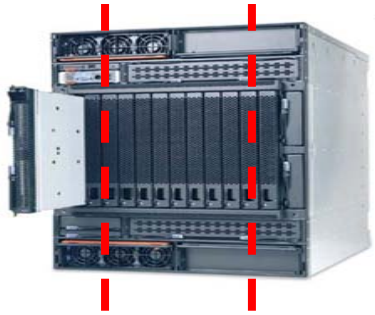

Cloud

SOUTHWEST VIRGINIA HIGHER EDUCATION CENTER

Outline

- Educational Computing
- Pilot in Rural K-8 School
- Deployment
- Conclusion

Challenges of Educational Computing

- Cybersecurity of Computing Resources
- Scarcity of Innovative Software
- Lack of Integration into Curriculum
- Insufficient Teacher Training
- Lack of Funding
- Digital Divide
  - Metropolitan schools typically “have”
  - Rural schools often “have not”

Address many of these challenges with cloud computing.

Outline

- Educational Computing
- Pilot in Rural K-8 School
- Deployment
- Conclusion

Pilot: Rural K-8 School

- Accredited Cooperative
  - Meet the needs of children and the needs of community
  - Founded by parents desiring more active role in education
  - Students have done well in transitioning to public high school
- Challenges
  - Limited funding
  - Limited computing capacity
  - Limited expertise (volunteers)

- Approach
  - Leverage cloud computing to overcome challenges
    - Adaptable environment for younger users, particularly K-8.

Teach Programming

- Class Profile
  - Fourteen 2nd and 3rd grade students
- Logo
  - Created to teach programming as early as 2nd grade

Fascination with turtle graphics

Classroom Experience with Logo

- Abstraction is tough.
- Make Logo commands concrete.
  - Bee-Bot: a modern mechanical “turtle”
  - “Kid”-Bot: take turns being turtle and programmer

From “Concrete” Logo to “Computer” Logo

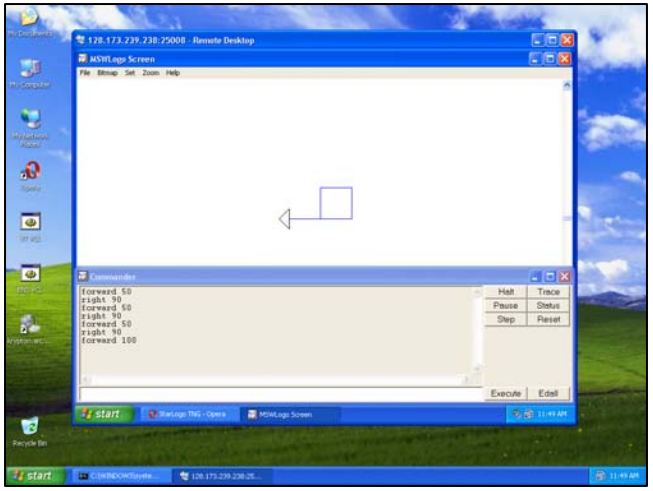

- MSW Logo, StarLogo, KTurtle, XLogo
- Preliminary results are encouraging

Storytelling Alice

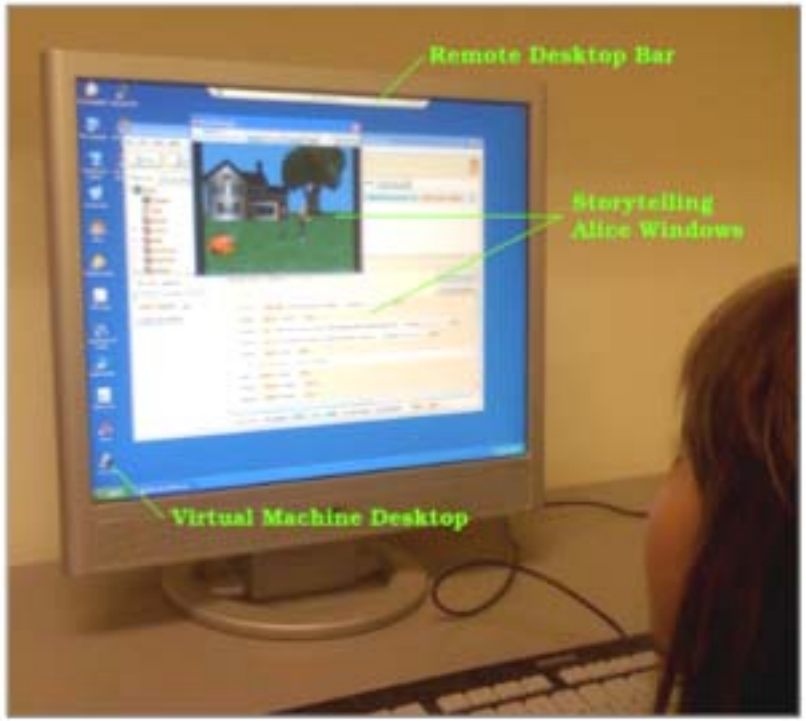

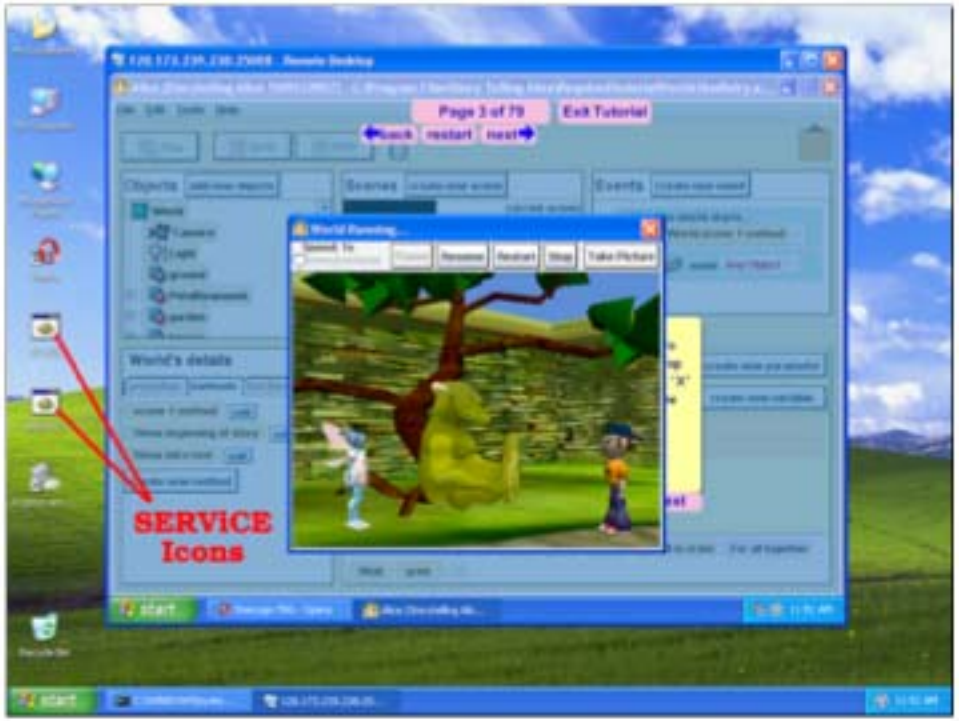

Outline

- Educational Computing
- Pilot in Rural K-8 School
- Deployment
  - User Interface Issues
  - Networking and Cybersecurity
  - Experiences
- Conclusion

Double-Click to Invoke Local Application

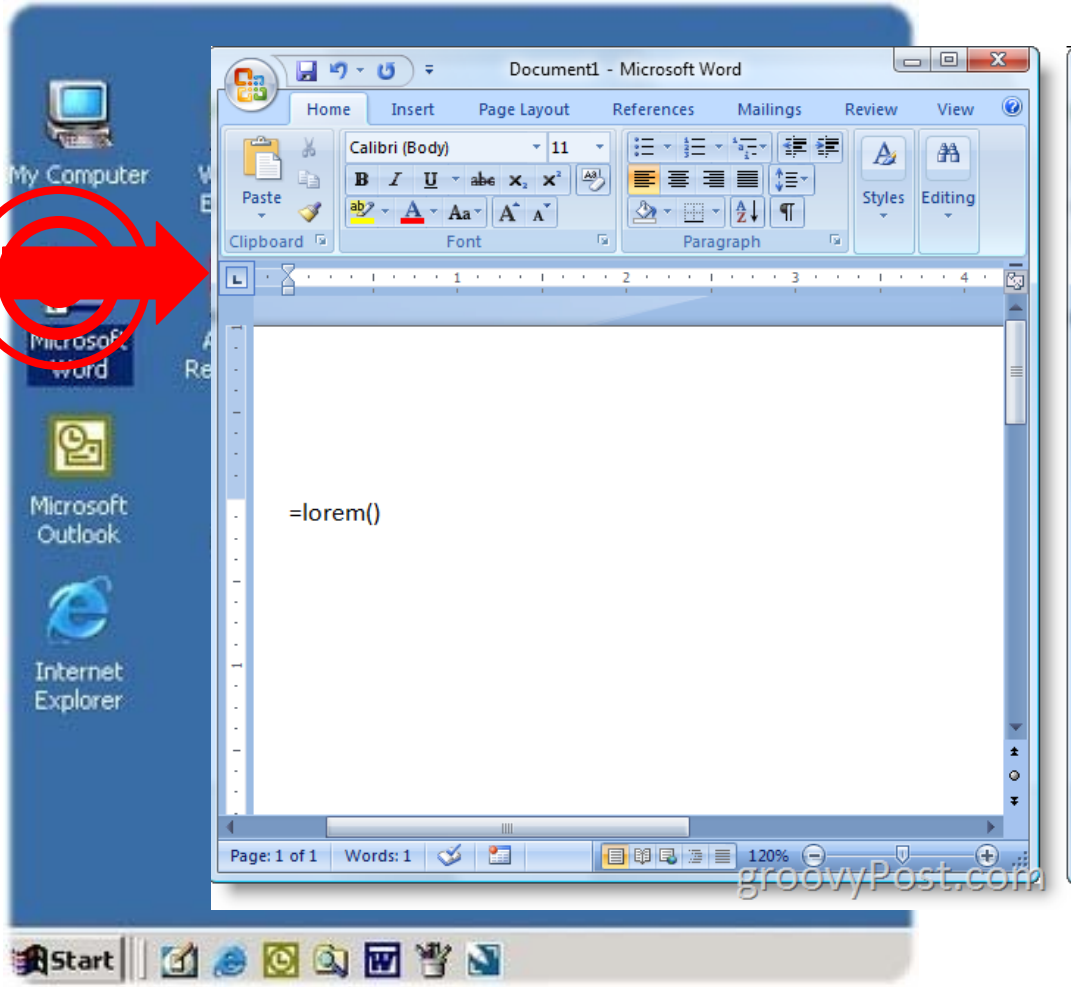

Double-Click to Invoke Virtualized Application

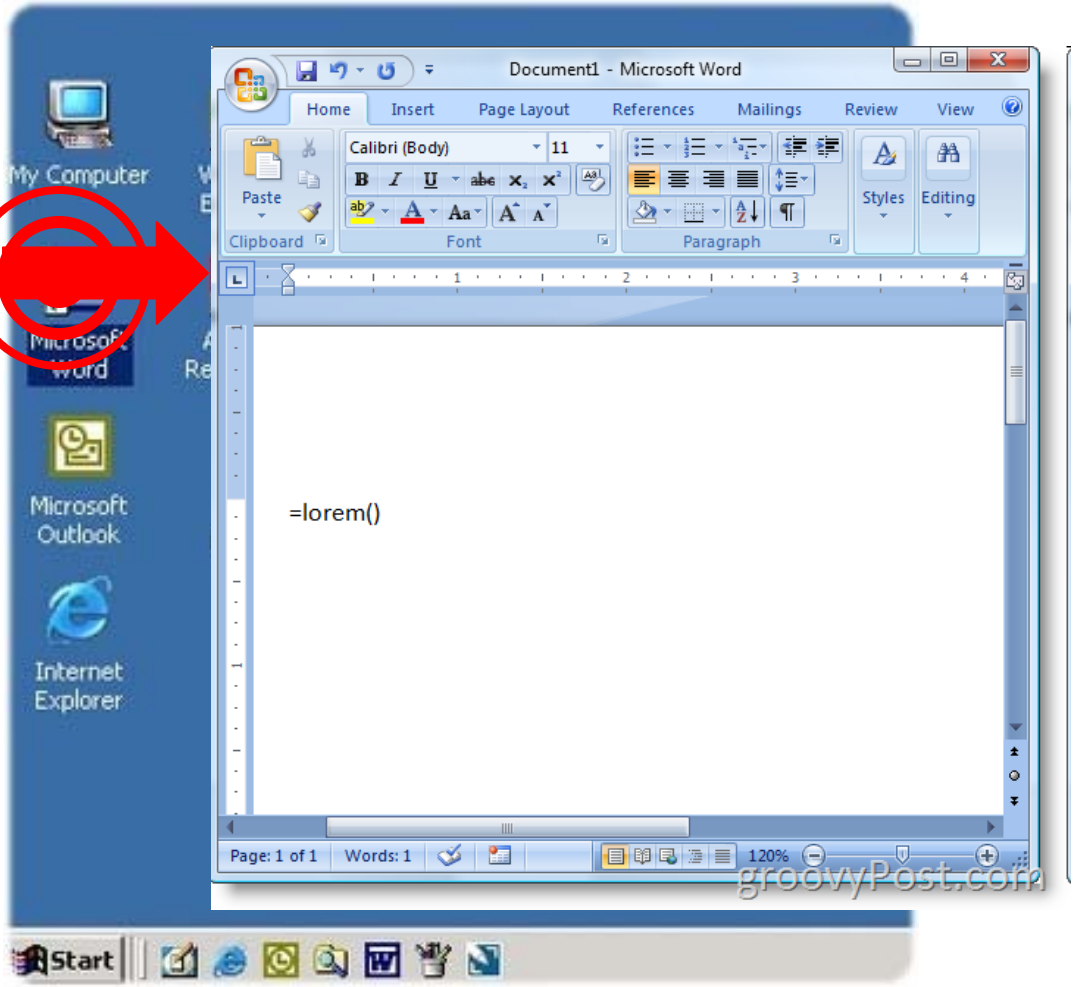

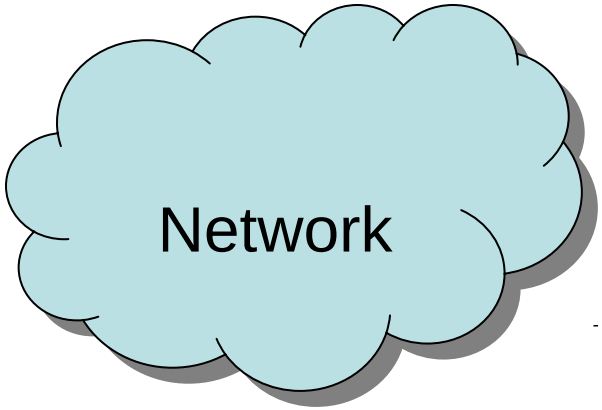

Outline

- Educational Computing
- Pilot in Rural K-8 School
- Deployment Issues
  - User Interface
  - Networking and Cybersecurity
  - Experiences
- Future Work
- Conclusion

Choice: Public vs. Private IP Addresses

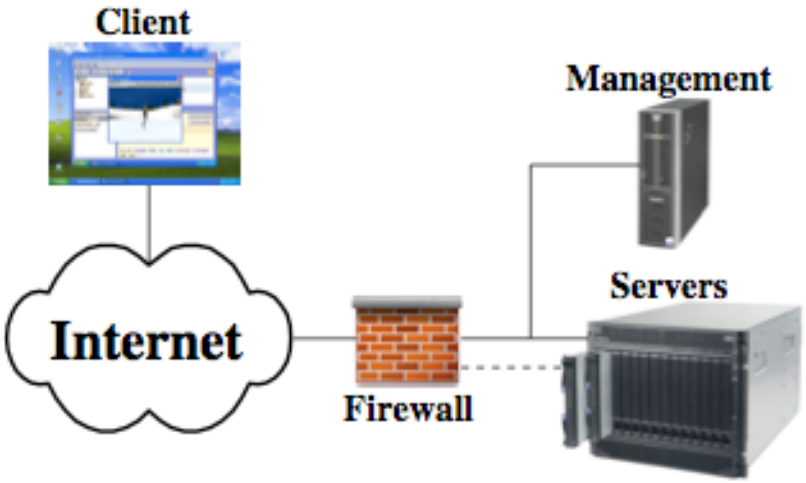

- Traditionally use public IP addresses
- Must use private IP addresses
  - Not enough public addresses
  - Security policy limits forward-facing machines
- Dynamically “punch hole” in firewall

A Solution

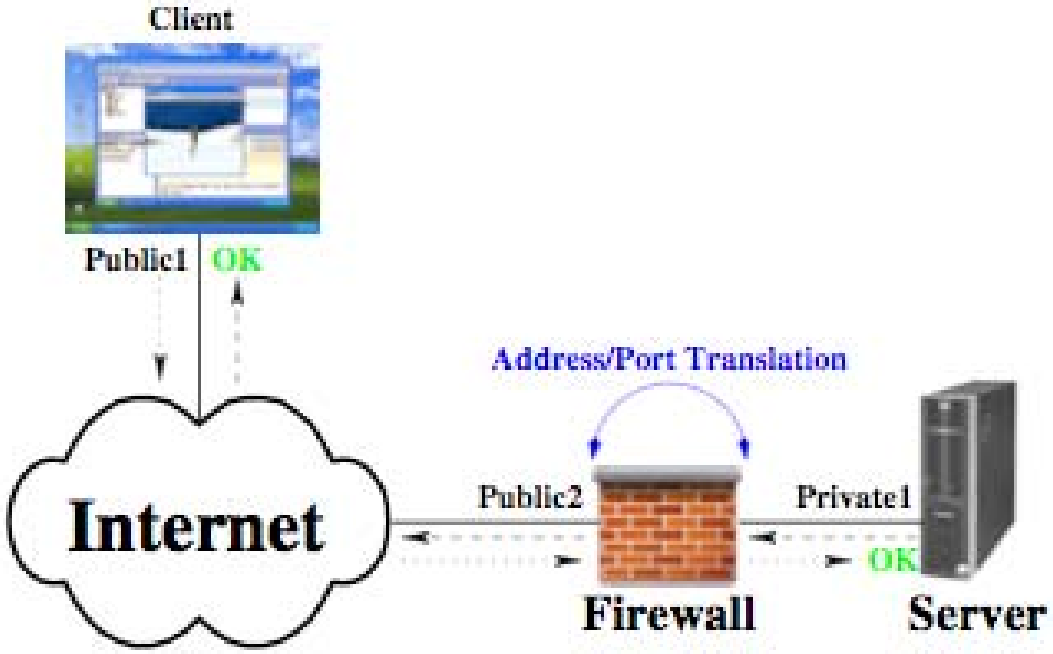

- Network Address/Port Translation (NAPT)
- Dynamically map addresses to/from Internet

Outline

- Educational Computing
- Pilot in Rural K-8 School
- Deployment
  - User Interface Issues
  - Networking and Cybersecurity
  - Experiences
- Future Work
- Conclusion

Ongoing Deployment

- Enterprise
  - 83+ Servers Completed
    - 1-4 Servers per node
    - Often memory limited
    - Predominantly x86, x86-64
    - Xen, ESX, Hyper-V, Linux Vserver
  - 400-700 Servers Remaining
    - 1-8x expansion (“easy” / “cheap” ➔ rearchitect)
- HPC
  - Also Looking at iDataplex
- Academic
  - K-8, Middle School
  - Using BladeCenters

Additional Issues

- **Funding**
- **Networking**
  - Insufficient IPv4 Address Space (also exploring IPv6)
  - Lack/Immaturity of IPv6 products
  - Starting VLAN work for isolating VMs
- **Standardization**
  - Heterogeneity in Hardware
    - x86 vs. x86-64 vs. Power
  - Heterogeneity in Hypervisors
    - Driven by application requirements
    - Ex: Oracle only certified on Oracle EL and Oracle VM
    - Ex: Microsoft only supports Windows on Hyper-V
    - Oracle VM, Hyper-V, Vmware ESX, Xen, Linux Vserver
- **Resources**
  - Leverage NC State model and code; IBM partnership
  - Utilize students (giving them experience)

Outline

- Educational Computing
- Pilot in Rural K-8 School
- Deployment
- Conclusion

Conclusion

- HPC, Academic, and Enterprise Cloud
- An “icon-ified” cloud computing laboratory
  - Delivering compute resources to rural K-12 schools
  - Logo Programming for 2nd and 3rd Graders
    - Deployed with initial integration in the computing curriculum
  - Storytelling Alice for 4th - 8th Graders
    - Deployed but not integrated with the curriculum
- “Native Desktop” Metaphor
  - Reduces cognitive load of “web browser” metaphor for younger students
  - Transition to standard web portal as they mature
- Production Deployment Ongoing
  - Welcome collaboration to solve issues

This work was generously supported by IBM.

Questions?

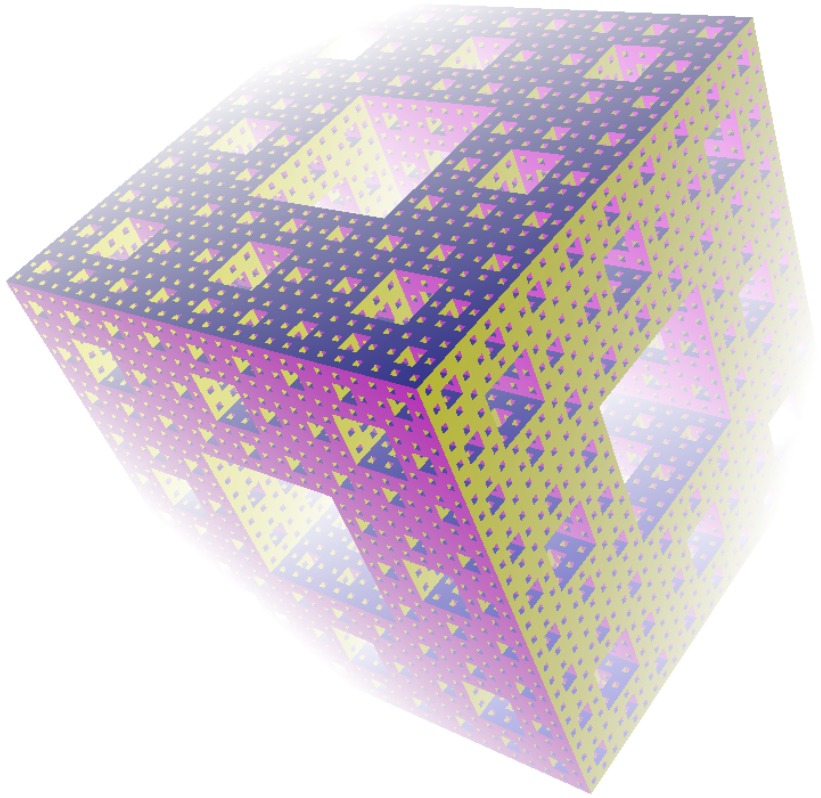

Supporting Transformative Research Through Regional Cyberinfrastructure (CI)

Gary Crane, SURA Director IT Initiatives

IBM Cloud Computing Seminar
North Carolina State
March 26, 2009

SURA Mission

SURA is a 501(c)3 university association with 63 member institutions whose mission is to:

- Foster excellence in scientific research
- Strengthen the scientific and technical capabilities of the nation and the Southeast
- Provide outstanding training opportunities for the next generation of scientists and engineers

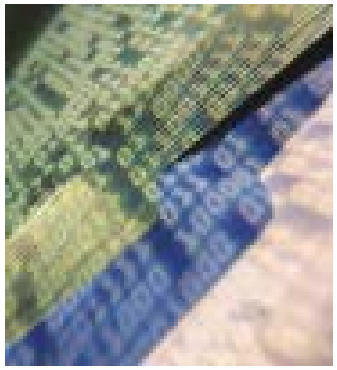

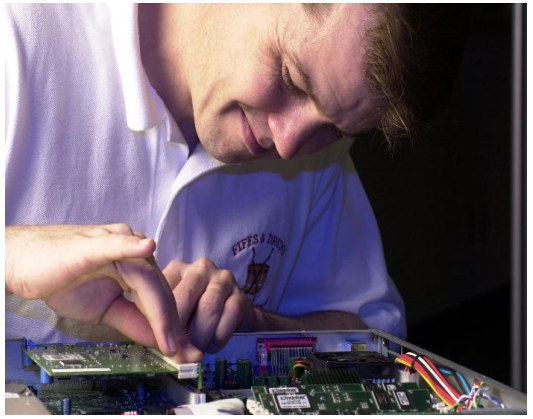

SURAgrid

www.SURAgrid.org

SURA Region

- **37%** of the US population
- **10** EPSCoR states
- **95%** of the nation's Historically Black Colleges and Universities (HBCUs)
- **22%** of the nation's Hispanic Serving Institutions (HSIs)

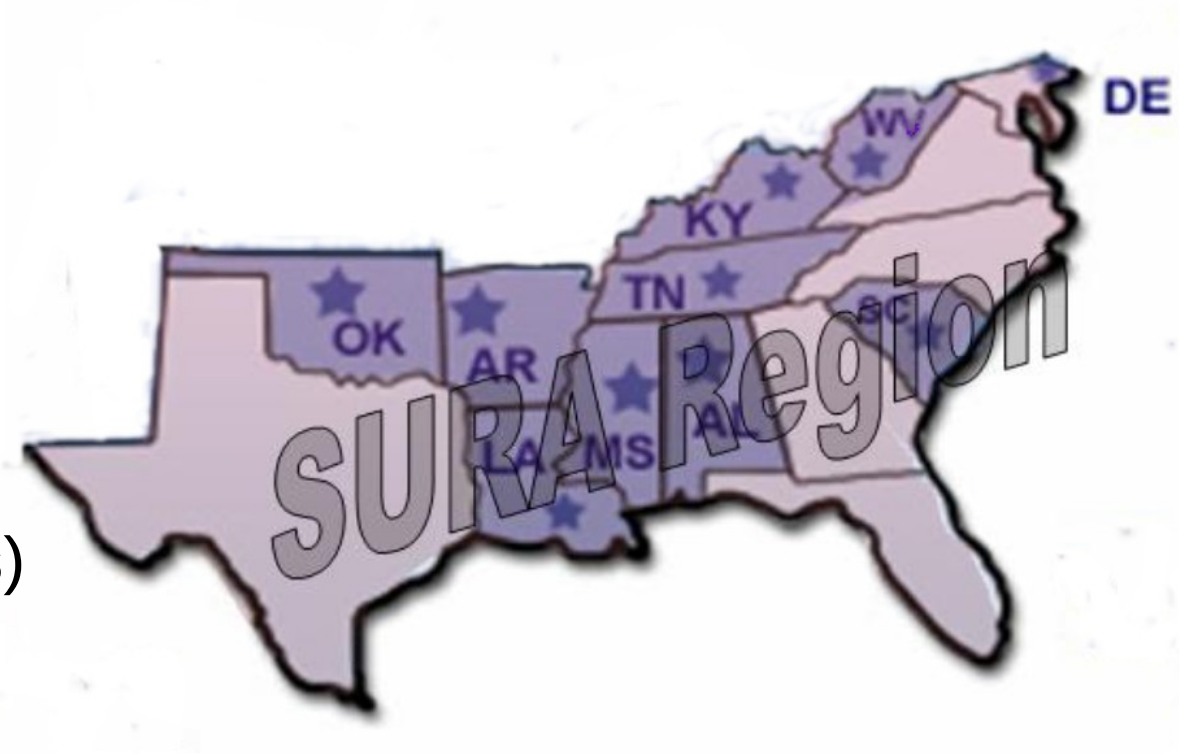

SURA Programs

Jefferson Lab – DOE Office of Science – to probe nucleus of atom and study quark structure of matter

SCOOP – DOD Office of Naval Research/NOAA – to provide IT "glue" to integrate coastal research components

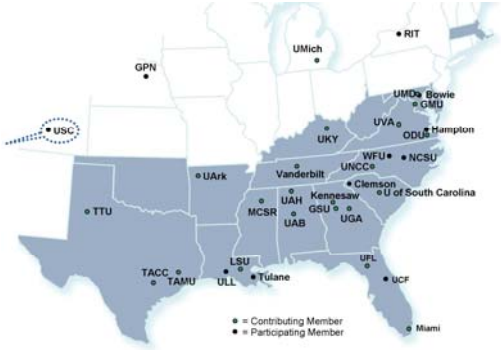

Information Technology - to build cyberinfrastructure foundation (the integration of high performance computing and networking) to support SURA's scientific and research programs

Relations – to formulate and sustain internal and external relations strategy and support for SURA's scientific and research programs

SURAgrid

SURAgrid Lowering Barriers for Deploying and Utilizing CyberInfrastructure

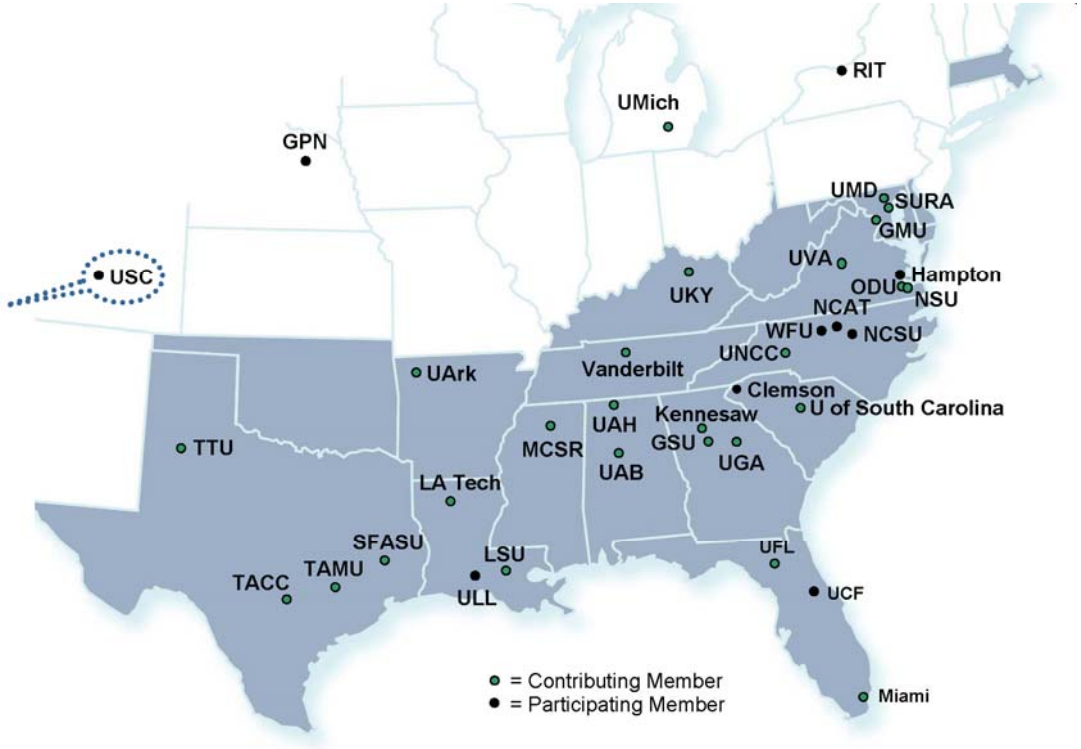

- 36 Institutions
- Shared accessible grid computing environment
- Enabling CI supported research & education
- On-Ramp to National CI
- Access to group negotiated discounted HPC systems

Members *(26 contributing, 10 participating)*

Contributing Members

1 University of Alabama at Birmingham
2 University of Alabama in Huntsville
3 University of Arkansas
4 University of Florida
5 George Mason University
6 University of Georgia
7 Georgia State University
8 Kennesaw State University
9 University of Kentucky
10 Louisiana State University
11 Louisiana Tech University
12 University of Maryland
13 University of Miami
14 University of Michigan ATLAS Computing
15 Mississippi Center for SuperComputing
16 Norfolk State University
17 University of North Carolina, Charlotte
18 Old Dominion University
19 University of South Carolina
20 Southeastern Universities Research
21 Stephen F. Austin State University
22 Texas A&M University
23 Texas Advanced Computing Center
24 Texas Tech
25 Vanderbilt University
26 University of Virginia

Participating Members

1 University of Central Florida
2 Clemson University
3 Great Plains Network
4 Hampton University
5 University of Louisiana at Lafayette
6 North Carolina State University
7 Rochester Institute of Technology
8 University of Southern California
9 Wake Forest University
10 North Carolina A&T

Active Governance Structure

SURAgrid Vision

Promote excellence for research and education enterprises by fostering collaborative engagement in cyberinfrastructure across the SURA region.

SURAgrid Mission

SURAgrid provides a community for collaborative development and use of cyberinfrastructure services to support the research and education missions of our membership.

SURAgrid Goals

- Develop a research outreach program to identify new users and new uses for the evolving regional and national computational and collaborative cyberinfrastructure available to the SURA region.
- Plan, manage and support the SURAgrid infrastructure to provide a solid foundation for the evolution of SURA region research and education programs.
- Develop a communications strategy for SURAgrid.
- Develop a sustainability model for SURAgrid.
- Strengthen existing and develop new corporate and organizational partnerships focused on improving regional use of Cyberinfrastructure (CI) services.

Major Areas of Activity

- Application Discovery & Deployment
- Outreach & Community Building
- Access Management Services
- Grid-Building
- Corporate Partnerships
- Collaborative Funding Efforts

Community Cyberinfrastructure

	Institutions	Resources	CPUs	Peak TFlops	GBytes Memory	GBytes disk
September 2005	9	11	490	1.3	548	4,755
October 2006	14	18	910	3.1	950	8,020
November 2007	12	16	2,041	12.6	3,626	56,310
November 2008	**12**	**16**	**2,426**	**18.2**	**4,484**	**54,292**

Over 6TF of capacity are being added by new and expanded systems coming onto SURAgrid by early Summer ’09

SURAgrid Access Management Model

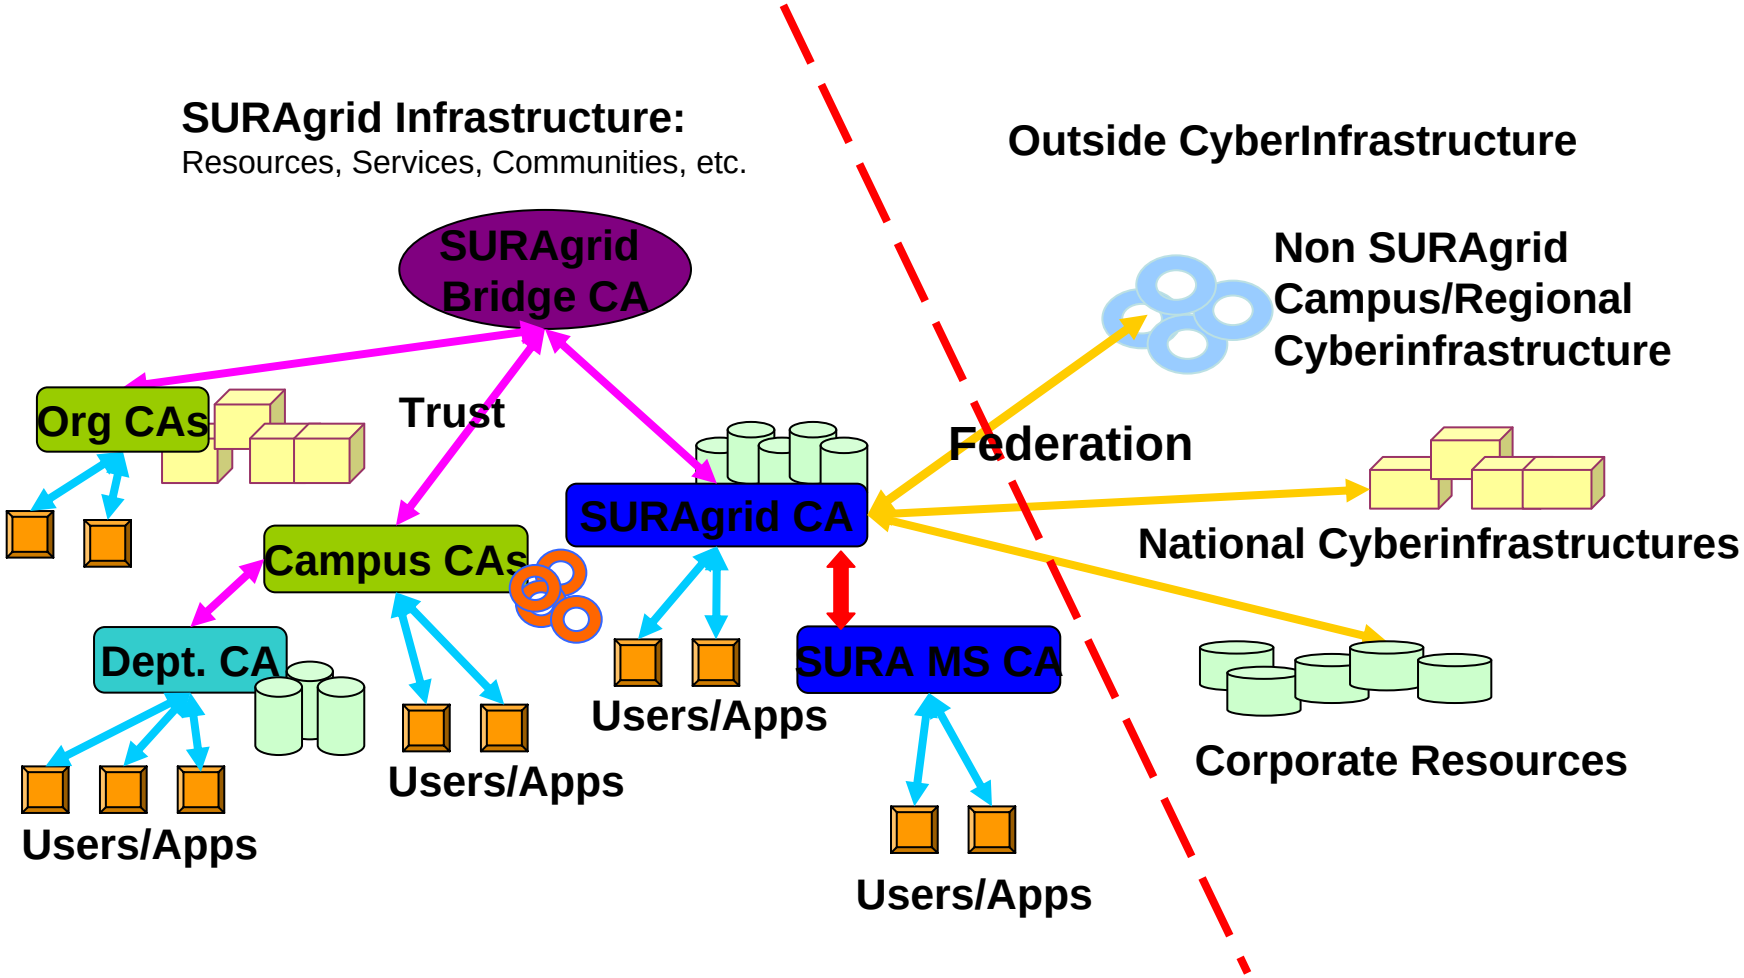

	Applications Deployed on SURAgrid	Current Status	Date Deployed
1	**ODU Turbulence Studies (Miami IBM P5)**	**Active**	**Nov-08**
2	**UIUC Geosciences Spatial Interpolation (Microsoft HPC Test-bed)**	**Active**	**Oct-08**
3	**ODU Options Pricing (tested Hybrid MPI-OpenMP on IBM P5s)**	**Ran Summer 2008**	**Mar-08**
4	**TTU GROMACS (GSU IBM P5)**	**Active**	**Feb-08**
5	**VCU Virtual Parasite (GSU IBM P5 & Microsoft HPC Test-bed)**	**Active**	**Sep-07**
6	**GSU Virtual Screening for Computational Chemistry (IBM P5)**	**Active**	**Sep-07**
7	**University of Delaware Climate Modeling with CAM3 (IBM P5)**	**Active**	**Jun-07**
8	UFL CH3D Storm Surge Monitoring System with Grid Appliance	On Hiatus	Sep-06
9	SURA SURAgrid Teaching Environment	Ran Spring/Summer 2007	Aug-06
10	ODU Bio-Sim: Bio-electric Simulator for Whole Body Tissue	On Hiatus	Aug-06
11	**LSU Wave Watch 3 for SCOOP**	**Ran during 2008 Hurricane Season**	**Mar-06**
12	NCSU Simulation-Optimization for Threat Management in Urban Water Systems	Ran 2006 - Migrated to TeraGrid	Mar-06
13	**UNC Storm Surge Modeling with ADCIRC**	**Ran during 2008 Hurricane Season**	**Jun-05**
14	UABgrid Dynamic BLAST	On Hiatus	May-05
15	GSU Multiple Genome Alignment on the Grid	On Hiatus	Dec-03

SURA Corporate Partnerships

- Significant product discounts
- Owned and operated by SURAgrid participants
- Integrated into SURAgrid with 20% of capacity available to SURAgrid pool
- IBM p575 – 1 and 2 TF configurations
- IBM e1350 Linux– 1 rack 3 TF and 2 rack 6 TF configurations
- Dell PowerEdge 1950– Single rack 2TF configuration
- Microsoft funded Windows HPC Server Pilot Program

Call for Collaboration

- SURAgrid offers the following benefits to collaborators:
  - Access to established SURAgrid community working groups in various areas of CI development and deployment to demonstrate a broader perspective and ability to extend the scope and impact of your proposed work to a larger, regional community;
  - Access to an existing pool of distributed high performance computing resources to show sufficient capability to support your proposed work. SURAgrid offers an effective test bed environment for CI development activities;
  - Assistance with identifying collaborators throughout the SURA region;
  - Outreach to new and underserved communities to broaden the reach and impact of your proposed work;
  - Direct assistance (team facilitation, call bridge, agency information) from SURA in facilitating proposal development and team communications for proposals that involve SURA or SURAgrid.

Q & A

Gary Crane, gcrane@sura.org

IBM University Relations: Opportunities for Partnership

Andy Rindos, Head, RTP CAS &
WW CAS Coordinator

The Larger Benefits of a Partnership with IBM

- **IBM University Relations programs provide many opportunities to partner with IBM development and IBM Research.**
- **These programs have sustained many NC State-IBM collaborations (educational programs, research projects, proofs-of-concept, etc.) around VCL and cloud computing (including IBM Blue Cloud solutions).**
- **In addition to providing faculty grants, student fellowships, IBM hardware and free software, some of these activities have also brought NC State and other universities in contact with the vast IBM customer and business partner ecosystem.**
- **And as evidenced by the recent WW announcement about the VCL Apache incubator project (October 2008), such collaborations can also provide universities access to IBM corporate PR machinery – with its WW reach.**
- **IBM hires thousands of new employees each year. – and is therefore eager to partner with universities in the development of curricula that provides graduates with highly marketable 21st century skills.**

Worldwide University Programs: moving towards a seamless virtual organization

Ctrs. for Advanced Studies | University Relations | Academic Initiative

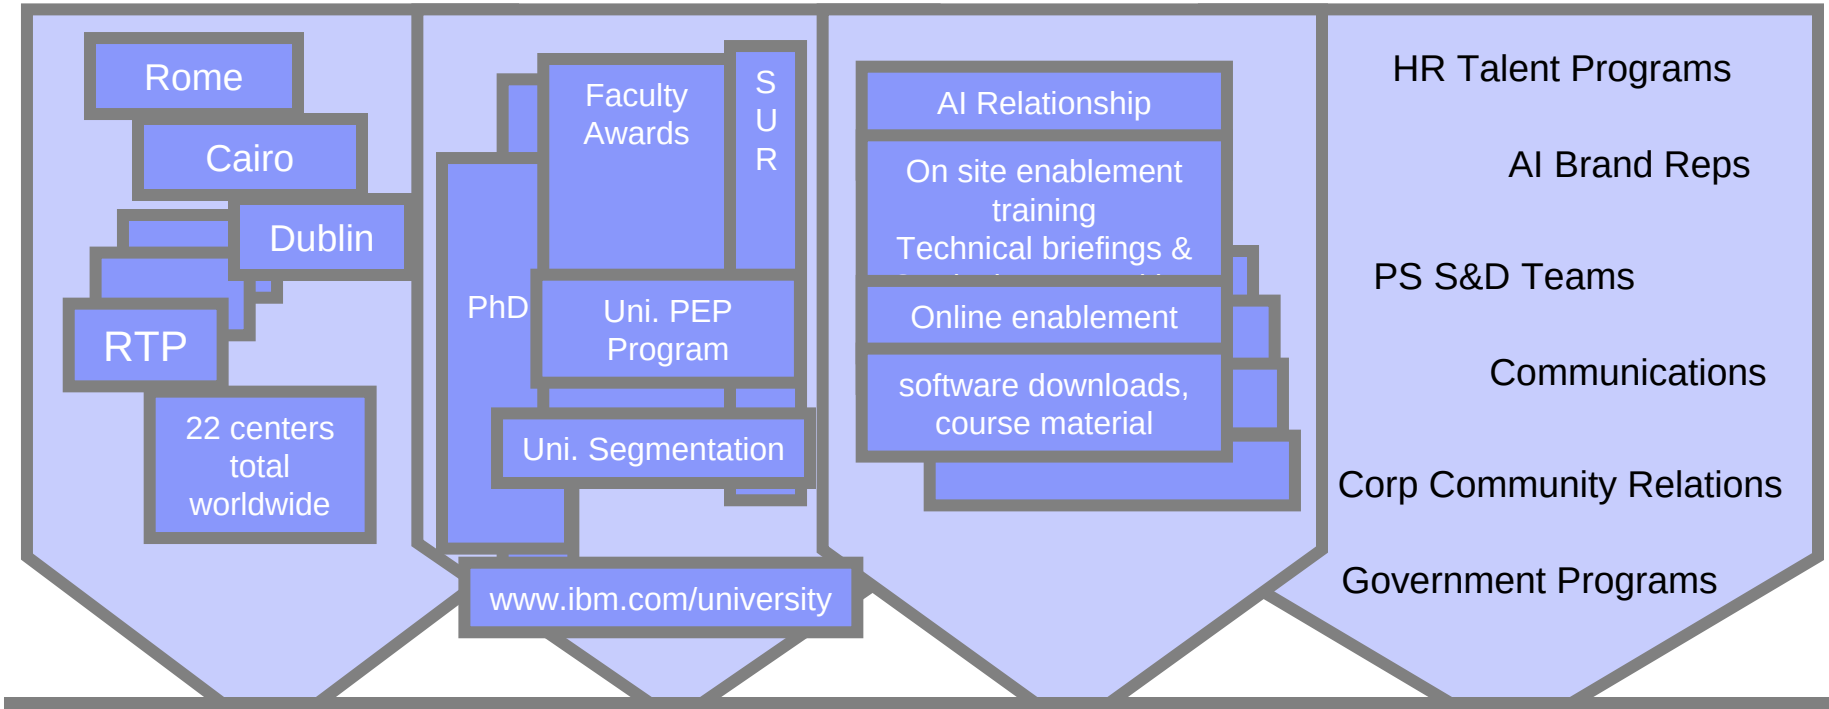

Worldwide University Programs

CAS is an official corporate University Relations program, representing the high-touch local UR mission, with an overall goal of integrating the 4 pillars of UR activities – research (a CAS specialization), recruiting, skills (AI).

Shared University Research (SUR)

- SUR Program Executive: Lilian Wu
- Typically 3 submission deadlines annually
- Grants written internally - require Partnership Exec, recruiting & account team (sales) support; defined Technical Rep
- Grant provides SUR $ budget to buy specified hardware at internal cost – now requires 2/3 matching

Faculty Awards

- Program Director: Jeff Brody
- Provides monetary grants in support of university research
- Up to $40K gift per faculty per year
- Numerous Faculty Awards Representatives that university can contact Several submission deadlines each year (2 for matching)
- Also Innovation Awards – targeted technologies

IBM Ph.D. Fellowship Program

- Tuition & fees + stipend
- Submissions due near year's end
- Limited to 2 new submissions per year per department (excluding renewals)
- Highly competitive

IBM Academic Initiative (AI)

- Director: Kevin Faughnan
- Relationship Mangers (by region) & AI Ambassadors
- IBM partnership with universities to assist in preparing students for today's high technology job market – 21st century skills.
- Promotion of open standards, open source, J2EE/Java/ Eclipse, etc. – and a special System z education program (with a substantial customer ecosystem)
- Provides consultation, hardware access
- All IBM software is free for use in the classroom or research

IBM Centers for Advanced Studies

23 centers WW

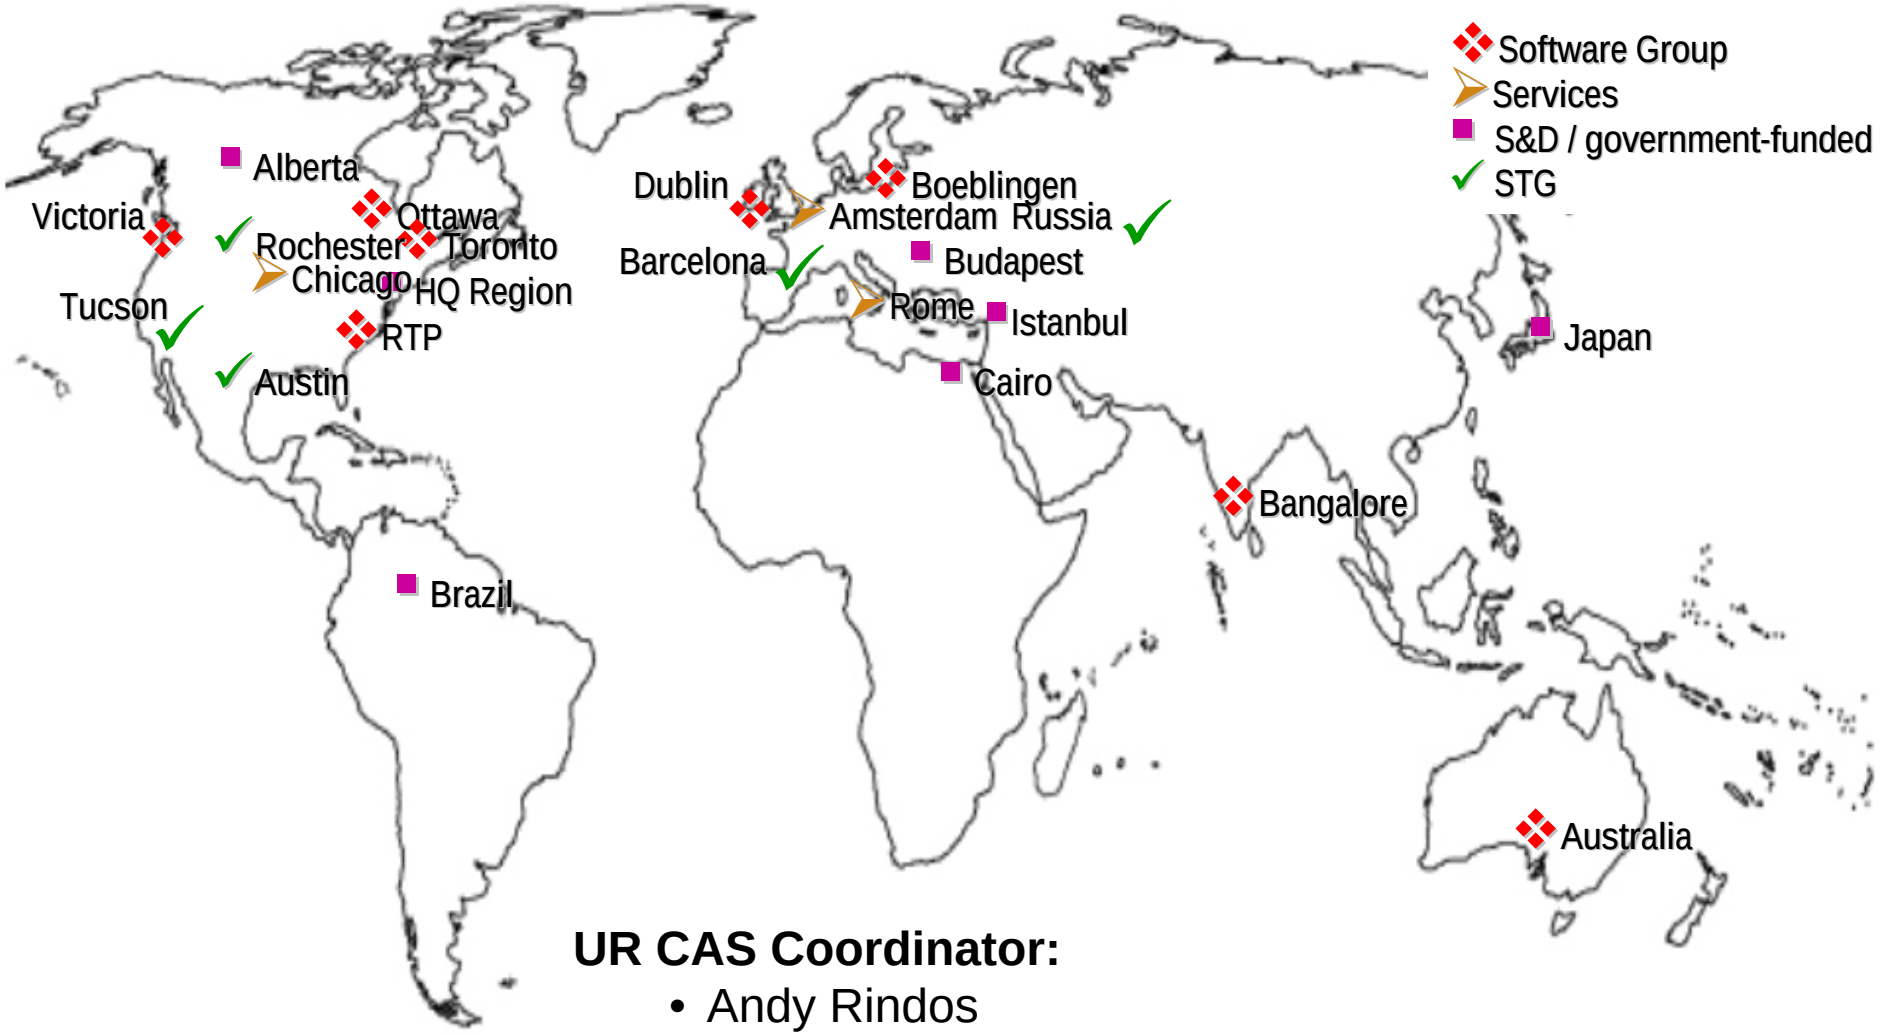

UR CAS Coordinator:

- Andy Rindos

The VCL Apache Incubator Project and The Virtual Computing Initiative (VCI)

- **The VCL code is freely available from the Apache website – and all are welcome to participate in its development through the Apache incubator project**
  - **http://incubator.apache.org/projects/vcl.html**
  - Code developers include Aaron Peeler, Josh Thompson, Andrew Kurth
  - The IBM WebSphere Technology Institute (WSTI) sponsored the VCL project at Apache (Matt Hogstrom)
  - Participation by your staff, faculty and students in creating a viable open source community around VCL will help the project graduate from “incubator” to “top-level” status
- **The VCL virtual appliance (integrating Apache web server + MySQL db + xCAT) can be freely downloaded at UNC-CH’s ibiblio (Brian Bouterse)**
  - **http://www.ibiblio.org/vclvm/**
- **The Virtual Computing Initiative (VCI) seeks to create an education community around VCL**
  - Sponsored by blade.org; **http://blade.org/vci.cfm**
  - Encourage universities that have established VCL pilots or production systems to share best practices, participate in the VCL Apache project, share images, etc.
  - Working to establish a VCL image repository (open source or free IBM Academic Initiative SW) at ibiblio

IBM University Days

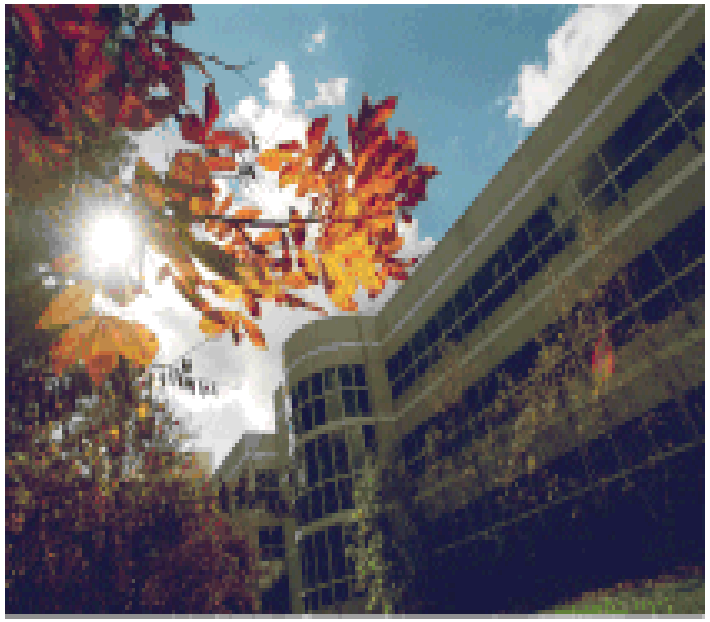

2nd International Conference on the Virtual Computing Initiative (VCI)

May 15-16, 2008

The IBM Employee and Activity Center (EAFC) ▼ Research Triangle Park, North Carolina

Over 300 university faculty, students, IT personnel attended from 25 universities

IBM Software Group (SWG) Proof-of-Concepts and VCL

- **IBM Tivoli Monitoring (ITM) agents have been integrated into VCL images for several years, together with IBM Tivoli Common Reporting (TCR)**
  - ITM has been providing VCL IT data for production purposes, as well as in support of several collaborative research projects
  - NC State Prof. Chris Healy received an IBM Faculty Award to improve the data presentation
- **Two important proof-of-concepts (POCs) for IBM SWG products will be launched shortly within the NC State VCL**
  - IBM **WebSphere Rainmaker**, a DataPower-based box, will manage and provision WebSphere/DB2 virtual clustered images into the VCL cloud (for use in NC State educational programs in SOA, etc.)..
  - The **Tivoli Image Management** solution will allow VCL users to more effectively navigate through the current NC State repository of over 600 VCL images.
- **Tivoli developers are working with the VCL team to integrate Tivoli Provisioning Manager (TPM) in various ways within VCL**
  - This includes the creation of **Tivoli Sservices Automation Manager (TSAM)** “sandbox” clustered images for research purposes (upgrading current **Request Driven Provisioning** clustered images at VCL).
  - **TPM for OS Deployment (TPMfOSD)** is being explored as a replacement or adjunct for xCAT in the VCL appliance.

IBM STG (hardware) and GTS (services) and VCL

- **Work is underway to integrate IBM Systems and Technology (STG) Blue Cloud Ensembles into VCL.**
  - VCL will then be able to provision IBM Blue Cloud ensembles on demand, and provide the energy, performance and availability management and guarantee capabilities that this technology will provide
  - Stay tuned for other forthcoming announcements
- **Through an STG-sponsored Faculty Award, NC State has developed a solution allowing VCL to reserve and access System z resources (from Marist College and other facilities) in support of local System z education programs**
  - Note that universities like NC State and NCCU who are providing substantial System z educational programs are given full access to our mainframe **customer ecosystem** – who then hire many of their students.
- **IBM Global Technology Services (GTS) VIA (Virtual Infrastructure Access) currently supports VCL as a customized offering – and may soon offer it as a fully supported solution (with IBM BladeCenter servers).**
- **IBM and NCSU are starting to work through the details to offer VCL as a service from our newest, greenest DC.**

NC State VCL-IBM Research Collaborations and Opportunities

1. **Image Repository Technology (Contacts: Glenn Ammons, Vas Bala, Giovanni Pacifici)**
   - Prototype image repository software developed by IBM Research (that will be at the core of the Tivoli Image Management solution) has been running on VCL servers for several months, as part of a joint research collaboration project.
     - NC State just received a $200K Open Collaborative Research (OCR) award from IBM Research for security-related research associated with this technology.
   - Educational image repository at ibiblio would use this. technology - a major showcase and real-world test
2. **Reservoir (Contacts: Yaron Wolfsthal)**
   - EU FP7-funded project involving multiple universities/vendors & IBM Research Haifa
   - Platform/front-end for integration of multiple cloud computing infrastructures/ varied architectures
   - Interest in bringing VCL “under” Reservoir, & bringing VCL to EU Reservoir universities
3. **Research Compute Cloud (Contacts: Mark VanderWiele et al.)**
   - Will provide RDP/TSAM “sandboxes” on demand (cloud within a cloud) to collaborate with university partners, i.e., an RCC testbed outside the IBM firewall
4. **Cloud Services Research (Contacts: Dan Dias et al.)**

Tivoli Software

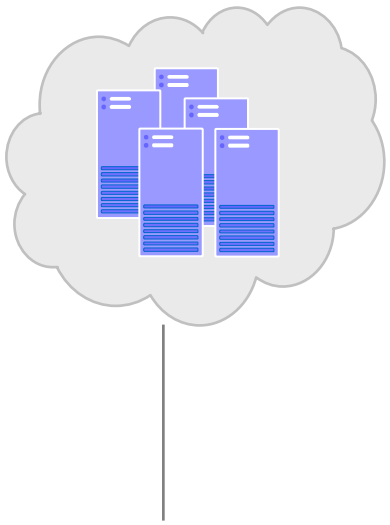

IBM's Perspective on Cloud Computing

NC State Cloud Computing Seminar

Pratik Gupta
STSM, Chief Virtualization Management Architect
IBM Software Group

pratikg@us.ibm.com

A Crisis of Complexity. The Need for Progress is Clear.

1.5x

Explosion of information driving 54% growth in storage shipments every year.

70¢ per $1

70% on average is spent on maintaining current IT infrastructures versus adding new capabilities.

85% idle

In distributed computing environments, up to 85% of computing capacity sits idle.

What is Cloud Computing?

A user experience and a business model

- Cloud computing is an emerging style of IT delivery in which applications, data, and IT resources are **rapidly provisioned** and provided as **standardized offerings** to users over the web in a **flexible pricing model.**

An infrastructure management and services delivery methodology

- Cloud computing is a way of **managing** large numbers of highly **virtualized resources** such that, from a management perspective, they resemble a single large resource. This can then be used to deliver services with **elastic scaling**.

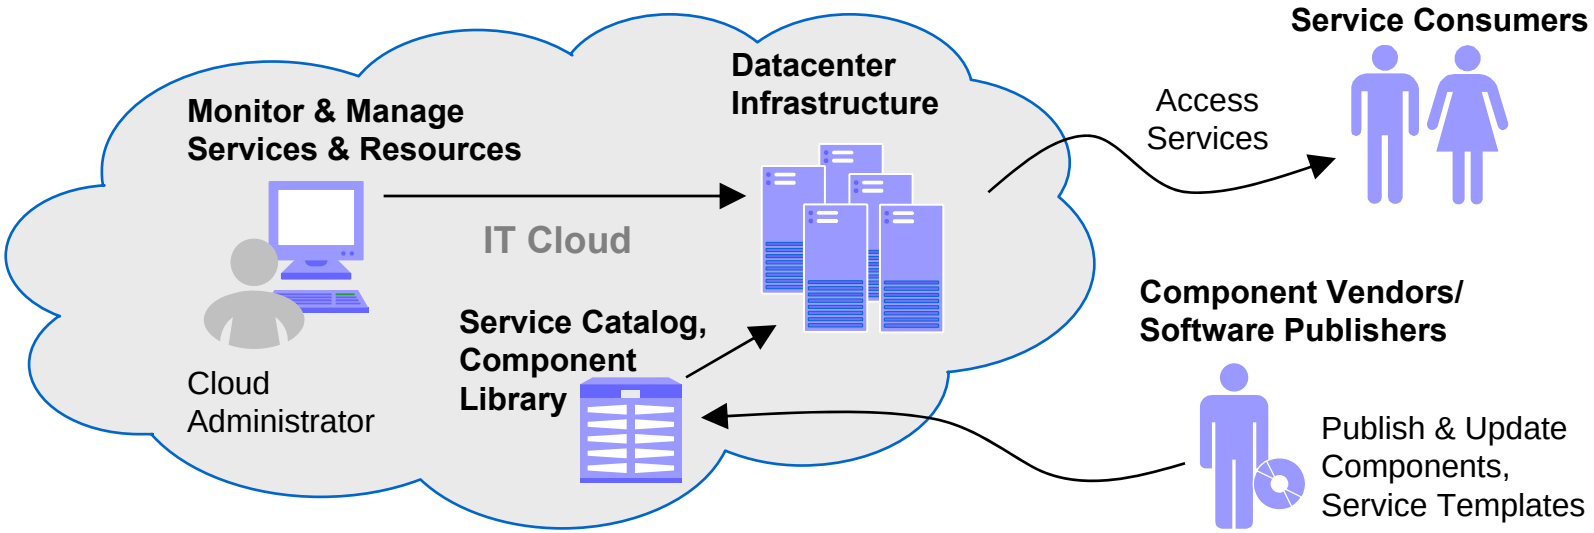

Cloud-onomics…

CLOUD COMPUTING

VIRTUALIZATION + ENERGY EFFICIENCY + STANDARDIZATION + AUTOMATION = Reduced Cost

….leverages virtualization, standardization and automation to free up operational budget for new investment

AGILITY + BUSINESS & IT ALIGNMENT + SERVICE FLEXIBILITY + INDUSTRY STANDARDS = OPTIMIZED BUSINESS

… allowing you to optimize new investments for direct business benefits

The future: Three co-existing delivery models

Over time, IT workloads will move to Cloud delivery models as applicable for the client.

Examples:

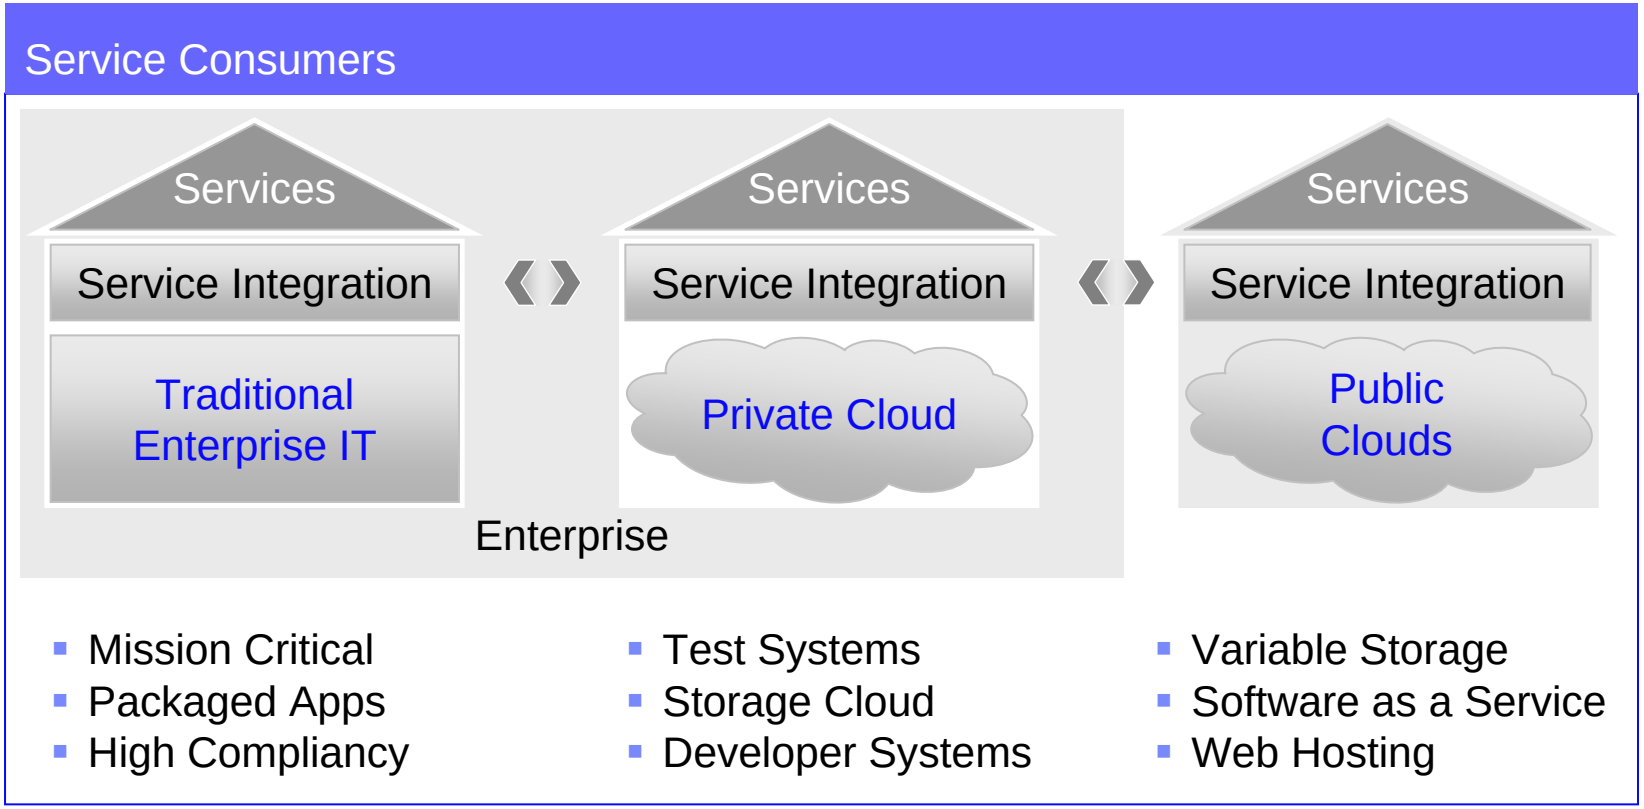

Cloud Computing deployments and services

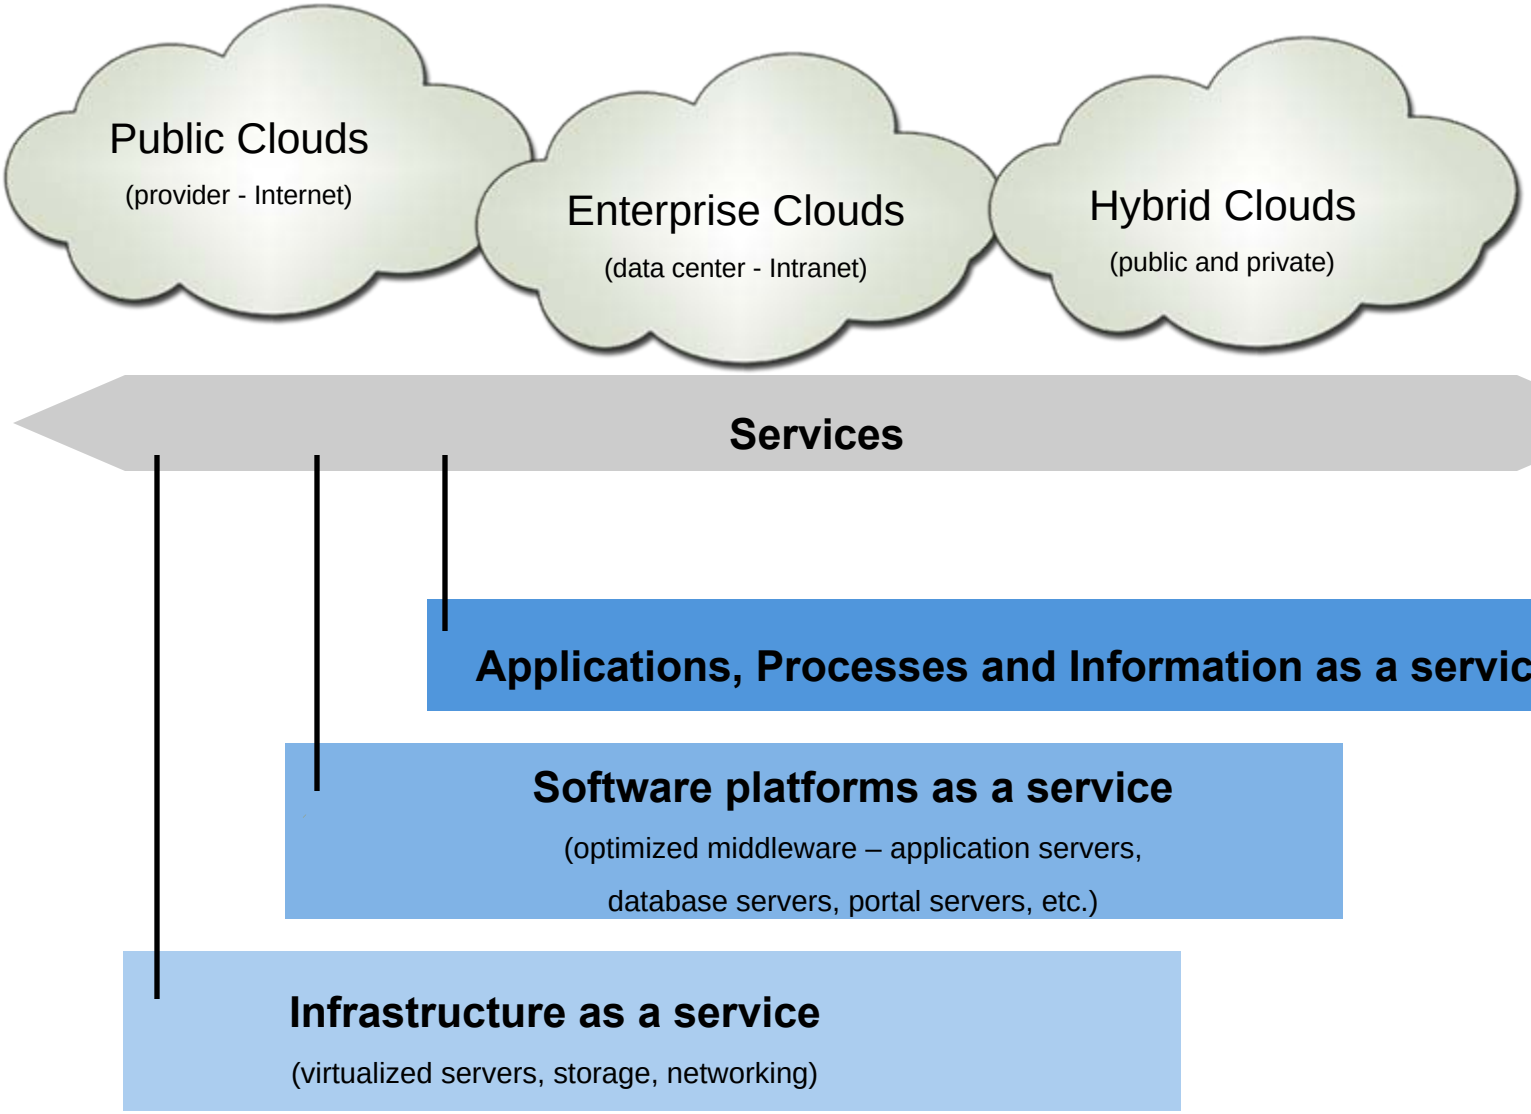

Architectural Model for Cloud Computing

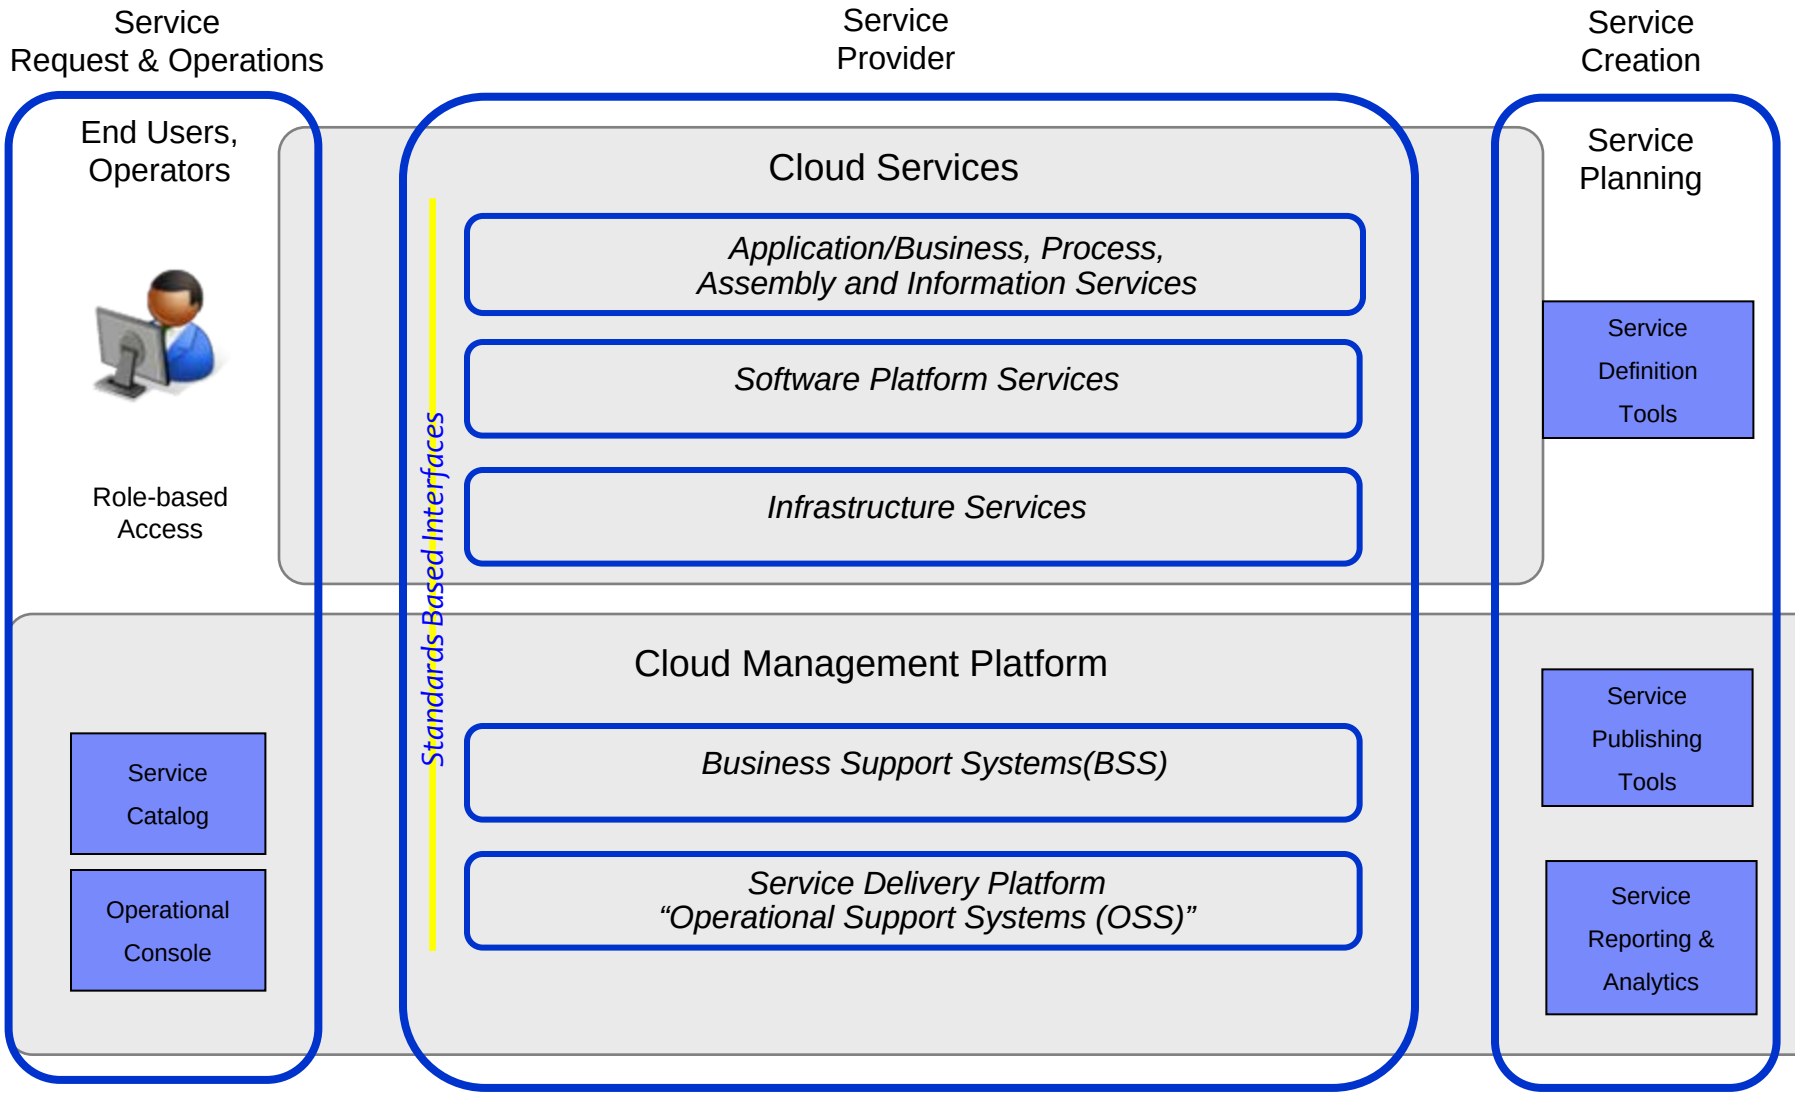

Workload Analysis

Workloads

- E-Mail, Collaboration
- Software Development
- Test and Pre-Production
- Data Intensive Processing
- Information Infrastructure

Service Management

- Service Catalog
- Request UI
- Operations UI
- Dynamic Scheduling
- Monitoring
- Capacity Planning SLA

Virtualization


Virtual Servers


Virtual Storage

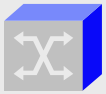
Virtual Networks


Virtual Applications & Middleware


Virtual Clients

Physical Layer


Non-IBM Servers


IBM System z Power Systems


System x, BladeCenter


IBM & Other Storage

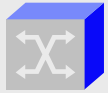
Networking

What workloads are we seeing move to Cloud delivery?

1. **Single virtual appliance workloads**
2. **Test and Pre-production systems**
3. **Mature packaged offerings, like e-mail and collaboration (see http://www.lotuslive.com)**
4. **Software development environments**
5. **Batch processing jobs with limited security requirements**
6. **Isolated workloads where latency between components is not an issue**
7. **Storage Solutions/Storage as a Service**
8. **Backup Solutions/Backup & Restore as a Service**
9. **Some data intensive workloads if the provider has a cloud storage offering tied to the cloud compute offering**

What workloads may not be ready for Cloud delivery today?

1. **Workloads which depend on sensitive data normally restricted to the Enterprise**
   - Employee Information - Most companies are not ready to move their LDAP server into a public cloud because of the sensitivity of the data
   - Health Care Records - May not be ready to move until the security of the cloud provider is well established
2. **Workloads composed of multiple, co-dependent services**
   - High throughput online transaction processing
3. **Workloads requiring a high level of auditability, accountability**
   - Workloads subject to Sarbanes-Oxley, for example
4. **Workloads based on 3rd party software which does not have a virtualization or cloud aware licensing strategy**
5. **Workloads requiring detailed chargeback or utilization measurement as required for capacity planning or departmental level billing**
6. **Workloads requiring customization (e.g. customized SaaS)**

Implementation

End Users

Service Portal

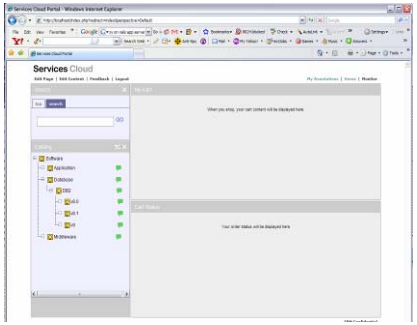

Service Request Catalog

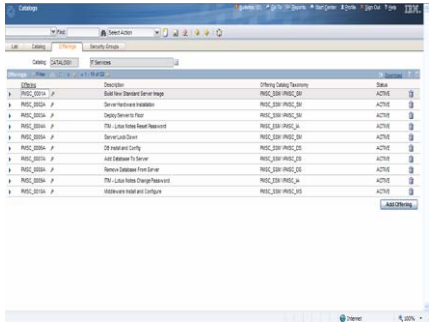

Provisioning Engine
- Workflows
- Expert Systems
- Scripts

Optional Service Modules
- e.g. Metering/ Usage Billing, Monitoring, etc.

Virtualized Cloud Infrastructure

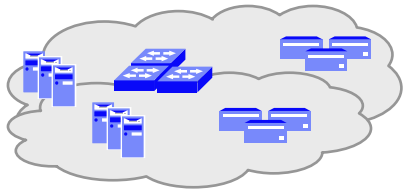

- Easy to access, easy to use Service Request Catalog
- Hides underlying complex infrastructure from user and shifts focus to services provided
- Enables the ability to provide standardized and lower cost services
- Facilitates a granular level of services metering and billing
- Workload standardization eases complexity

Architectural Model for Cloud Computing (OSS)

Service Request & Operations

IT Infrastructure & Application Provider

Service Creation & Deployment

End User Requests & Operators

. . .

Service Catalog Request UI
Operational UI

Datacenter Infrastructure

Service Management

User Request Management/Self Service Portal

Service Lifecycle Management

Service Catalog

Image Lifecycle Management | Provisioning | Performance Management | Availability/Backup/ Restore

Security: Identity, Access, Integrity, Isolation, Audit & Compliance | Usage Accounting | License Management

Cloud Administrator

Service Oriented Architecture | ***Information Architecture***

Optimized Middleware
(image deployment, integrated security, workload mgmt., high-availability)

Service Management

Standards Based Interfaces

Virtualized Infrastructure

Virtual Resources & Aggregations

Server Virt. | Storage Virt. | Network Virt.

System Resources

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

Virtual Image Management

Design & Build

Image Library (Store)

Deployment

Operational Lifecycle of Images

Service Management: Progressive Adoption of Capability

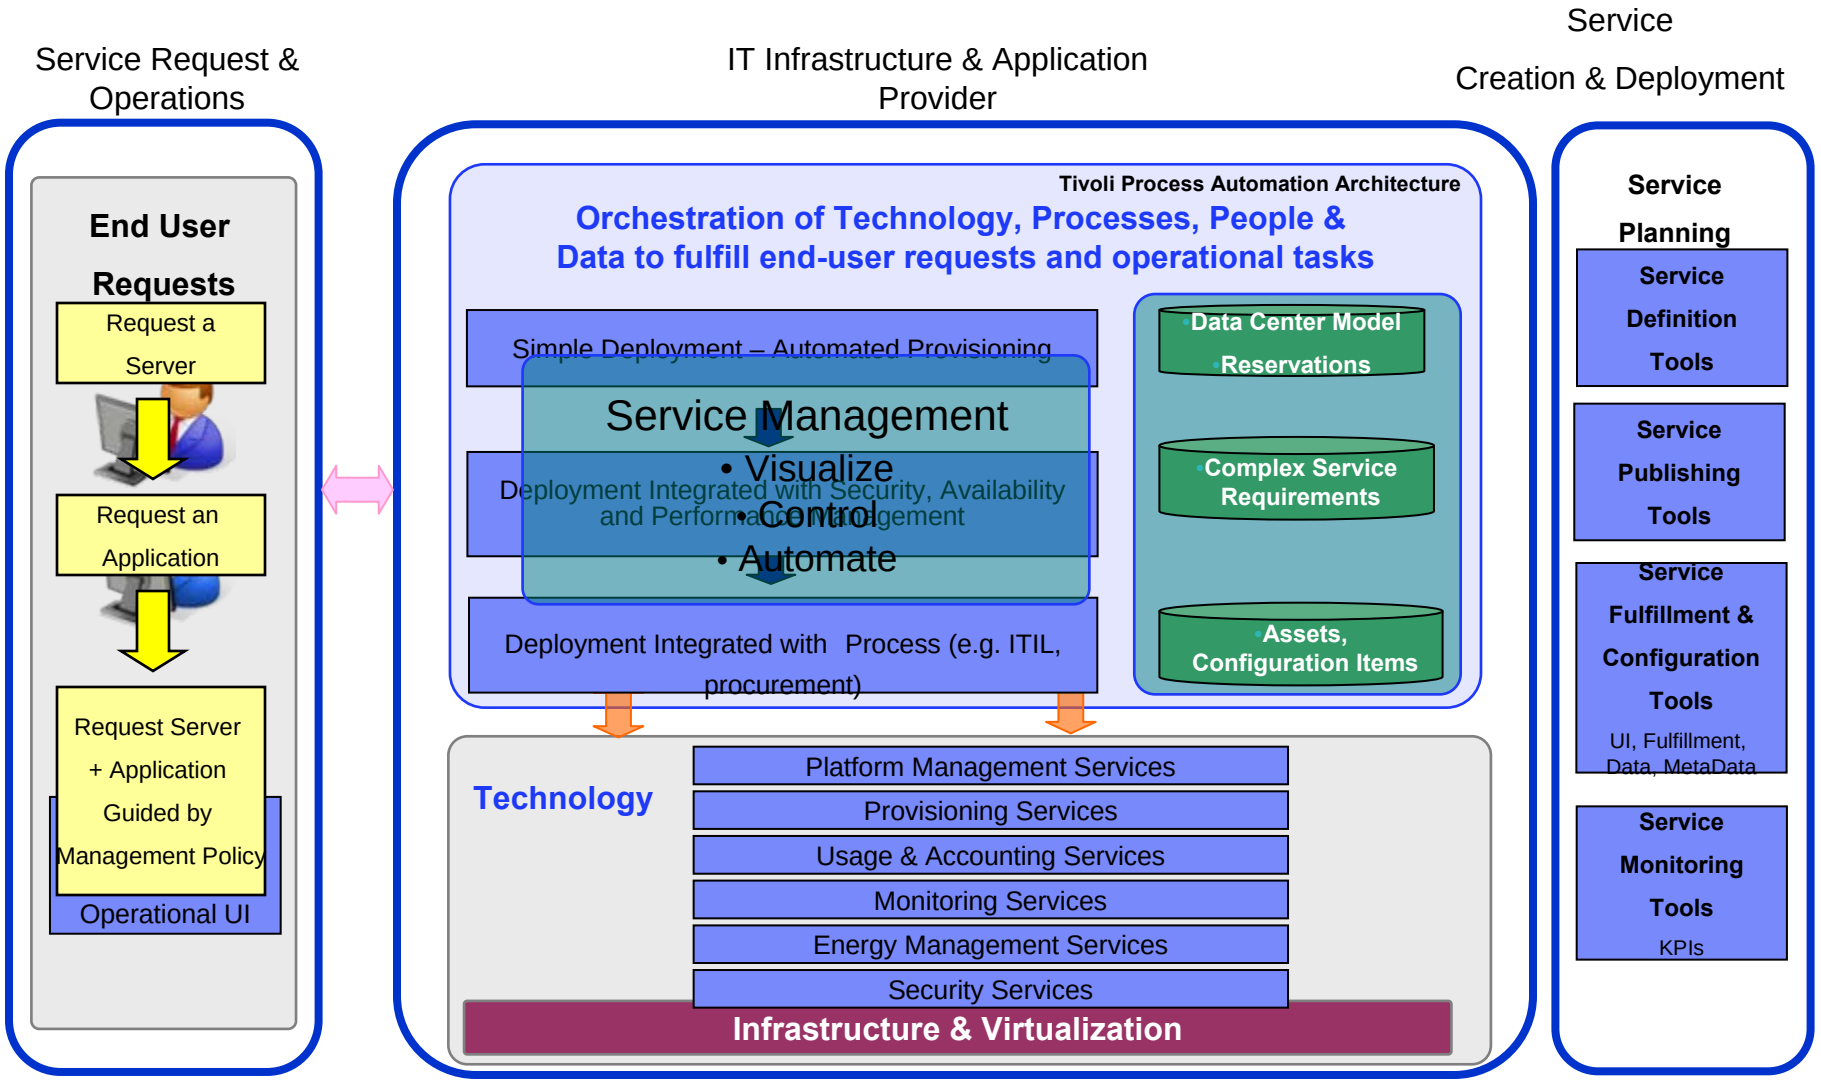

Lifecycle of a Cloud Service

Service Subscription

Subscriber

Administrator

Autonomic Execution of Management Plans

Service Offering

Service Catalog

Offering

Subscription & Instantiation

Production

Service Transition

Service Operations

Definition

Service Design

Termination

Cloud Service

Service Instance Termination

Subscriber

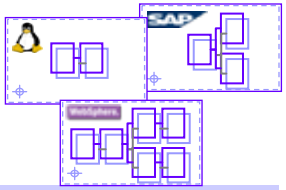

Service Definition

Service Creation

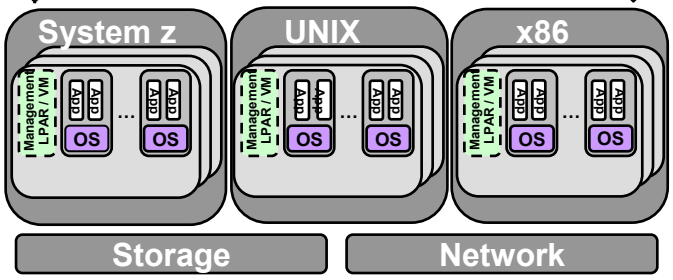

Tivoli Service Automation Manager (TSAM)

- Is an integrated service delivery platform
- Enables "as-a-Service" models:
  - Delivery and Management of Infrastructure Services
  - Delivery and Management of Platform Services
- Accelerates adoption of foundational capabilities for deploying & managing Cloud Services
- Enables dynamic instantiation and management of Cloud Services along their entire lifecycle
- Facilitates automation based on build & management plans including humans and management components

- **Request-driven provisioning and scheduling**
- **Application On-boarding through automation of middleware deployments**
- **Complete Lifecycle service management**

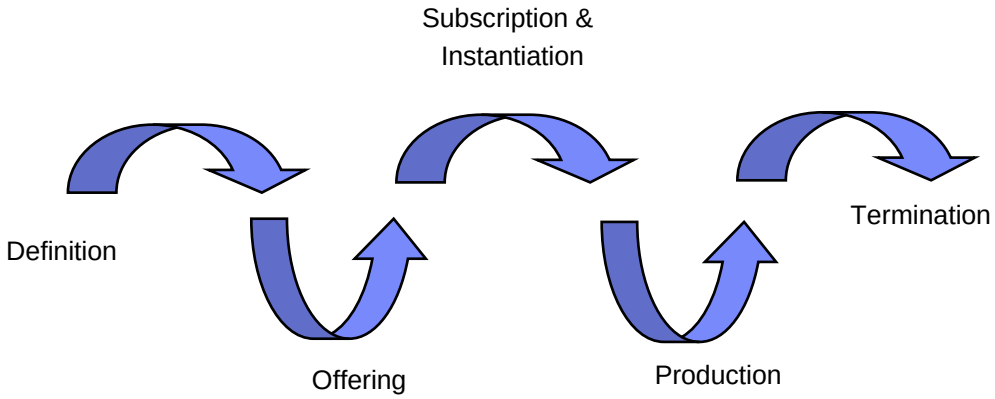

Request Driven Provisioning

Challenge: Self-service access to basic computing resources

- Self-Service GUI allows end users to request or reserve IT Resources, and optionally automatically fulfill that request for development or production environments.
- Simplified and repeatable deployment of Virtualized Operating systems and software stacks
- Automated configuration of systems and test tools
- Deployment of Service Management for various parallel projects (for example, monitoring)

Integrated Service Management for cloud

Challenge: How do I manage and control the cloud deployment throughout the service lifecycle?

- Appropriately manage growth, outages, changes, and other lifecycle aspects of the cloud
- Automate actions on the cloud based on monitoring metrics and threshold measurements
- Control the process workflow to be documented, approvals recorded, and measured by KPIs for SLA adherence
- Integrate configuration management, incident mgmt, change and release mgmt for your cloud

Application On-boarding

Challenge: Rapid provisioning of Middleware components to support a business application

- Eliminates the need to rely on manual deployment of middleware landscapes which are error-prone, slow, and labor expensive
- Best Practices and standard configurations are executed automatically without time intensive coordination of multiple internal teams
- Improve IT responsiveness to the business demands

Security in cloud computing

Authentication and role-based access control
- Federated Identity including single sign-on

Isolation Management
- Server, Storage and Network

Security for Image Management
- Security Metadata, Access Control, Authorization

Integrity management
- Virtual Image integrity

Risk and Compliance
- Auditing and Configuration Management
- Enterprise-level Regulatory Compliance

Policy Management

Threat Management

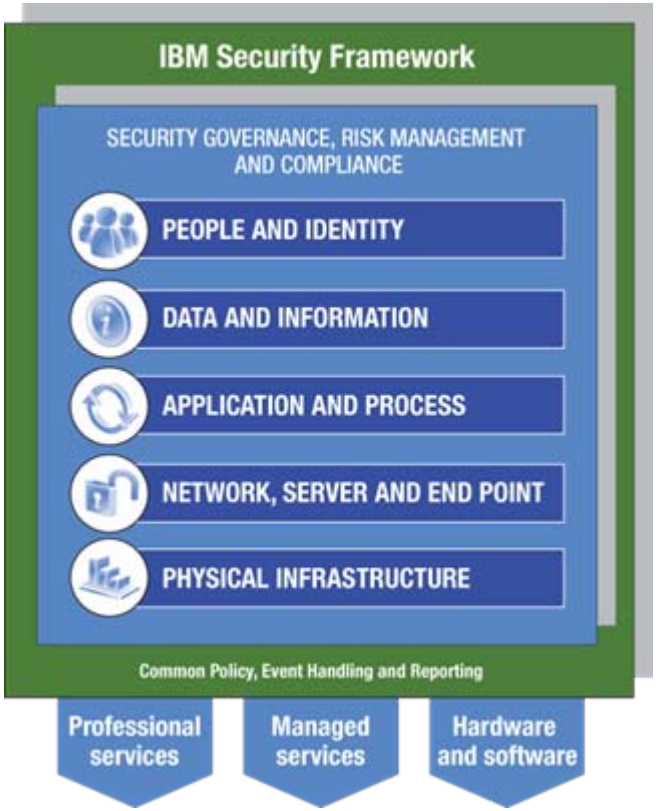

Image Lifecycle Management

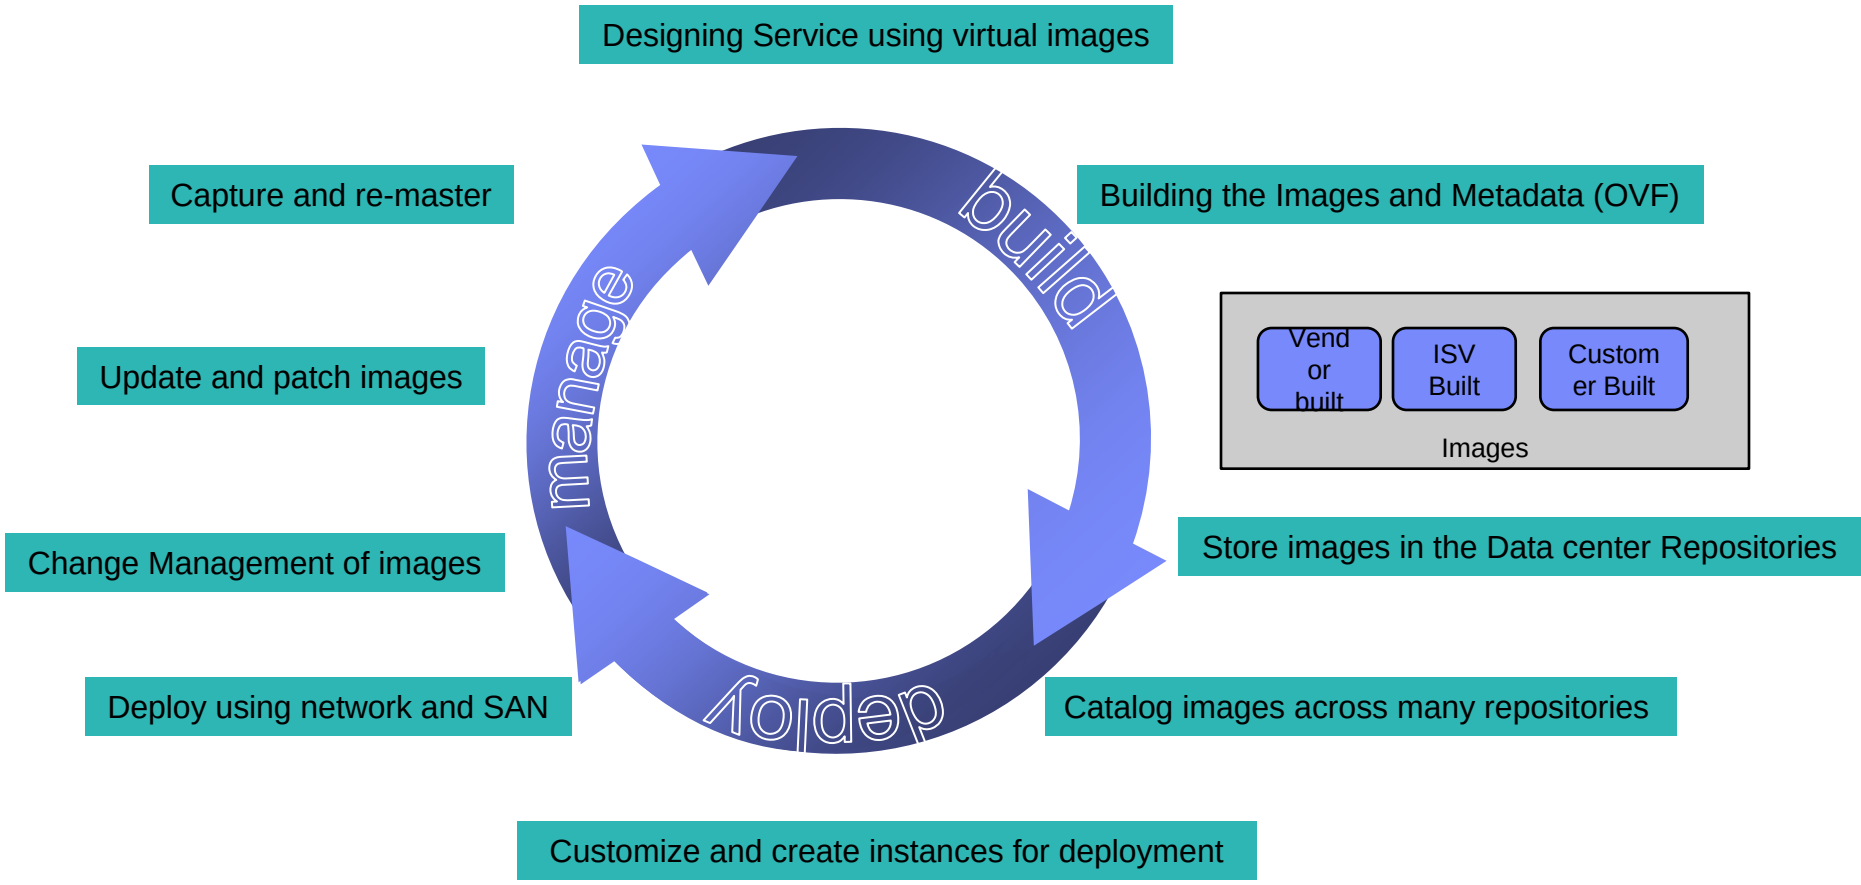

IBM Cloud Computing Services Offerings

A portfolio of leadership products and services for optimizing with cloud computing that continues to grow to support customers with cloud building and cloud delivered offerings.

Cloud Consulting

- Infrastructure Consulting Services for Cloud Computing
- Business Cloud Consulting Services
- Security and Resiliency Consulting Services for Cloud
- Resiliency Certification for Cloud Computing

Cloud Implementation

- Service Management for Cloud Computing
- Test and Developer Cloud Services
- Managed Security Services for Cloud Computing
- End User Cloud Services
- Scale out File Services

Cloud Delivered

- LotusLive
- Computing on Demand
- Information Protection Services
- Managed Data Protection for desktops and laptops
- DeveloperWorks on Amazon EC2

IBM Cloud Labs Mission & Customers

Drive IBM's leadership in cloud computing and act as core engine for all cloud activities

Getting started with Cloud Computing…

Develop a strategy ▶ *Best practices* ▶ *… think holistically*

Consolidate ▶ *Reduce from many to few* ▶ *… start with an inventory*

Virtualize ▶ *Assess and deploy* ▶ *… start now*

Manage ▶ *Gain and maintain control* ▶ *… modularity and standards are key*

In summary…

- **Cloud computing is a disruptive change to the way IT services are delivered**

- **Without a strategy, Cloud computing can be a threat to the CIO and IT team**
  - IT services delivered over the Internet
  - Perceived cost gap between a cloud service and traditional IT
  - “The next client/server”

- **With a strategy, Cloud computing is a huge opportunity for the CIO**
  - Lower cost of delivery for some workloads
  - More responsive IT
  - Ability to optimize delivery using traditional, private cloud, and public cloud
  - Greater visibility in billing / chargeback to LOBs
  - Greater range of available services, applications, and capabilities

Thank you!

For more information, please visit:
ibm.com/cloud

Or contact me at:
pratikg@us.ibm.com

WebSphere Clouds Strategy / Solutions

Bringing the appliance experience to WebSphere

Matt Hogstrom (hogstrom@us.ibm.com,
hogstrom@apache.org)

Agenda

- Why is IBM WebSphere looking to the clouds?

- Apache Incubator: Virtual Computing Lab

IT costs are increasing

- ***Costs to manage systems has doubled since 2000***
- ***Costs to power and cool systems has doubled since 2000***
- ***Devices accessing data over networks doubling every 2.5 years***
- ***Bandwidth consumed doubling every 1.5 years***
- ***Data Doubling every 18 months[1]***
- ***Server processing capacity doubling every 3 years[2]***
- ***10G Ethernet ports tripling over the next 5 years***

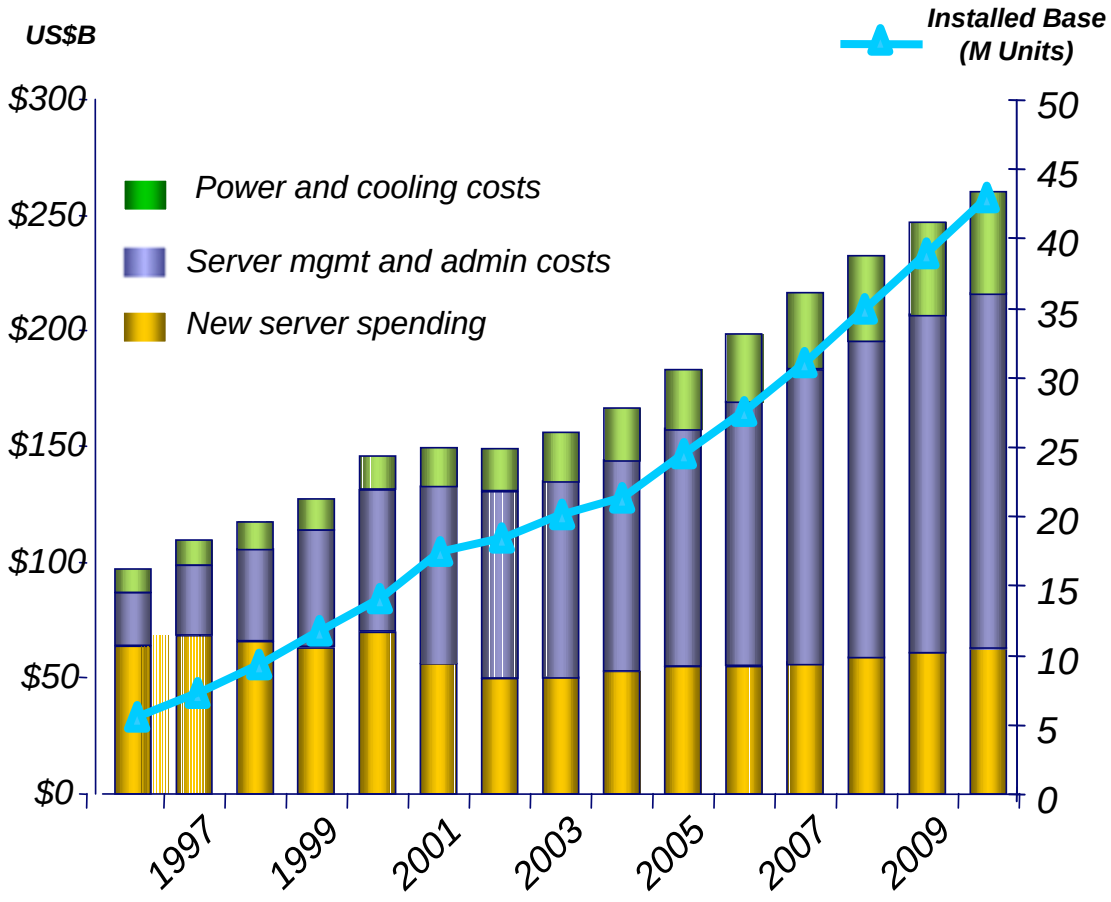

Source: IDC, 2008

[1]WW TB Capacity Shipped on Enterprise Disk Storage Systems

[2]Server processing consumption doubles every 3 years

Early Days

- IBM completed an internal PoC with the NC State VCL code modified to support WebSphere deployments
- Very much liked the model and it was extensible
- Examining issues related to productization, we needed to refactor the codebase

First things first … come up with a cool name

- An intangible but…
- Rainmaker is a cool name
  - SOA Deployer for IBM Software was not

Goals for Rainmaker

- Focus on topological deployments related to WebSphere types of use cases
- Provide a framework that allows for general purpose deployment of Virtual System Collections (VSC)
- Focus on using and shaping industry standards around Virtualization technologies and best practices
- Improve consumability of IBM products by leveraging simplification through virtualization
- Address system lifecycle by allowing for automated maintenance of deployed Virtual Systems

Customer requirement: make it work together

Create –

- *Install & config OS,*
- *Install &config WAS,*
- *Install Application,*
- *Install Patches, …*

Done!

Deploy

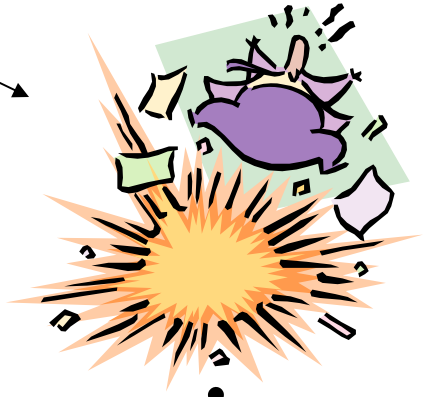

Start over…

Terminology

- Virtual Image
  - Collection of disks that can make up a single virtual machine
- Virtual Machine
  - Single OS instance running on virtualized hardware
- Pattern
  - Topological description of a collection of Virtual Machines
  - Includes configuration and possibly application artifacts
- Virtual System Collection
  - One or more Virtual Machines managed as a unit

Bringing the appliance experience to the WebSphere lifecycle

- Create
  - Design WebSphere topology patterns to match your application's requirements

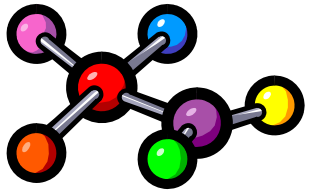

- Dispense
  - Create the virtual server instances to host the application
- Manage
  - Over the application's lifetime the application and the underlying infrastructure will need to be serviced
- Leverage server virtualization and application virtualization

Life Cycle Detail

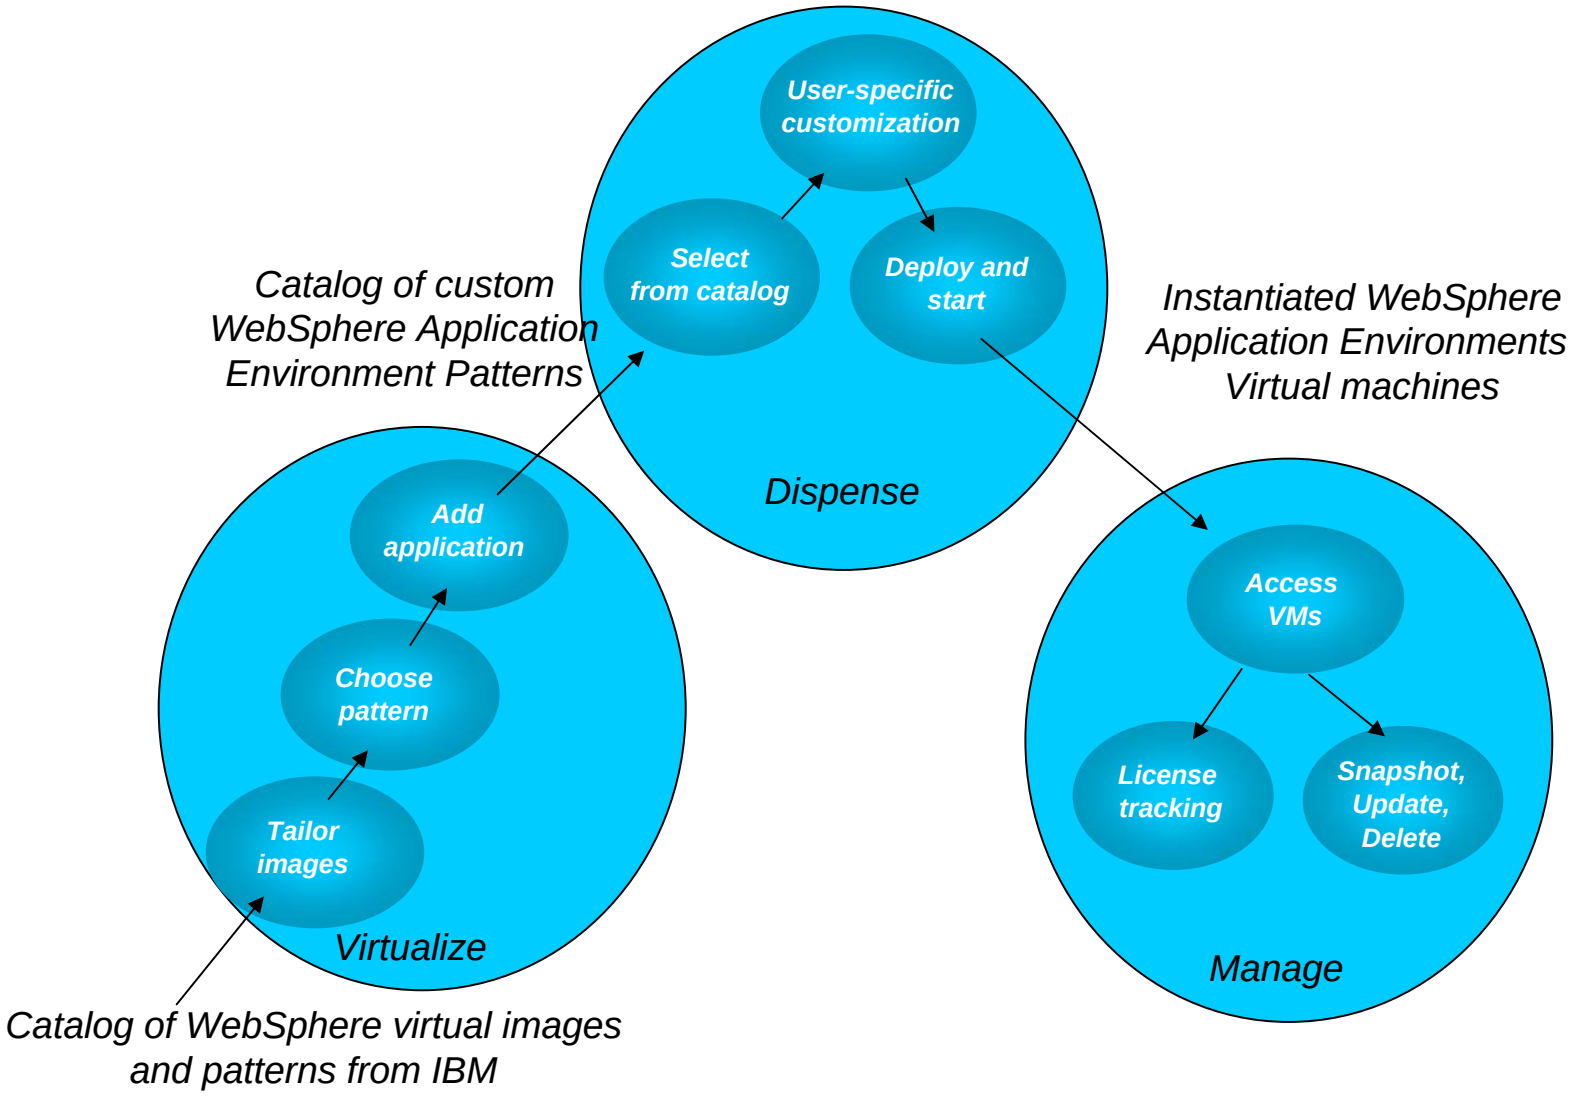

Technology trends – marrying appliances and virtualization

Appliances

DataPower

ISS Security

SAN Volume Controller

Virtualization

Easy to Use

Purposed

Fast Replication

Server Consolidation

Isolation

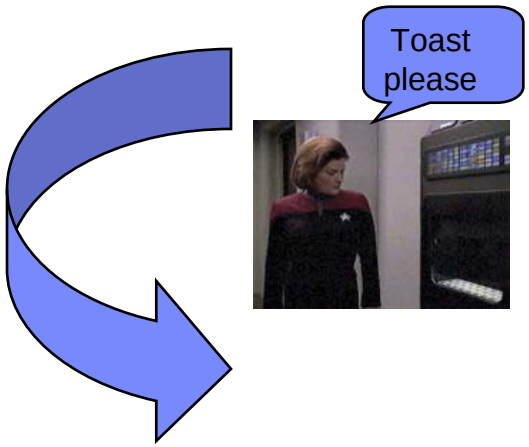

Virtual appliances

An emerging opportunity

VMware Virtual MarketPlace

Microsoft Run IT on VHD Program

WebSphere ND Image (basis for Rainmaking)

- **Base OS Installation**
- **Multiple Disk Design**
  - OS Disk: Linux
  - WebSphere Binary Disk
  - WebSphere IHS Disk
  - WebSphere Profiles Disk
- **WebSphere is pre-installed**
  - WAS Silent Install on separate disk
  - IHS Silent Install on separate disk
  - All 6 Profile types pre-created on separate disk
- **Activation engine provides user level customizations (all automatable)**
  - Operating System: Network and Passwords
  - WebSphere Configuration: choosing profile and constructing topology
- **Attended customization via YaST panels**
- **Open Virtual Format (OVF) 1.0 support**

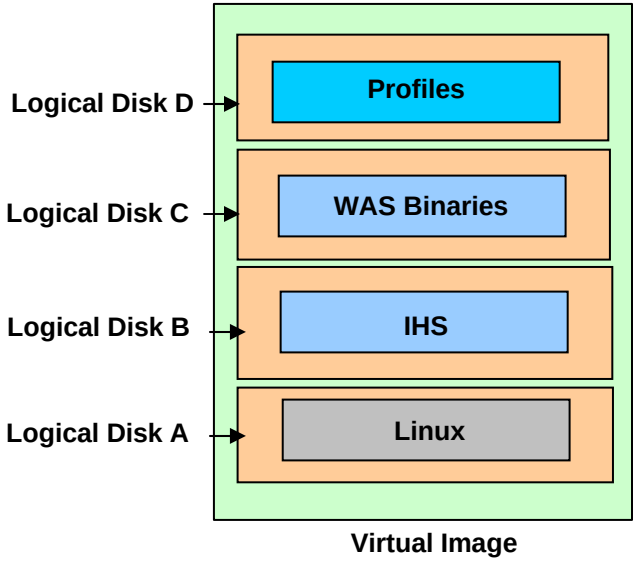

http://www.ibm.com/developerworks/downloads/ws/wasnd/learn.html

From image to deployed virtual systems

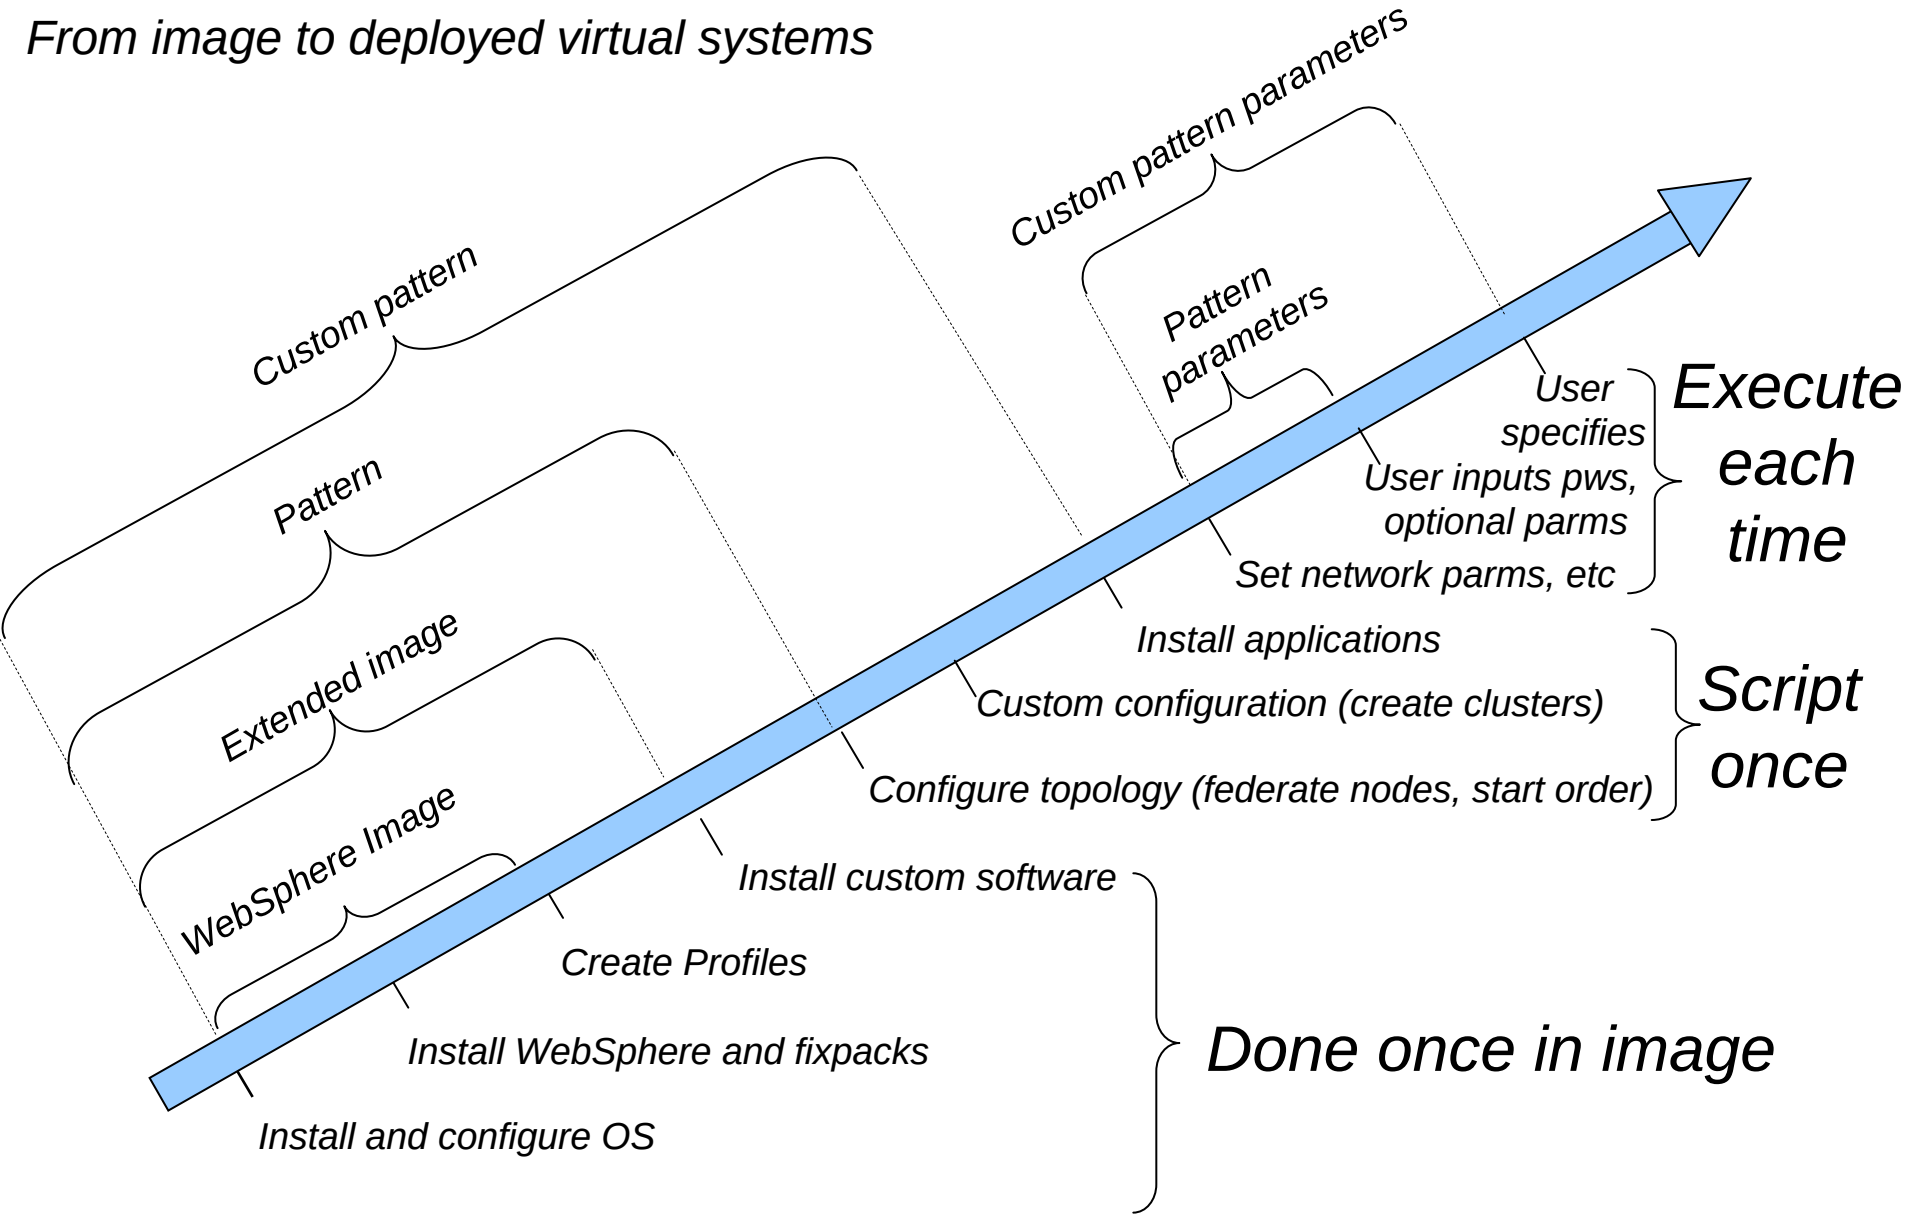

Patterns

- WebSphere is a collection of Virtual Machines that operate as a unit.
  - Each VM, or node, in a Collection has a specific purpose
    - Deployment Manager
    - Custom Node
    - Proxy Node

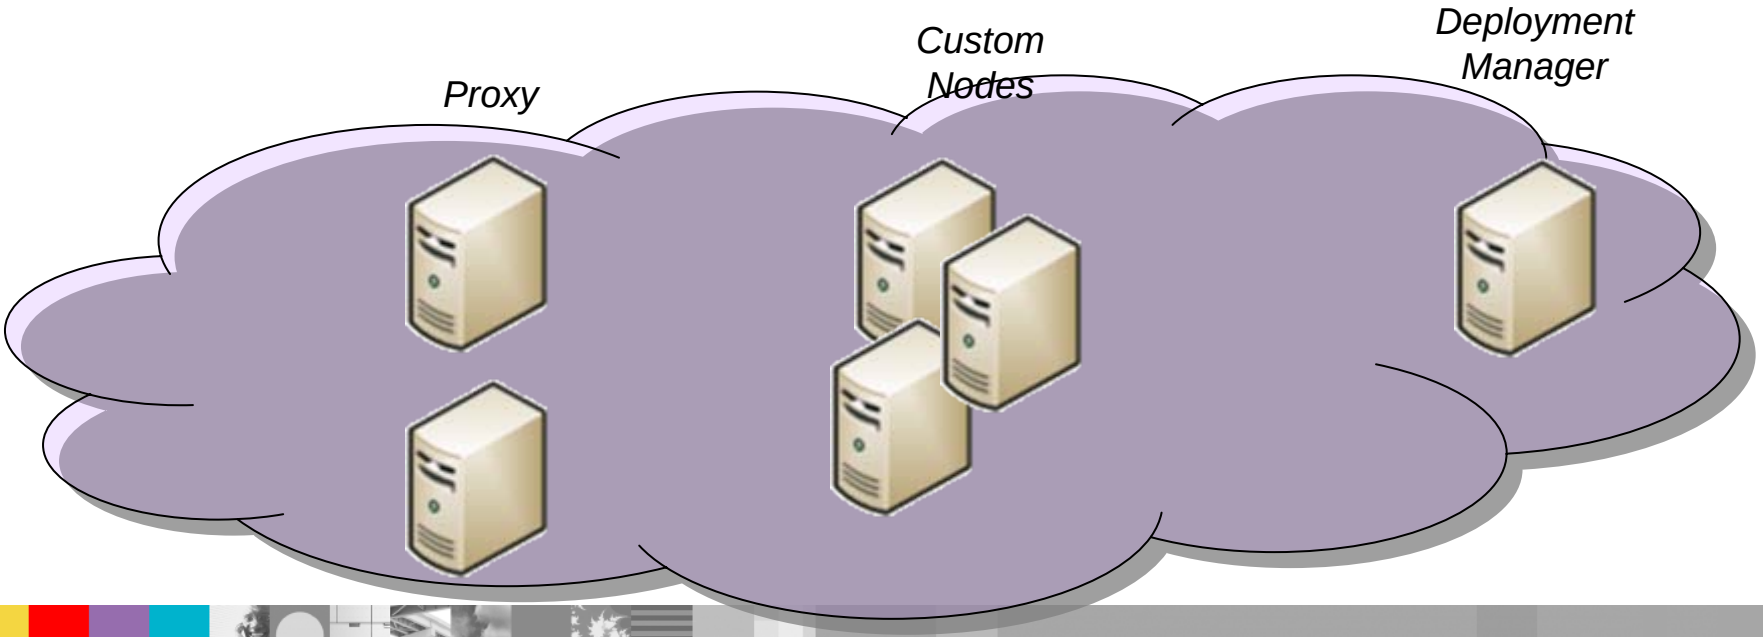

Rainmaker Flow

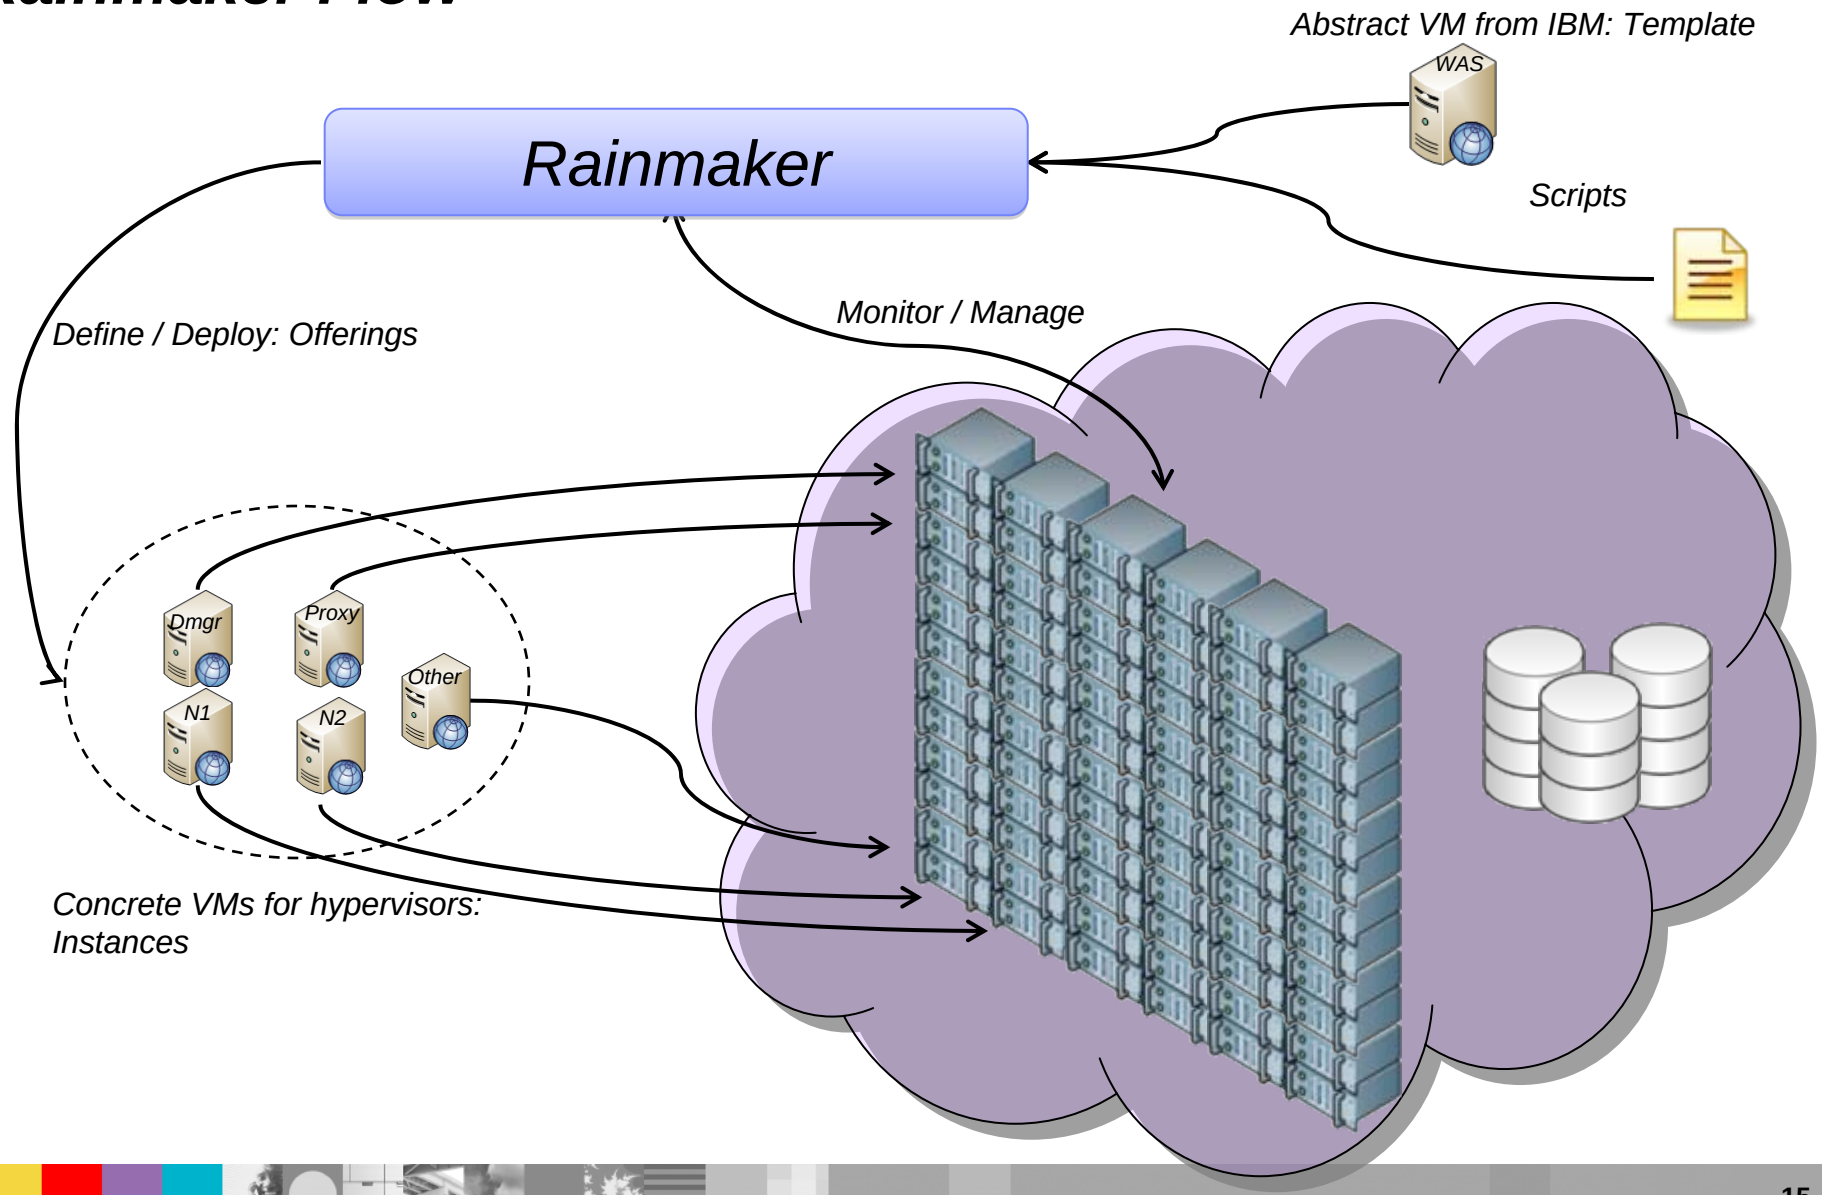

A cloud in a cloud

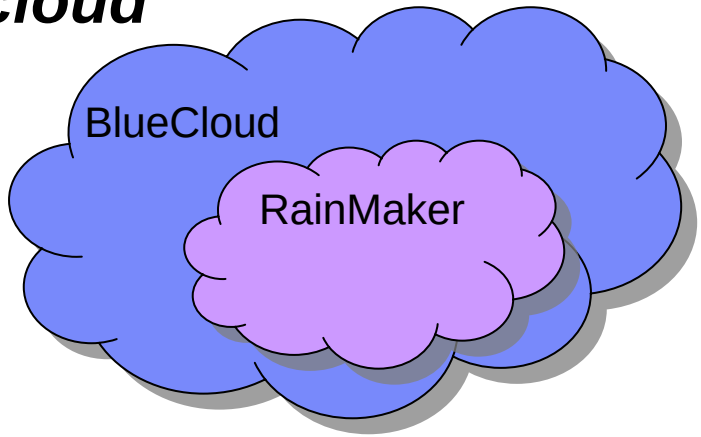

- WebSphere Clouds uses architectural concepts from IBM BlueCloud
  - WebSphere Ensemble Service Center
  - WebSphere Image Repository
  - WebSphere Middleware Ensemble
  - OVF metadata for template images and offerings
- Clouds in Clouds…
  - WebSphere Clouds are programmable and can be remote controlled using (REST) APIs, by IBM BlueCloud offerings including Tivoli Provisioning Manager.

IN THE WILD

Apache Incubator Virtual Computing Lab

- VCL donated to Apache VCL in November of 2008
- Initial committers from NC State, Duke, UNC, Virginia Tech
- Rules of the Road
  - People have opinions … opinions are not rules
  - Community oriented rather than technology oriented
  - Bottoms up … not top down
  - No permission necessary … simply get involved

- Mailing list vcl-dev@incubator.apache.org
  - Also vcl-user, vcl-commits
- Web Site http://cwiki.apache.org/confluence/display/VCL

Going Forward

- Looking to develop a community based site to have users and practitioners develop common practices for varying technologies
  - Operating systems
  - Databases
  - Application Stacks
- Validate Interoperability of diverse technologies in open source and development

IBM & NC State VCL Collaboration

Cloud Computing – Redefining IT Delivery

Dave Doria – IBM Systems and Technology Group -
Technical Strategy and Architecture Development

1 April 2009

Agenda

- **A Crisis of Complexity**
- **Addressing the Complexity**
  - The IBM Blue Cloud Initiative
  - Cloud Computing Defined
  - What is an IBM Ensemble
  - Ensemble built-in Capabilities
- **IBM / NC State “Cloud in a Box” Proof of Concept**
  - Integration of Server Ensembles into the VCL environment
  - Demonstrate improved IT management based on ensembles
  - Exploring Ensemble Advanced Optimizations

A Crisis of Complexity. The Need for Progress is Clear.

85% idle

In distributed computing environments, up to 85% of computing capacity sits idle.

70¢ per $1

70% on average is spent on maintaining current IT infrastructures versus adding new capabilities.

The IBM Blue Cloud™ Initiative

"Deliver ***Cloud Computing*** and ***IT Simplification*** to our customers, integrating the best of IBM's existing and future products **to *simplify the deployment and management of customer workloads*.**"

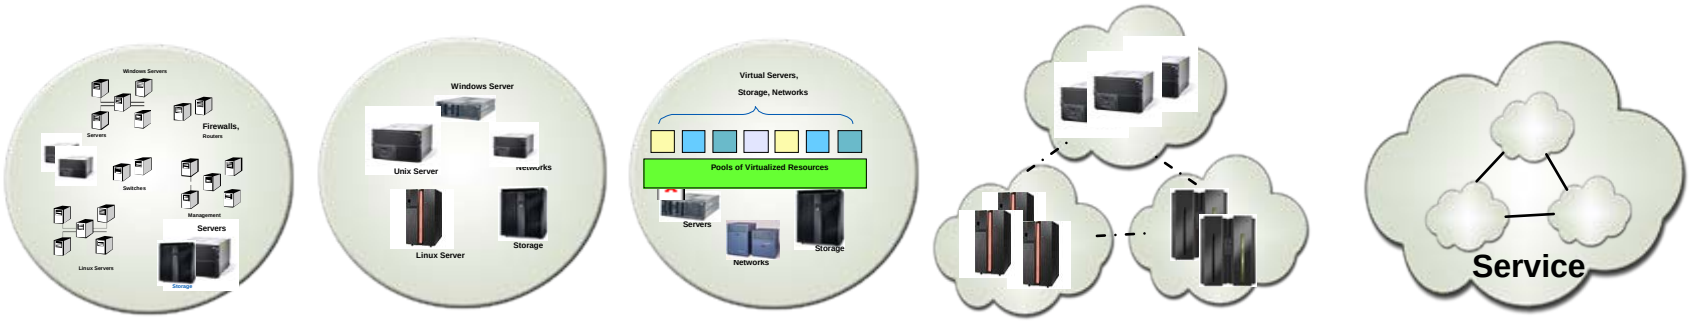

Complex → *Physical Consolidation* → *Virtualization* → *Ensemble* → *Cloud*

What is Cloud Computing?

A user experience and a business model

- Cloud computing is an emerging style of IT delivery in which applications, data, and IT resources are **rapidly provisioned** and provided as **standardized offerings** to users over the web.

An infrastructure management and services delivery methodology

- Cloud computing is a way of **managing** large numbers of highly **virtualized resources** such that, from a management perspective, they resemble a single large resource.

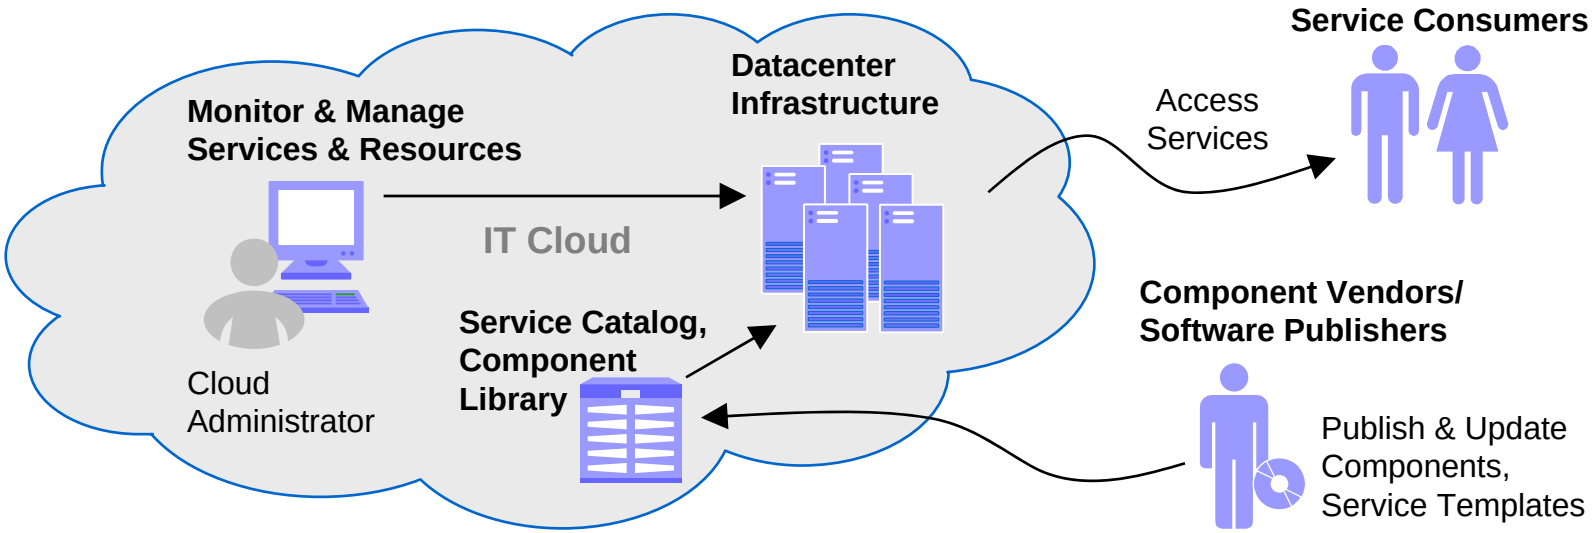

What is IBM Ensemble Technology

An ensemble is a pool of like systems that are managed as a single system

- Scale from a few to many thousands of virtual or physical nodes
- Reduce management complexity with integrated virtualization, management, and security software
- Allow workload optimization for maximum performance and efficiency

IT Infrastructure Evolution and the Expanding Uses Of Virtualization

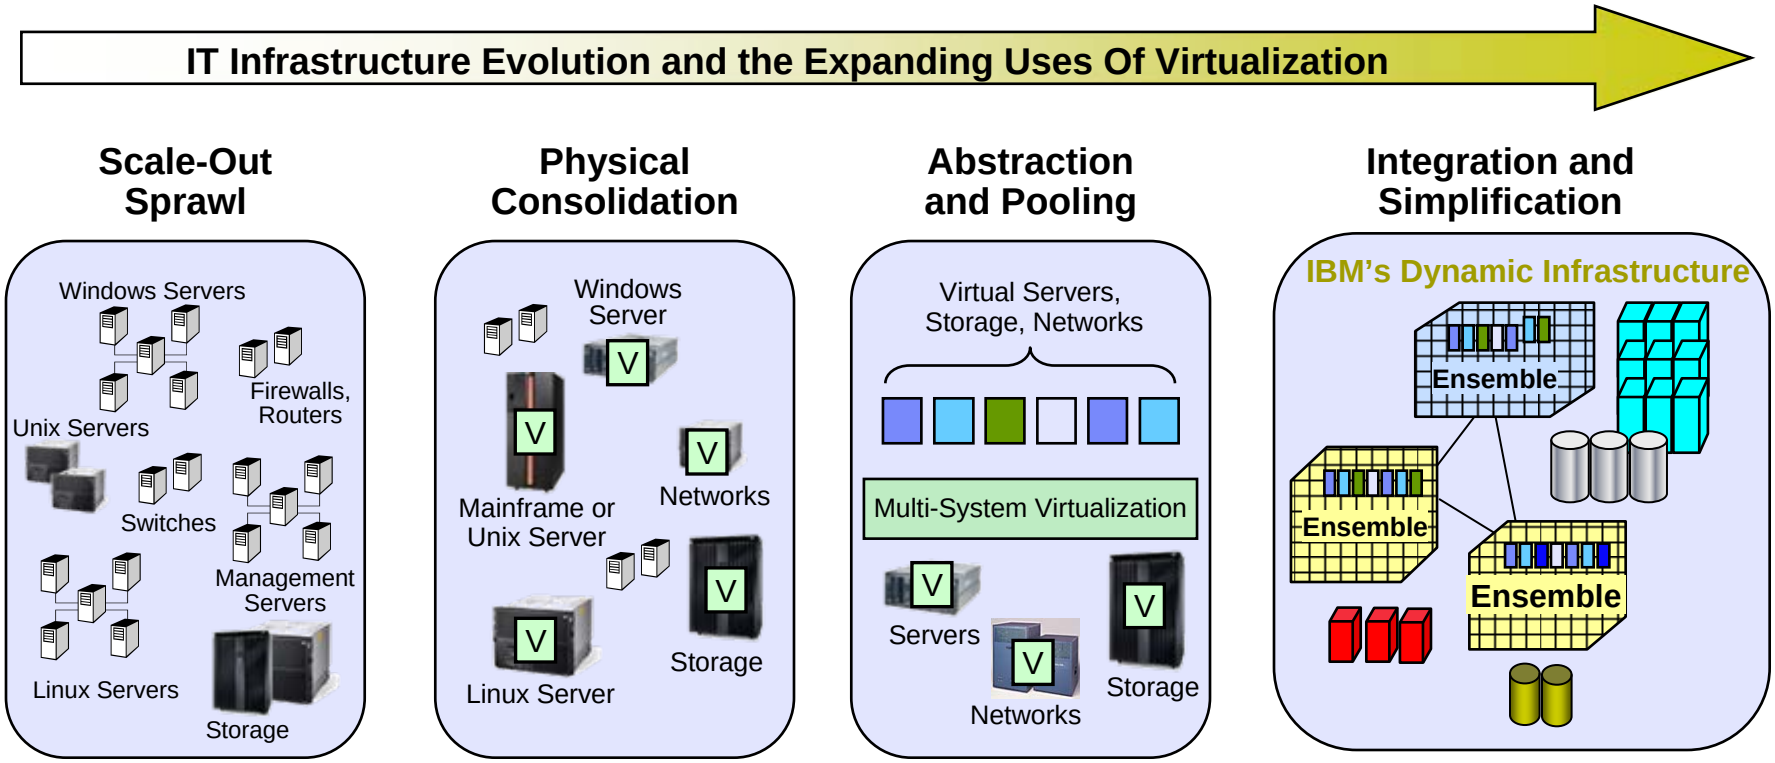

Workload Deployment and Management

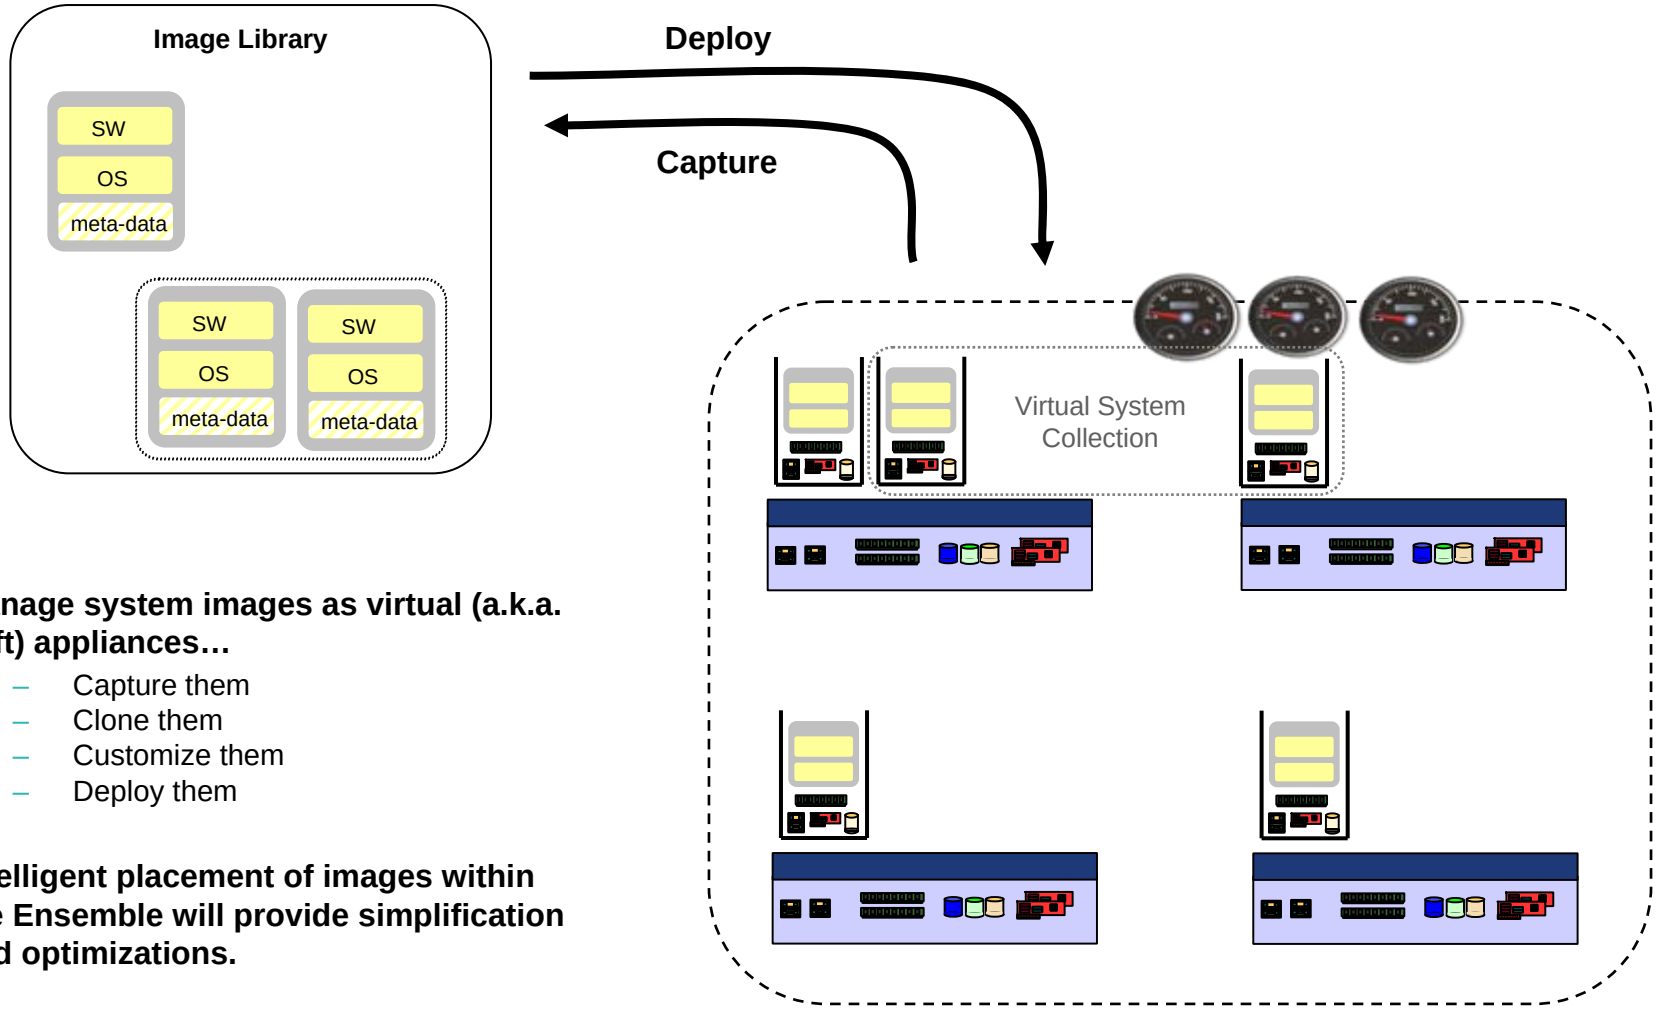

- **Manage system images as virtual (a.k.a. soft) appliances…**
  - Capture them
  - Clone them
  - Customize them
  - Deploy them
- **Intelligent placement of images within the Ensemble will provide simplification and optimizations.**

Virtualized System Pools (Ensemble)

View In Animation Mode

Adjusting Resource Allocations & Virtual Machine Mobility

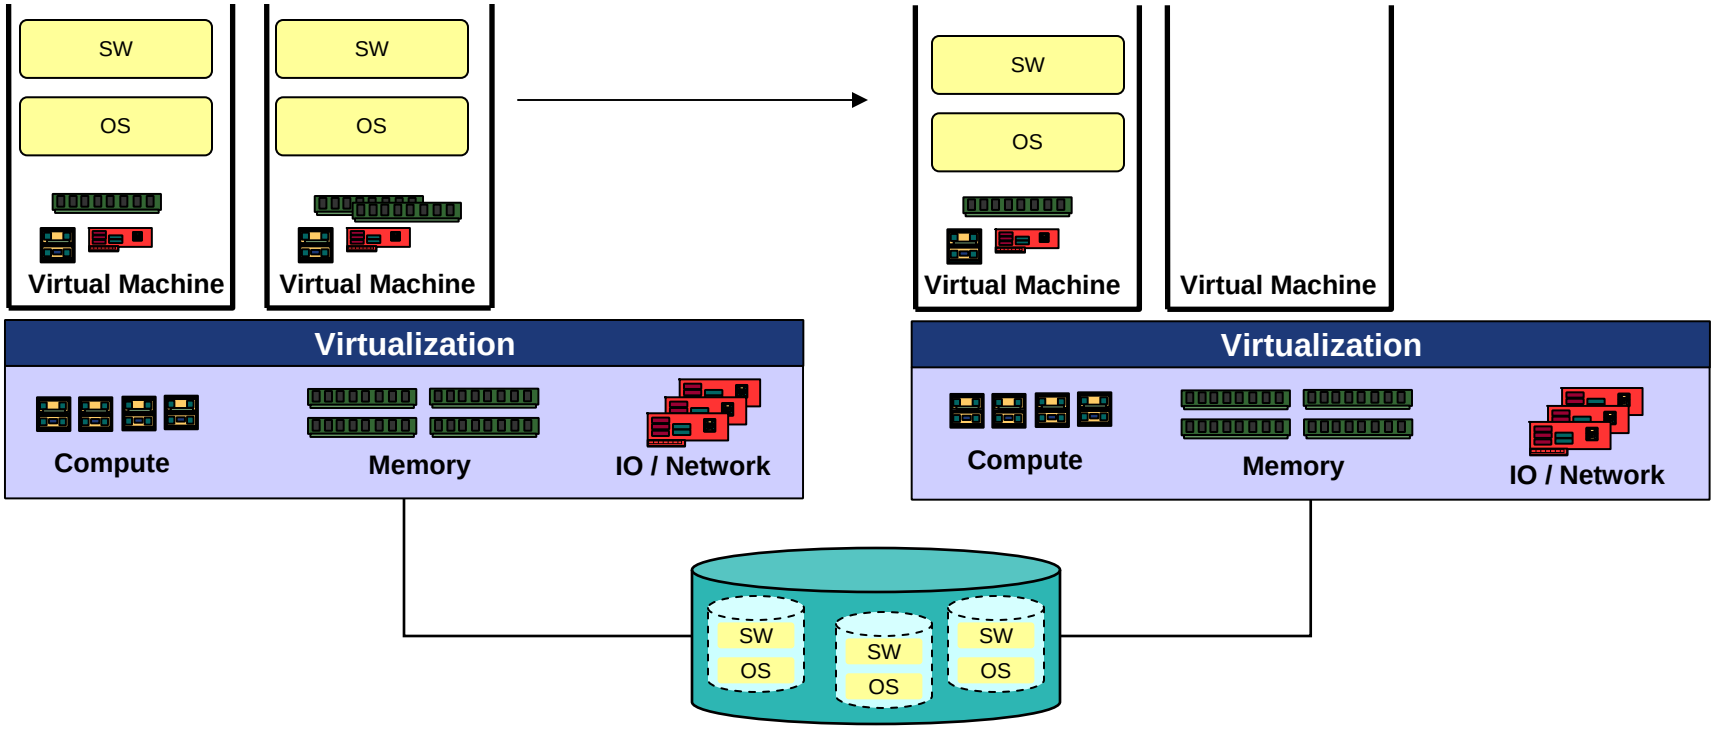

- **Dynamically adjust virtual machines resource allocations.**
  - Allowing unallocated resources to be used by a virtual machine.
  - Allowing resource allocation adjustments to be made between virtual machine.
- **Virtual server mobility between host systems.**
  - Allocate resources on the target host.
  - Move the virtual machine in-memory state to target host.
  - De-allocating resources on the source host.

View In Animation Mode

Non-Disruptive System Updates

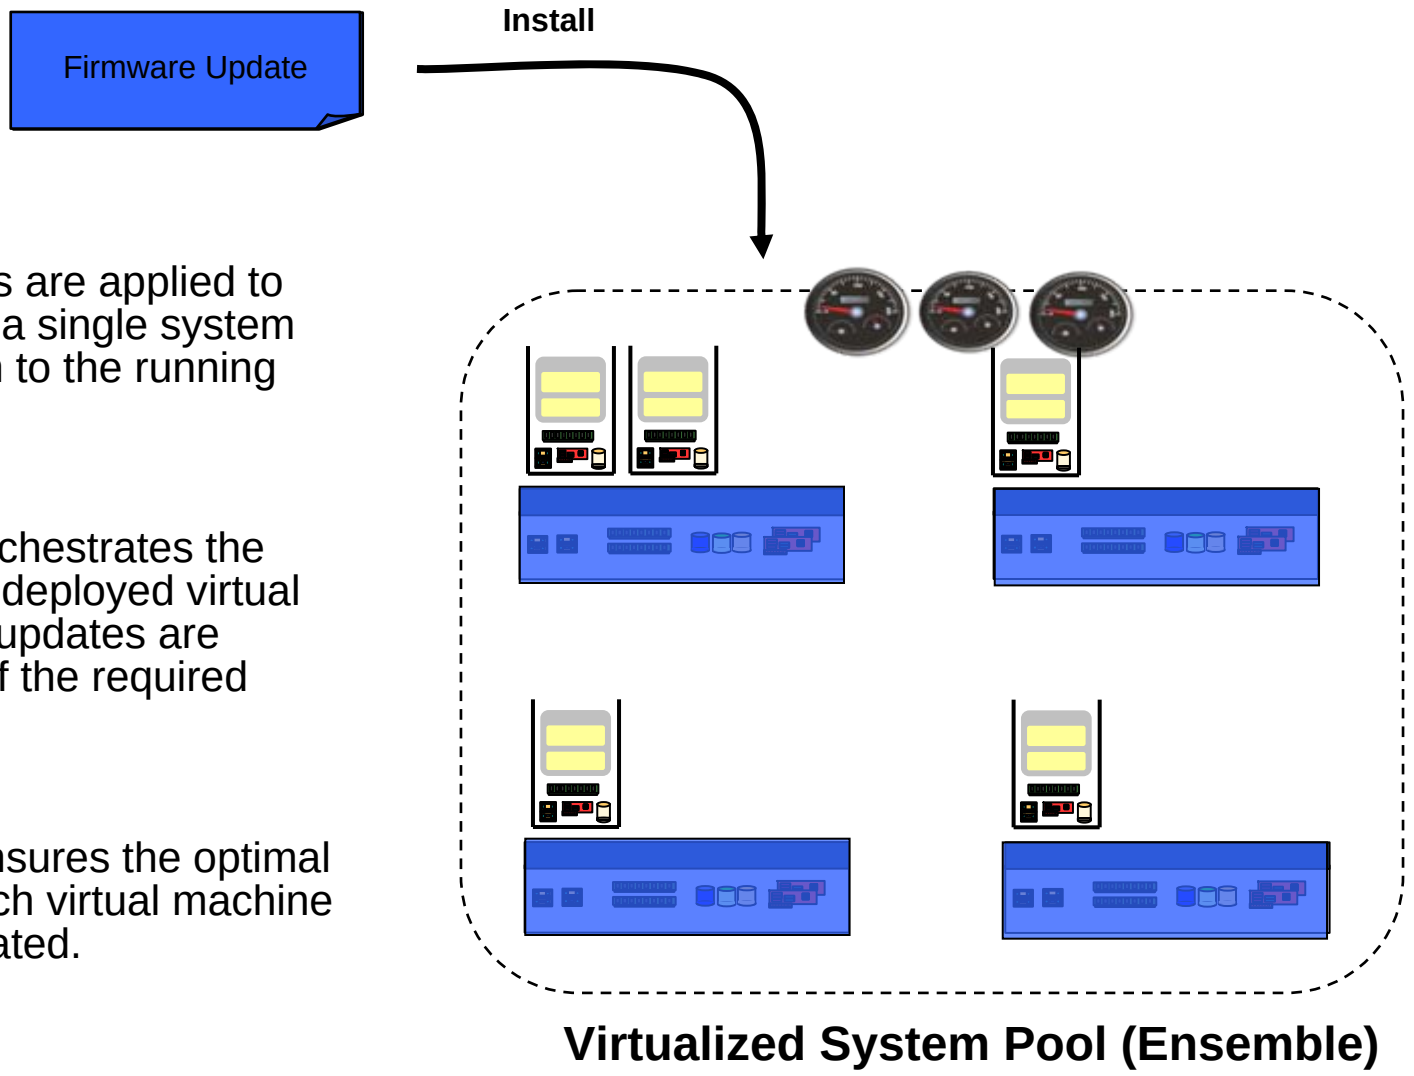

- Firmware updates are applied to the Ensemble as a single system with no disruption to the running workload.
- The Ensemble orchestrates the movement of the deployed virtual machines as the updates are applied to each of the required systems.
- The Ensemble ensures the optimal placement for each virtual machine as they are relocated.

View In Animation Mode

Automated Energy Optimizations

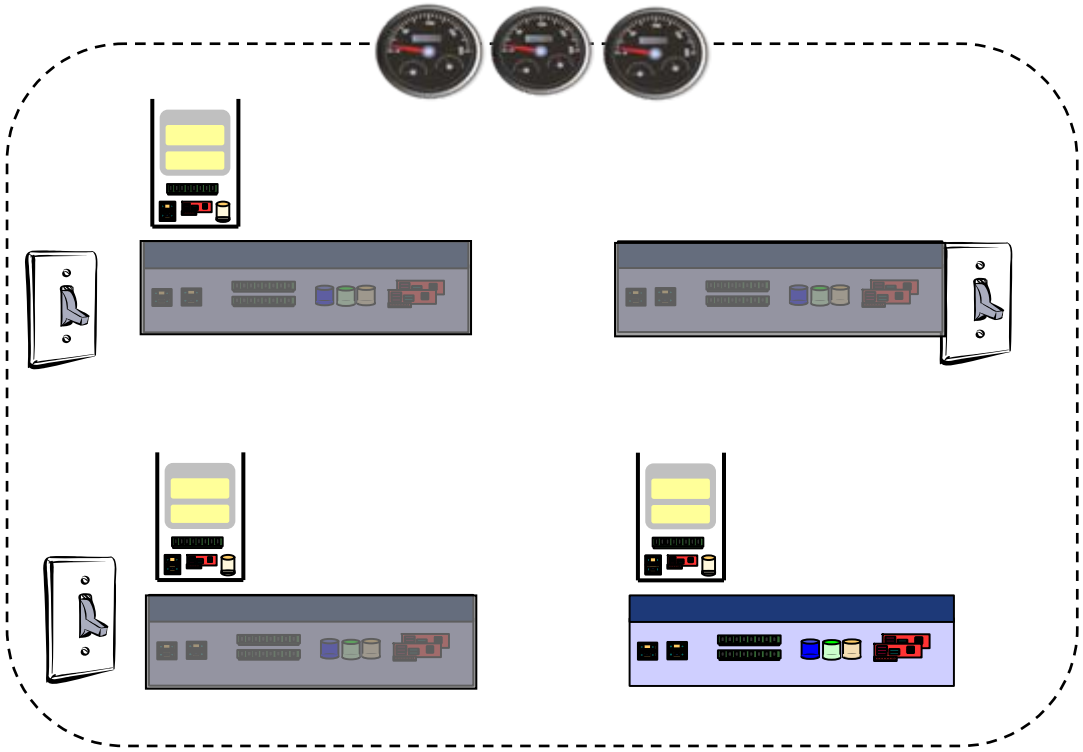

Virtualized System Pool (Ensemble)

- **Consolidate Virtual Servers on a fewer number of host systems.**
  - Move using ‘live’ virtual machine mobility (relocation).

- **Power Off / Suspend host systems that are currently not required.**

View In Animation Mode

Automating Workload Availability for Unplanned Downtime

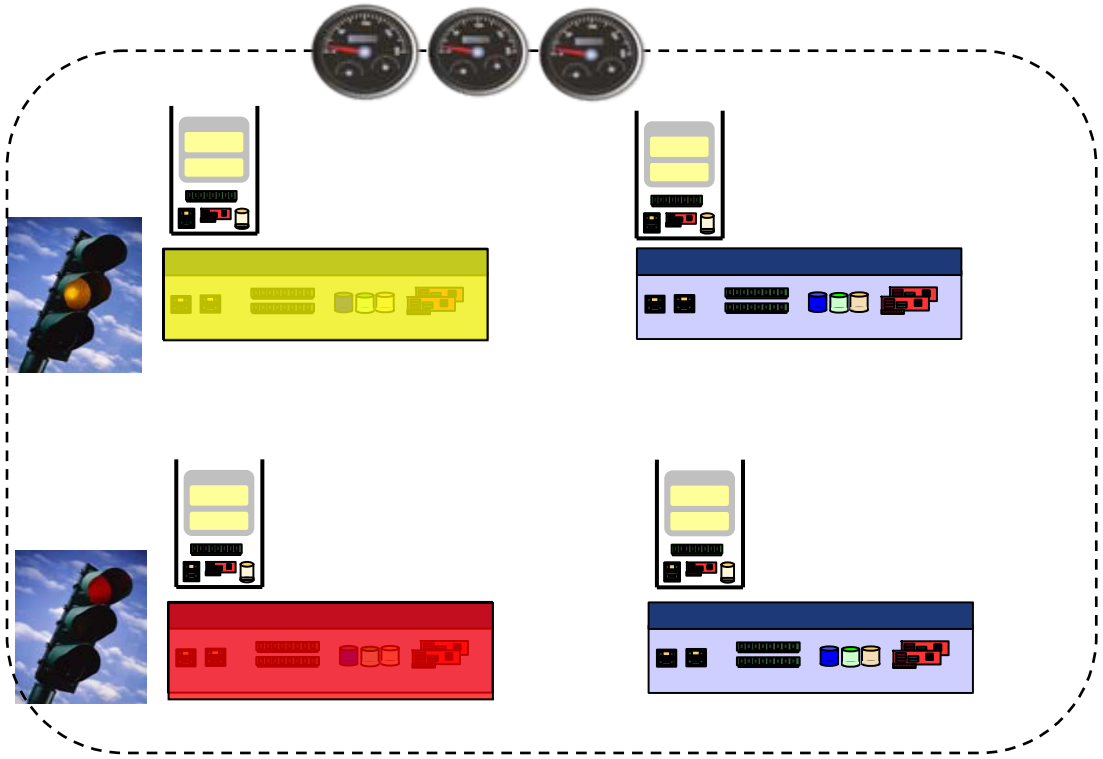

Virtualized System Pool (Ensemble)

- **Move virtual servers away from a failing host system.**
  - Use of HW Predicted Failure Analysis to drive ‘live’ virtual machine mobility.
- **Restart virtual servers when a host system fails.**
  - Restarting a virtual server (possibly from checkpoint) on another server is a form of ‘static’ migration.

View In Animation Mode

NC State VCL “Cloud in a Box” Proof of Concept

Overall Objectives

- Validate IBM cloud computing architecture in the VCL Environment
- Establish a BladeCenter based “Cloud in a Box” proof of concept to validate the ensemble pattern in the VCL environment.
- Feedback on the integrated management for image deployment and virtual systems management
- Validate advanced ensemble capabilities over time for security, availability, performance and energy management
- Enable the delivery of education solutions through VCL and ensembles

VCL Architecture – Overview

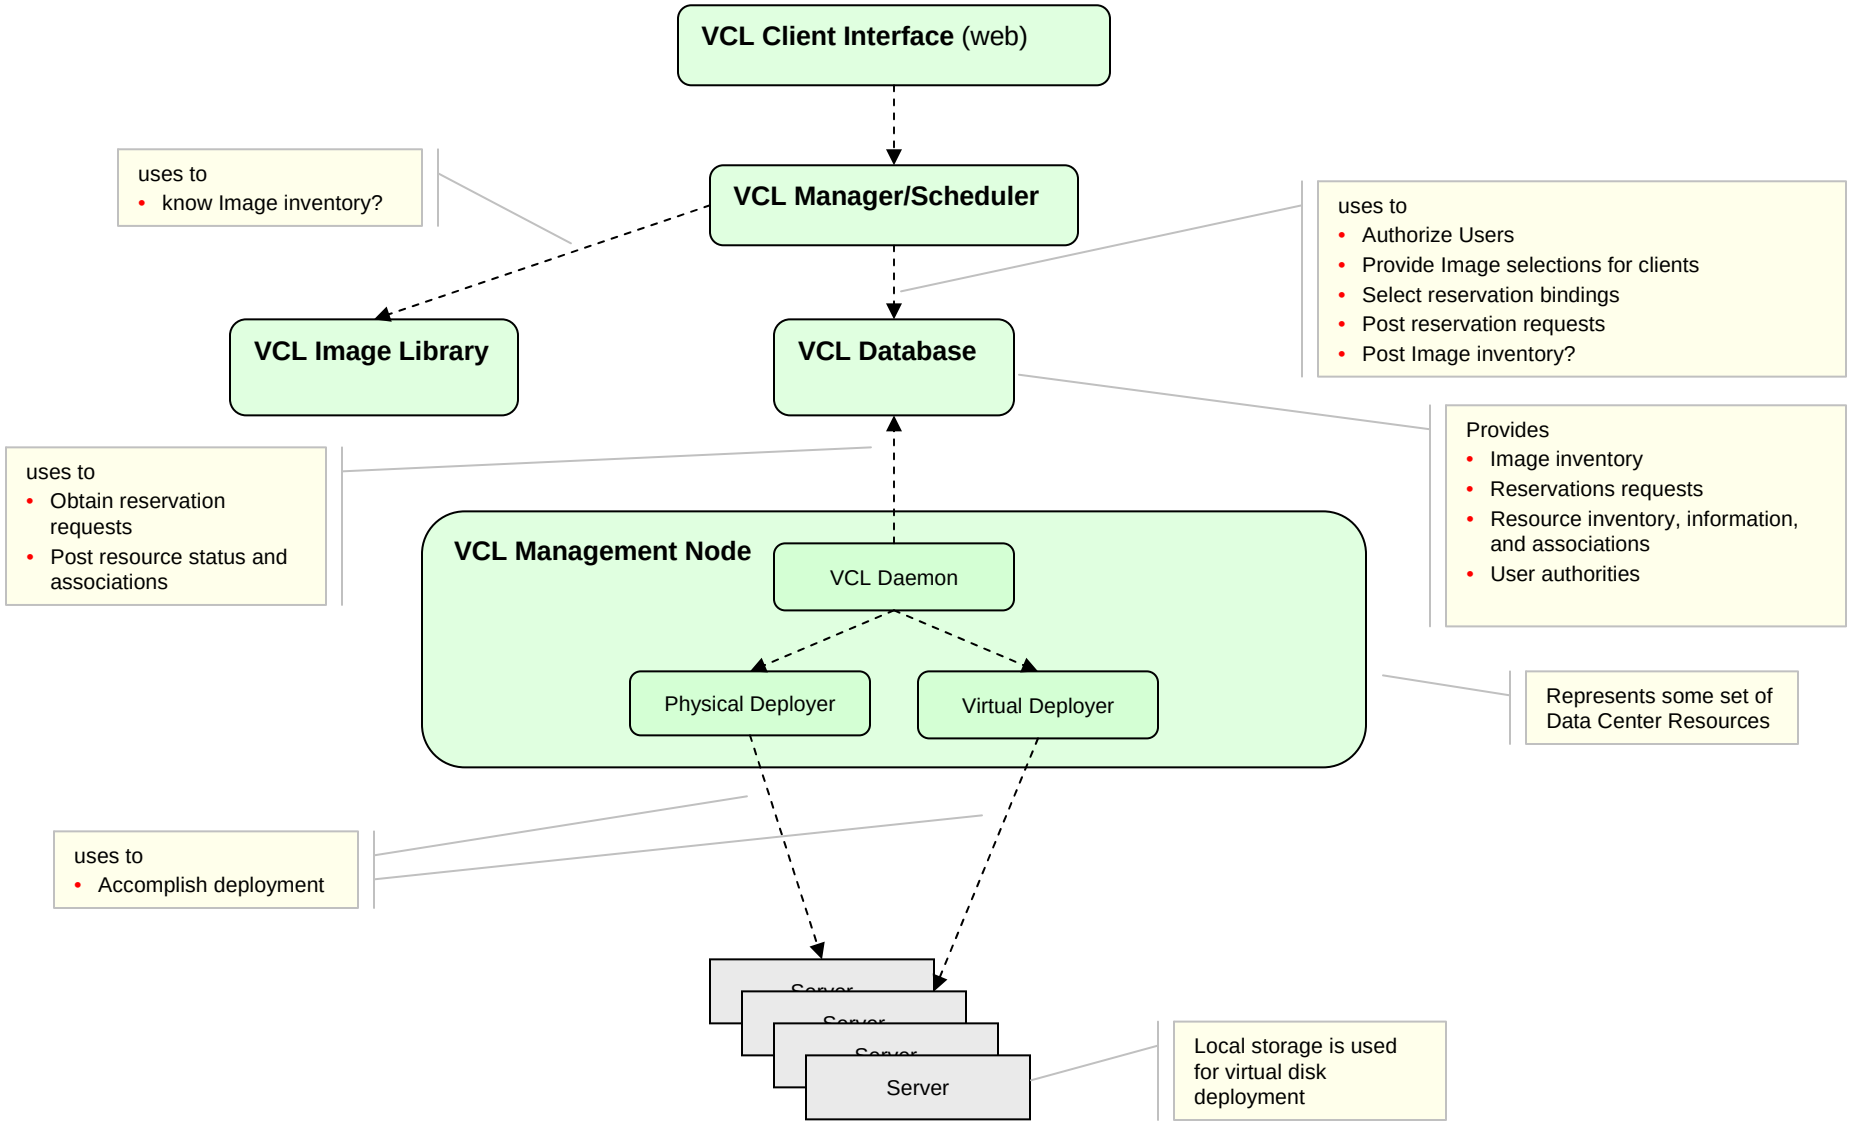

VCL Architecture – Ensemble Add, Overview

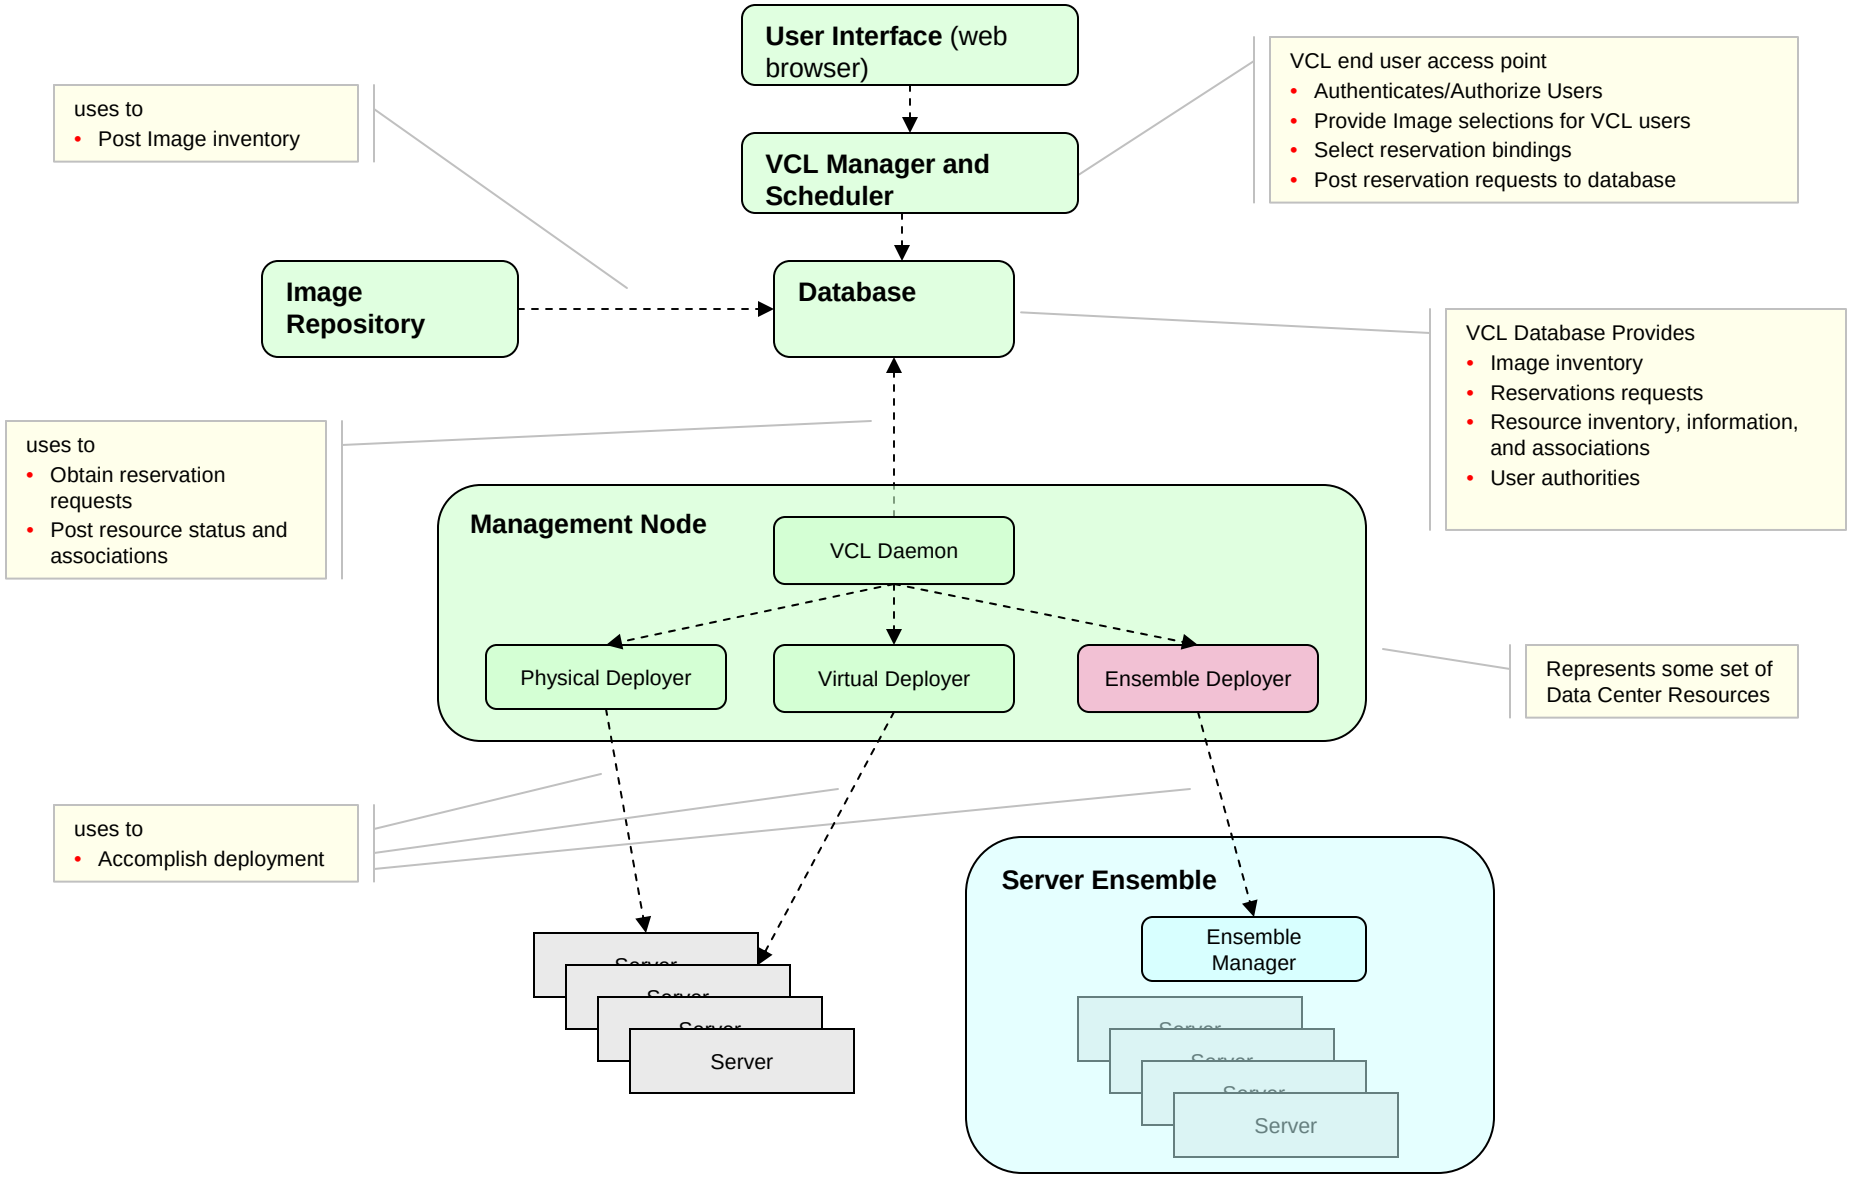

VCL Architecture – Ensemble and NAS Add, Overview

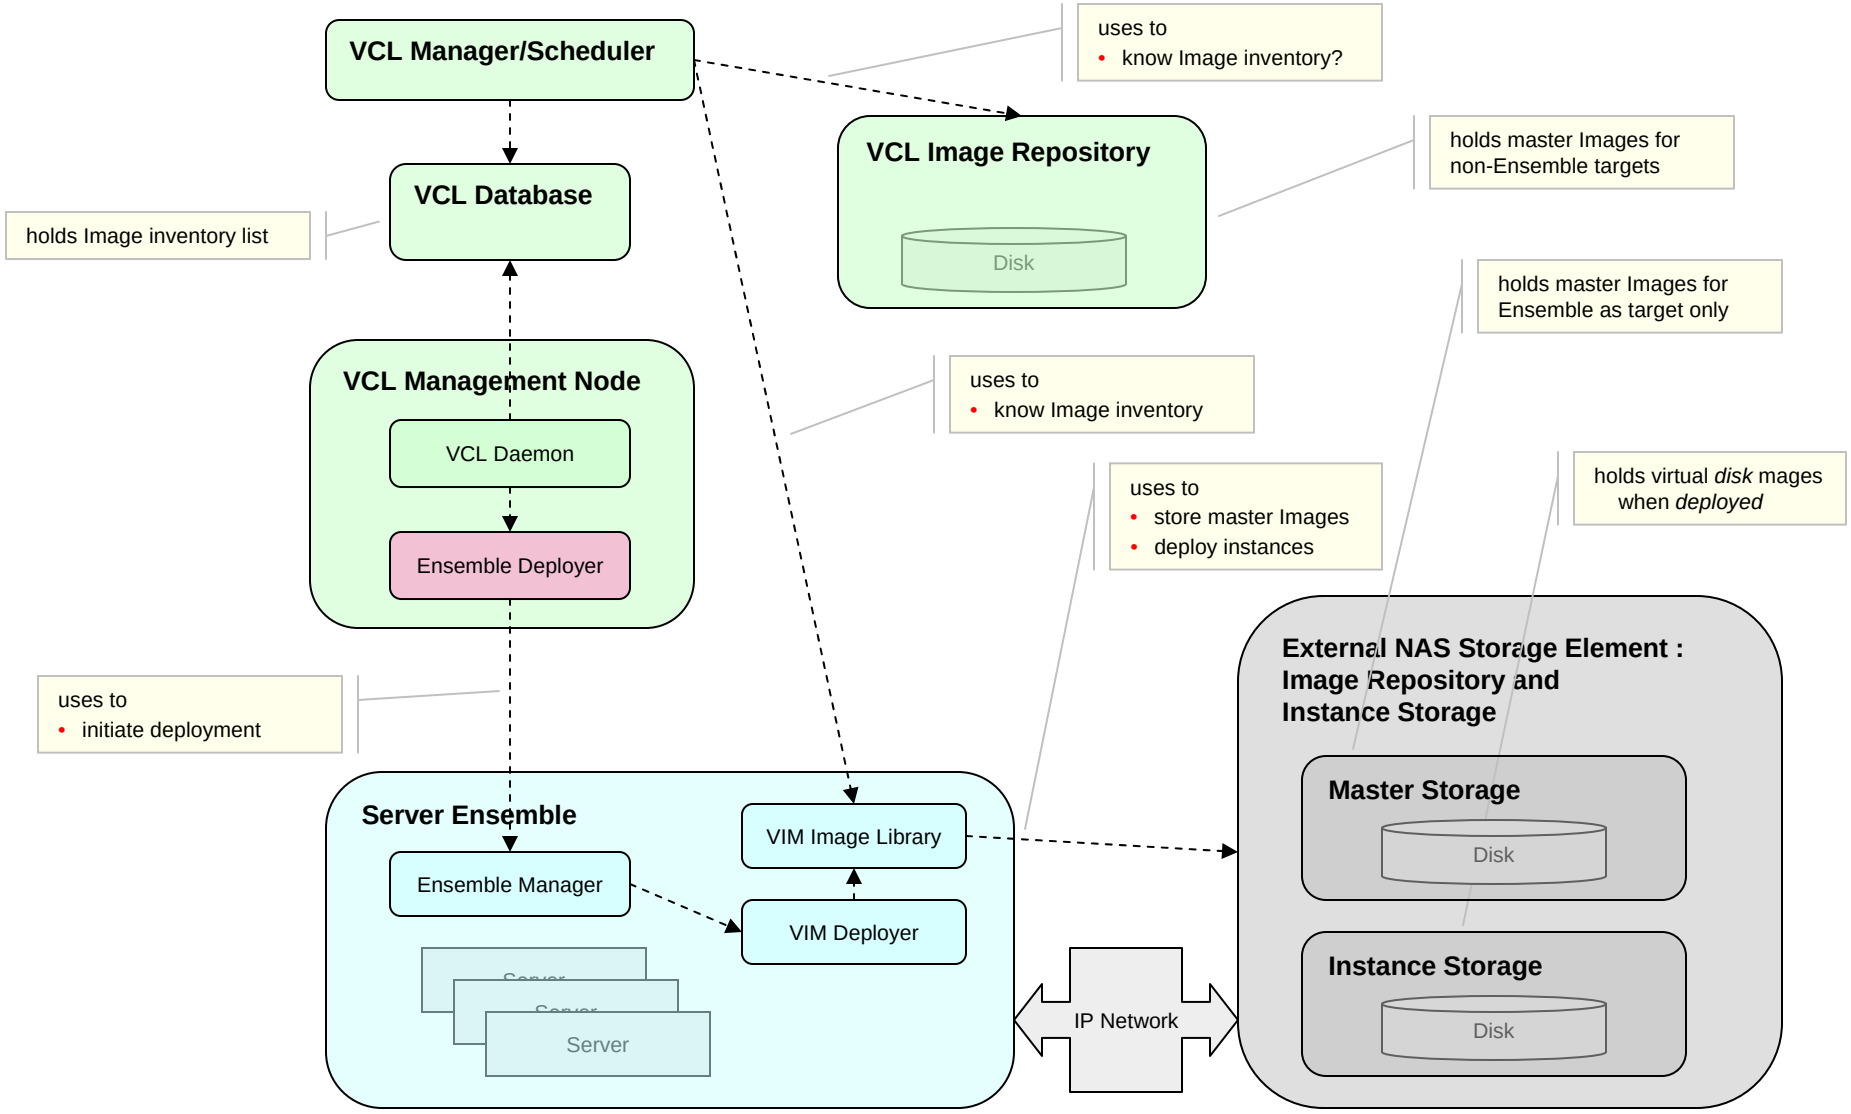

IBM / NC State VCL Collaboration

Future: Explore Advanced Optimizations and Automation

Set “Goals” at the Ensemble and Image level:

- Security (with SOSI initiative at NCSU)
- Availability (applying constraints and optimizations for VM placement / movement)
- Energy Management (optimizing performance and energy usage?)

Questions?

References and Useful Links

- IBM Think
- IBM Smarter Planet (Index)
- Wikipedia: Cloud Computing
- IBM Cloud Computing
- UC Berkeley paper on Cloud Computing

Trademarks and disclaimers

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

Trademarks of International Business Machines Corporation in the United States, other countries, or both can be found on the World Wide Web at http://www.ibm.com/legal/copytrade.shtml.

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.
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.
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.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

The customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance 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 or performance improvements equivalent to the ratios stated here.

Prices are suggested U.S. list prices and are subject to change without notice. Starting price may not include a hard drive, operating system or other features. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Photographs shown may be engineering prototypes. Changes may be incorporated in production models.

Backup

NORTH CAROLINA STATE UNIVERSITY

NC STATE UNIVERSITY

IBM

IBM | Global Technology Services

Cloud Computing – IBM Services Overview

New Complexities. New Risks. New Opportunities.

Craig Nygard, IBM Corporation

Everyone has an opinion …

In answer to the question, explain what cloud computing is? "It's hard to explain, you can talk to anyone here and they will give you a different version."

Shane Robison, CTO, Hewlett-Packard

"The interesting thing about cloud computing is that we've redefined cloud computing to include everything that we already do. We'll make cloud computing announcements. I'm not going to

Cloud computing is the use of networked infrastructure software and capacity to provide resources to users in an on-demand environment. … clouds provide a set of typically virtualized computers which can provide users with the ability to start and stop servers or use compute cycles only when needed, often paying only upon usage.

-http://www.vmware.com/technology/virtual-datacenter-os/cloud-vservices/faqs.html

Where are we now?

Our world is becoming

INSTRUMENTED

Our world is becoming

INTERCONNECTED

Virtually all things, processes and ways of working are becoming

INTELLIGENT

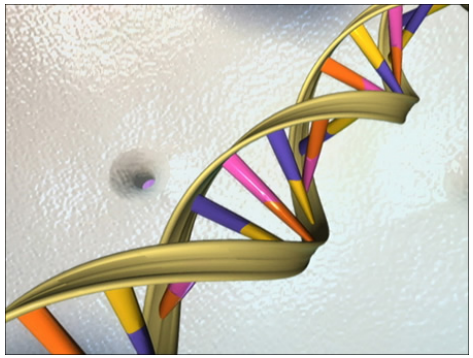

Initiatives that spring from becoming smarter will demand cloud services

Smarter planet: Thinking and acting in new ways to make our systems more efficient, productive and responsive.

SMART IS
Integrating all sources of knowledge about a student to provide a performance and results view.

SMART IS
Providing dynamic infrastructures to meet the everyday needs of a premier education institution.

SMART IS
Leveraging technology and scale to provide a better service at a lower cost for low value add applications

SMART IS
Reducing energy demand through virtualization technology.

SMARTER CLASSROOM:
New Intelligence for Student Success

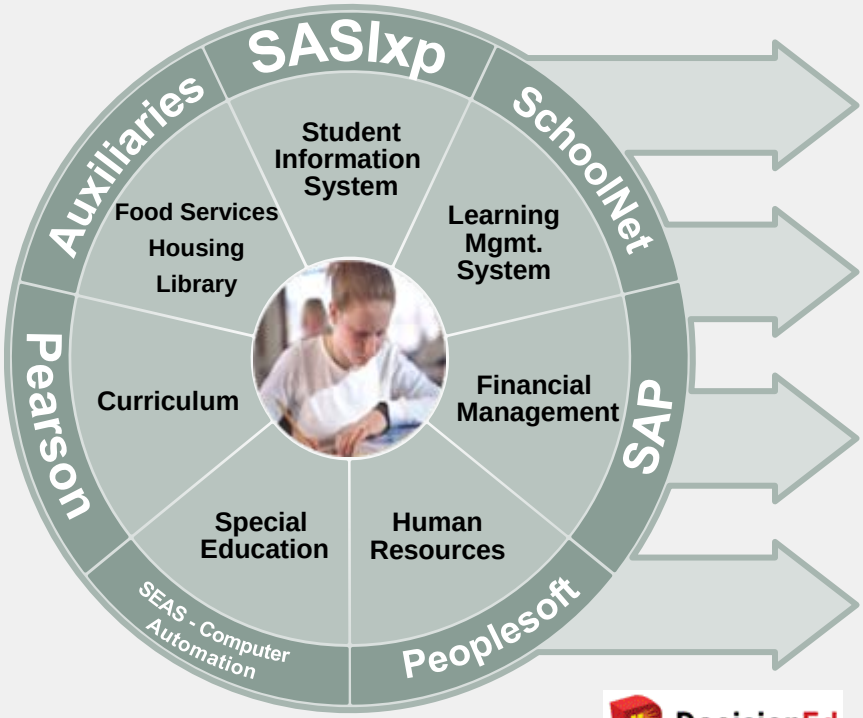

DecisionEd

Educational Performance Management with Cognos:
Invisible (Data) to Visible (Relevant Information)

METRICS

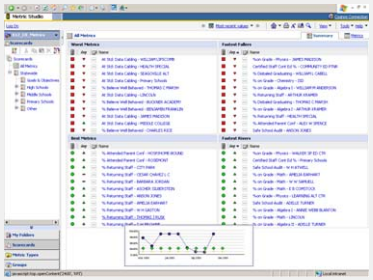

DASHBOARDS

ANALYTICS

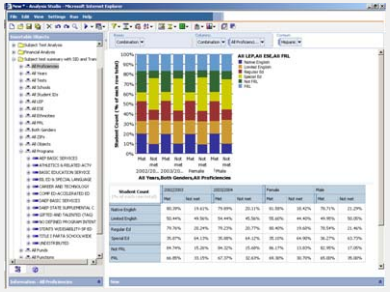

MANAGED REPORTING

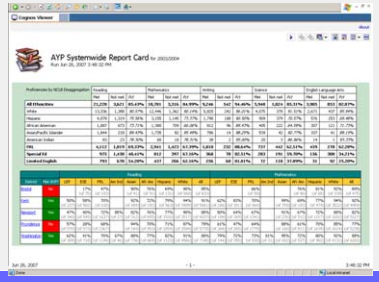

SMARTER CLASSROOM:
A Dynamic Infrastructure with Cloud Computing
North Carolina State University

North Carolina State University operates 13 off-campus regional research and extension centers and nine field laboratories.

Fact: The average number of servers is expected to grow another 50 percent by 2010.

What's smart?

- Supporting K-20 students statewide with a single cloud-based, virtual computing lab environment
- Leveraging open source, community developed software

Smarter Educational Outcomes:

- Lowers burden on support staff to support large number of users. Employs 5 staff members for 60,000 virtual images

SMARTER CLASSROOM:
Green technology through desktop virtualization

Wake Technical Community College

Wake Tech, in Raleigh, NC, is the state's flagship community college with over 57,000 students

Fact: The acquisition of a typical PC or laptop is only 25% of its total expense over its lifetime

What's smart?

- Deployed first of over 100 virtual labs with 25 clients for student classrooms
- Used Wyse V10L thin clients with VMware virtual desktop environment installed on the IBM BladeCenter

Smarter Educational Outcome

- Desktop virtualization with thin clients can reduce energy consumption reduction up to 45 percent

YACD*

The **INSTRUMENTED** world requires ***stream*** computing solutions.

The **INTERCONNECTED** world requires ***resilient*** computing solutions.

The **INTELLIGENT** world requires ***transformational*** solutions.

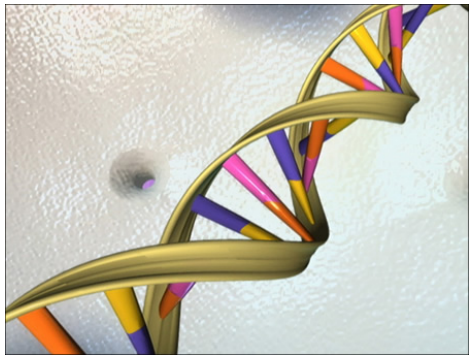

* Yet Another Cloud Definition

Smarter Educational systems leverage 21st Century technology to improve operations, increase the user experience and lower costs

Student access and administration built around **virtual desktops**

Virtualized computer resources of legacy desktop applications and services, using **Open Source** to lower costs.

Open Source eLearning tools and content to provide personalized learning pathways

Virtualized Cloud Services centrally supports a distributed set of campuses and classrooms

IBM hosted delivery as an option

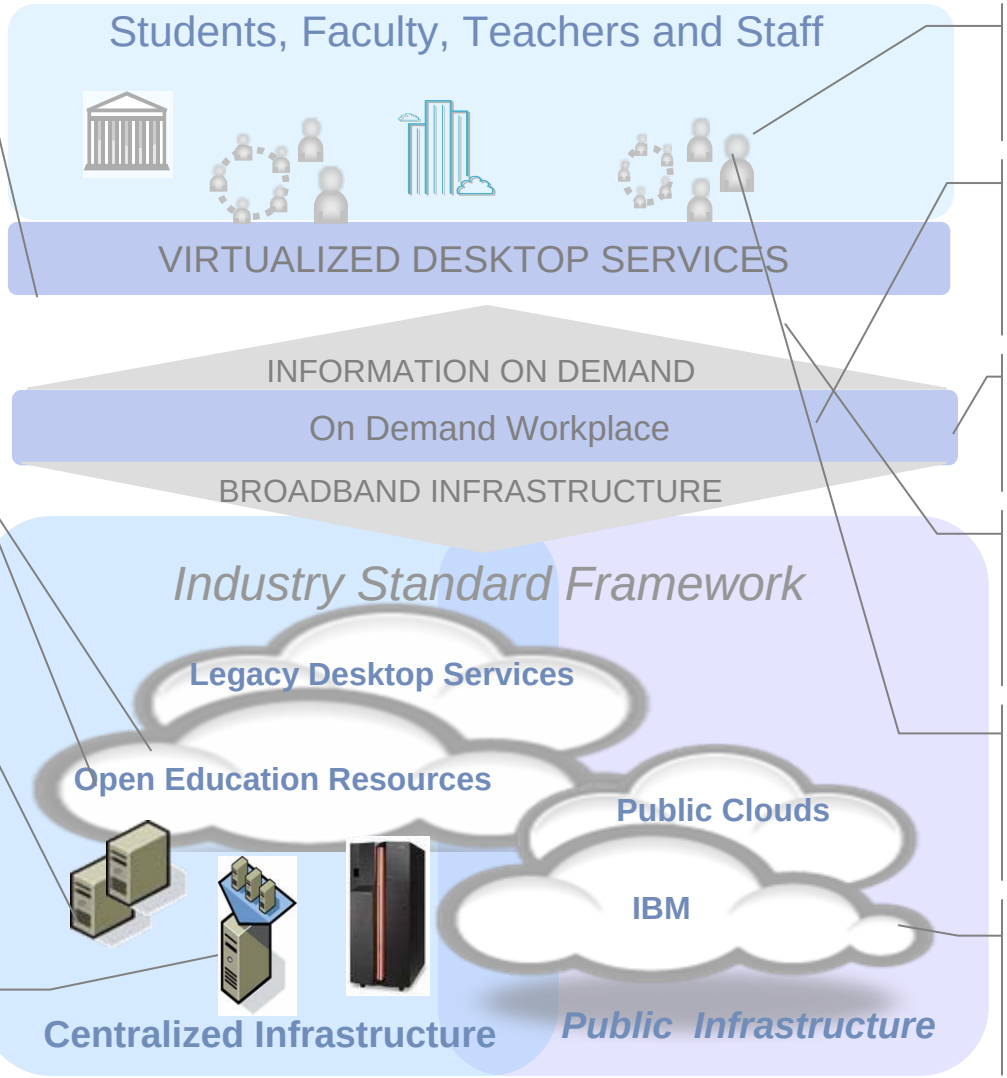

Thin Clients and **Mobile Devices** allow every user easy accessy

Business Intelligence provides insights on student and operational performance

Integrated Portal provides consolidated access to systems

Administrative Services provide for management of resources and assets

Network Services provide high speed connectivity between thin clients and servers

Web Services from IBM and others for collaboration and productivity

Cloud-onomics…

CLOUD COMPUTING

VIRTUALIZATION + STANDARDIZATION + AUTOMATION = Reduced Cost

….leverages virtualization, standardization and automation to free up operational budget for new investment

AGILITY + BUSINESS & IT ALIGNMENT + SERVICE FLEXIBILITY + INDUSTRY STANDARDS = OPTIMIZED BUSINESS

… allowing you to optimize new investments for direct business benefits

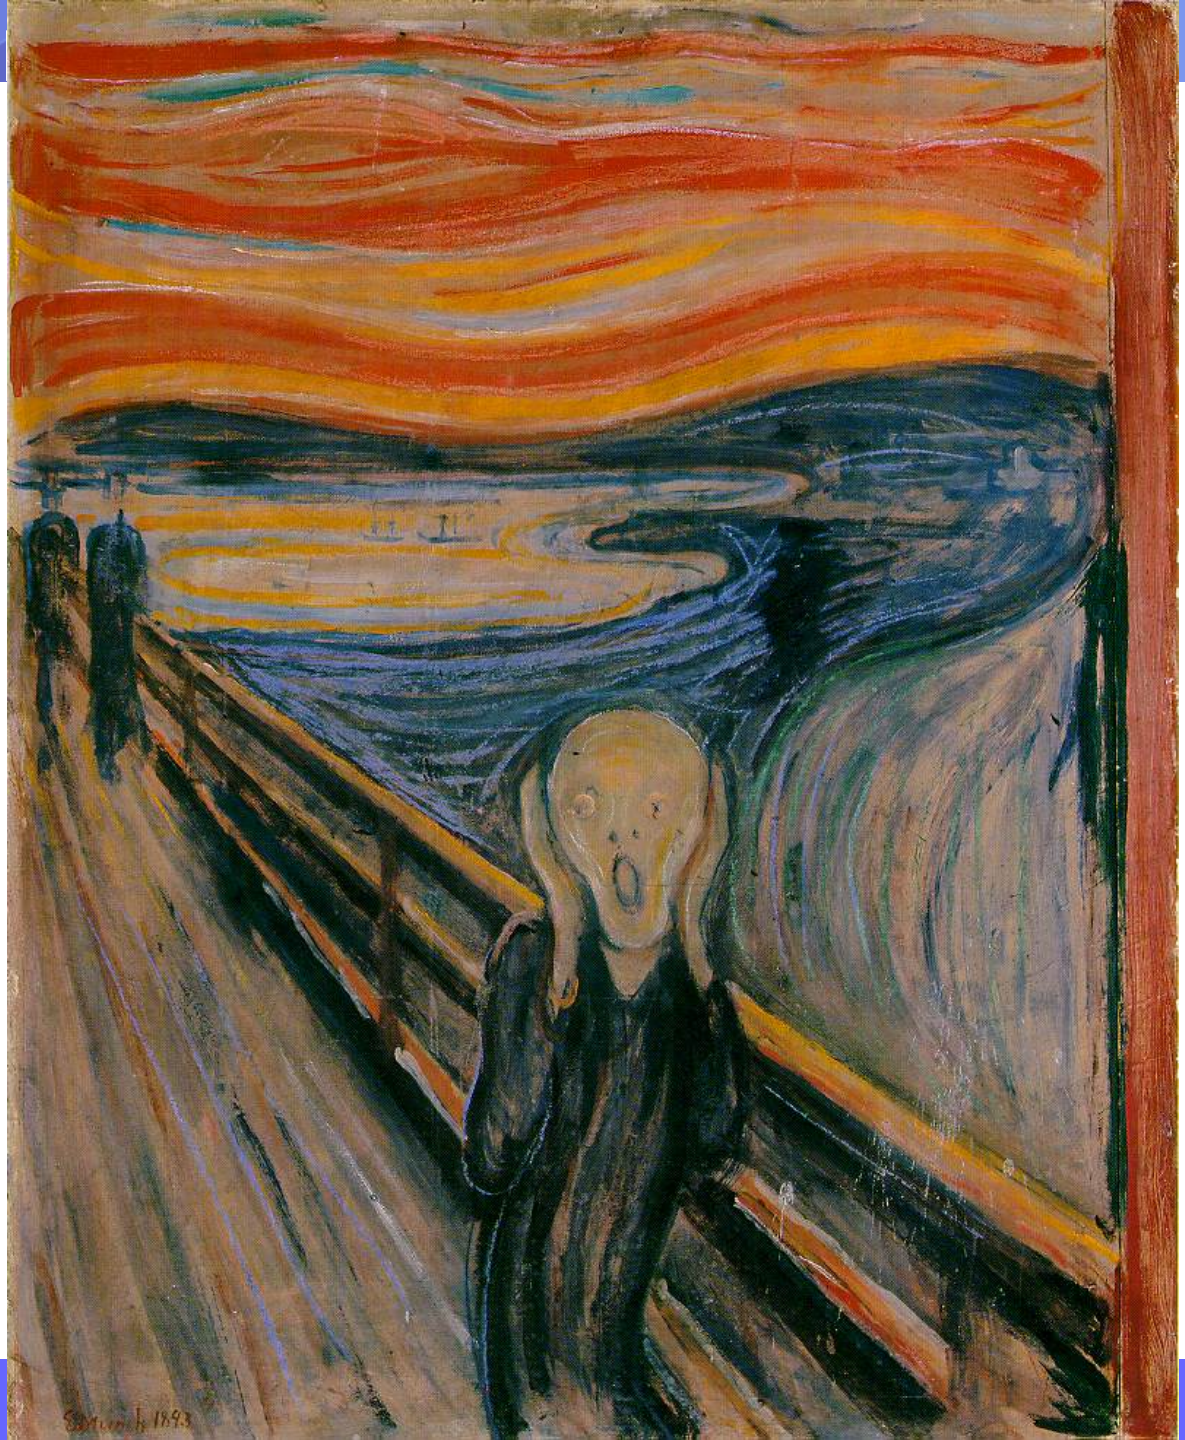

Cloud Computing: Enabling growth & innovation for your business

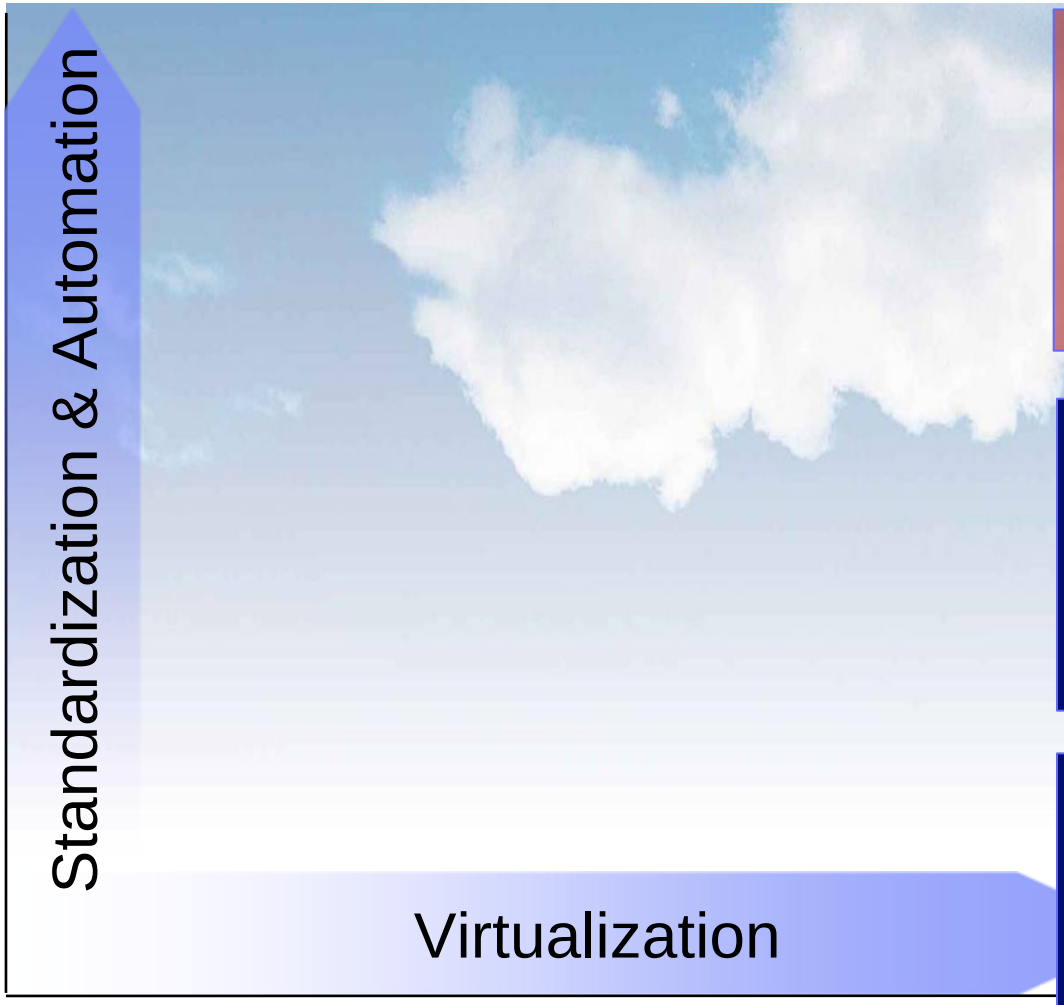

- ***Greater agility,*** ability to adjust to business requirements and market forces on demand
- ***Improved risk management*** through improved business resiliency
- ***More efficient pricing model,*** eliminating cost of excess capacity

- ***Better and flexible service*** for users, enabling self-service requests and delivering services more rapidly, with fewer errors.
- ***Improved time to market*** and acceleration of innovation projects

- ***Lower costs***, both capital and operational expenditures
- ***Free up skilled resources*** to focus on high value work and innovation projects
- Significantly ***improve energy efficiency and reduce idle time***

Where are clients implementing cloud computing?

Workloads where risk and migration cost may be too high and need elastic scalability:

- Database
- Transaction processing
- ERP workloads
- Highly regulated workloads

Workloads which can be standardized for cloud:

- Web infrastructure applications
- Collaboration infrastructure
- Development and test
- High Performance Computing

New workloads from existing or transformational applications made possible by cloud:

- High volume, low cost analytics – streaming data
- Collaborative Business Networks –
- Industry scale “smart” applications – scaled up

The future: Three co-existing delivery models

Over time, IT workloads will move to Cloud delivery models as applicable for the client.

Examples:

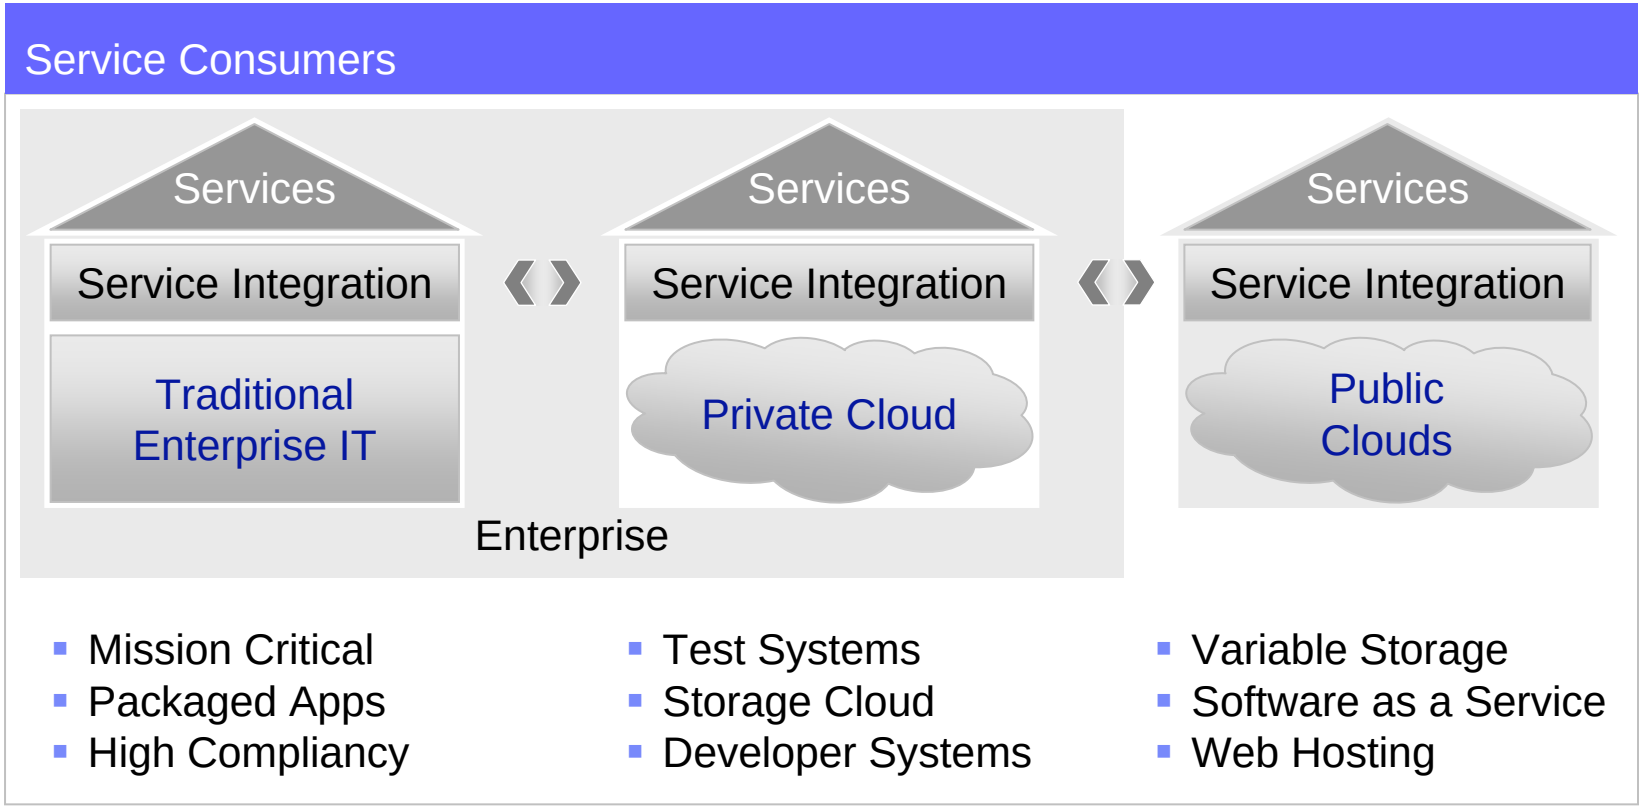

IT transformation includes Cloud Computing within IBM

IBM Technology Adoption Program (TAP)

Saving IBM over $2.5M per year

Self-service, on demand IT delivery solution for 3,000 IBM researchers across 8 countries

Enterprise class utility computing solution for clients

Systems platform testing for Enterprise Clients, SMBs, & ISVs

Cloud computing solution for IBM Learning Centers in Europe

Business Case Results : IBM TAP Greenfield Cloud Deployment

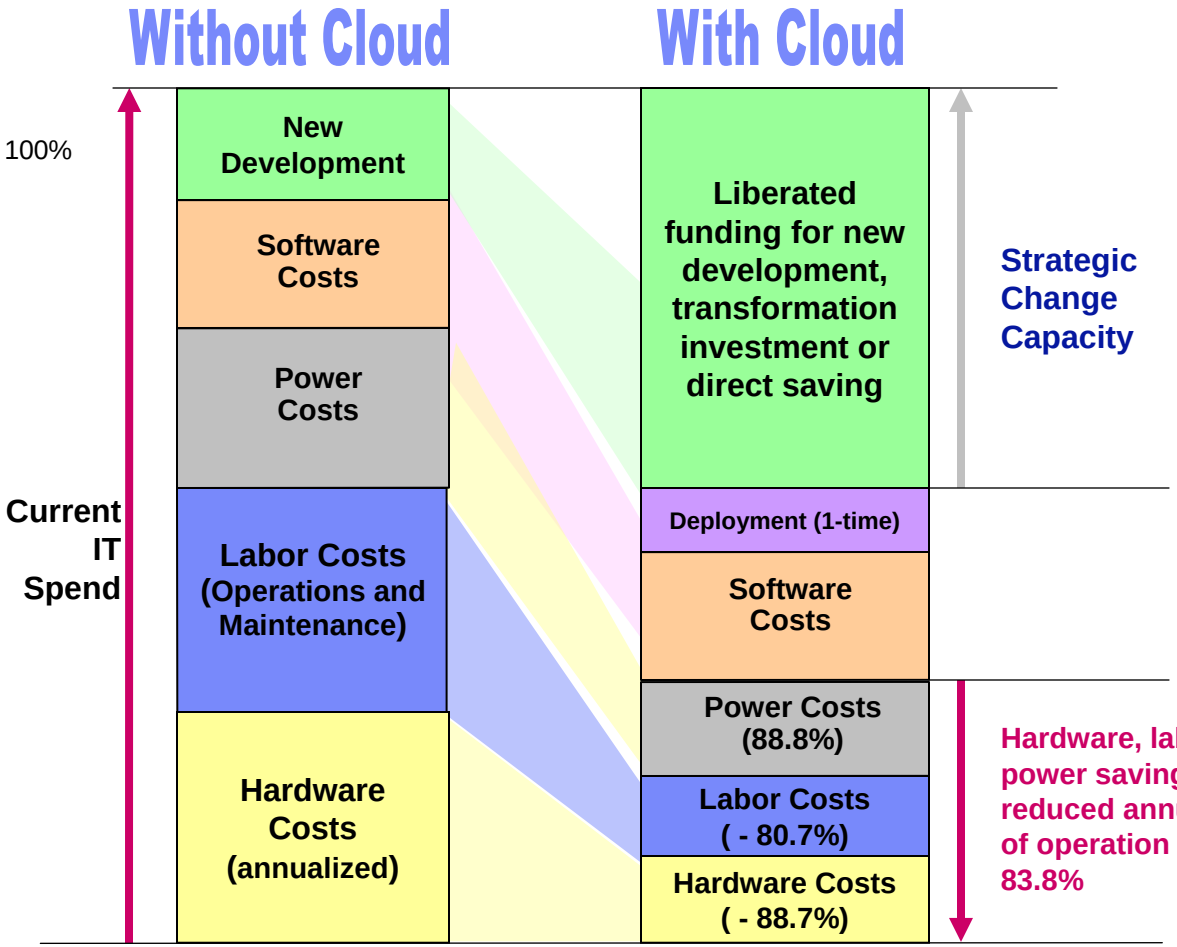

Note: 3-Year Depreciation Period with 10% Discount Rate

Business Case Results

Annual savings: $3.3M (84%)
$3.9M to $0.6M

Payback Period: 73 days

Net Present Value (NPV): $7.5M

Internal Rate of Return (IRR): 496%

Return On Investment (ROI): 1039%

IBM Research Computing Cloud (RC2)

A living lab to advance Research strategies

1. **Provides self service “on demand” delivery solution for research computing resources**

2. **Zero touch support for the full life cycle of service delivery**
   - Order creation
   - Approval process
   - E-mail notification
   - Automated provisioning
   - Monitoring

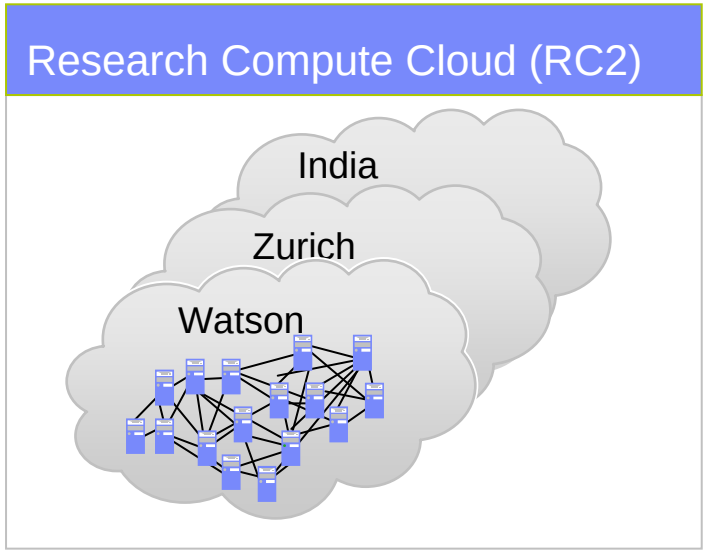

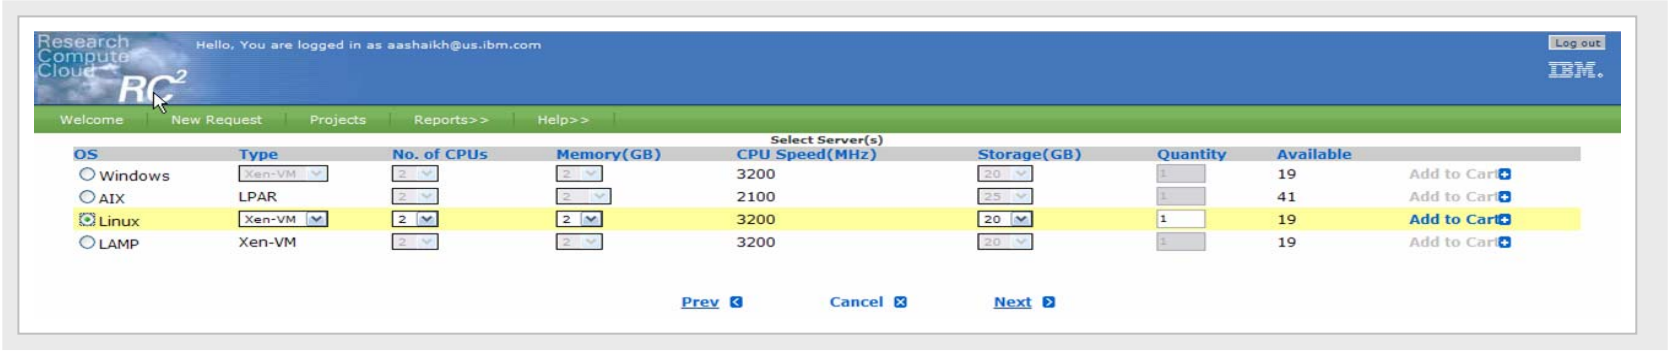

Data Centers

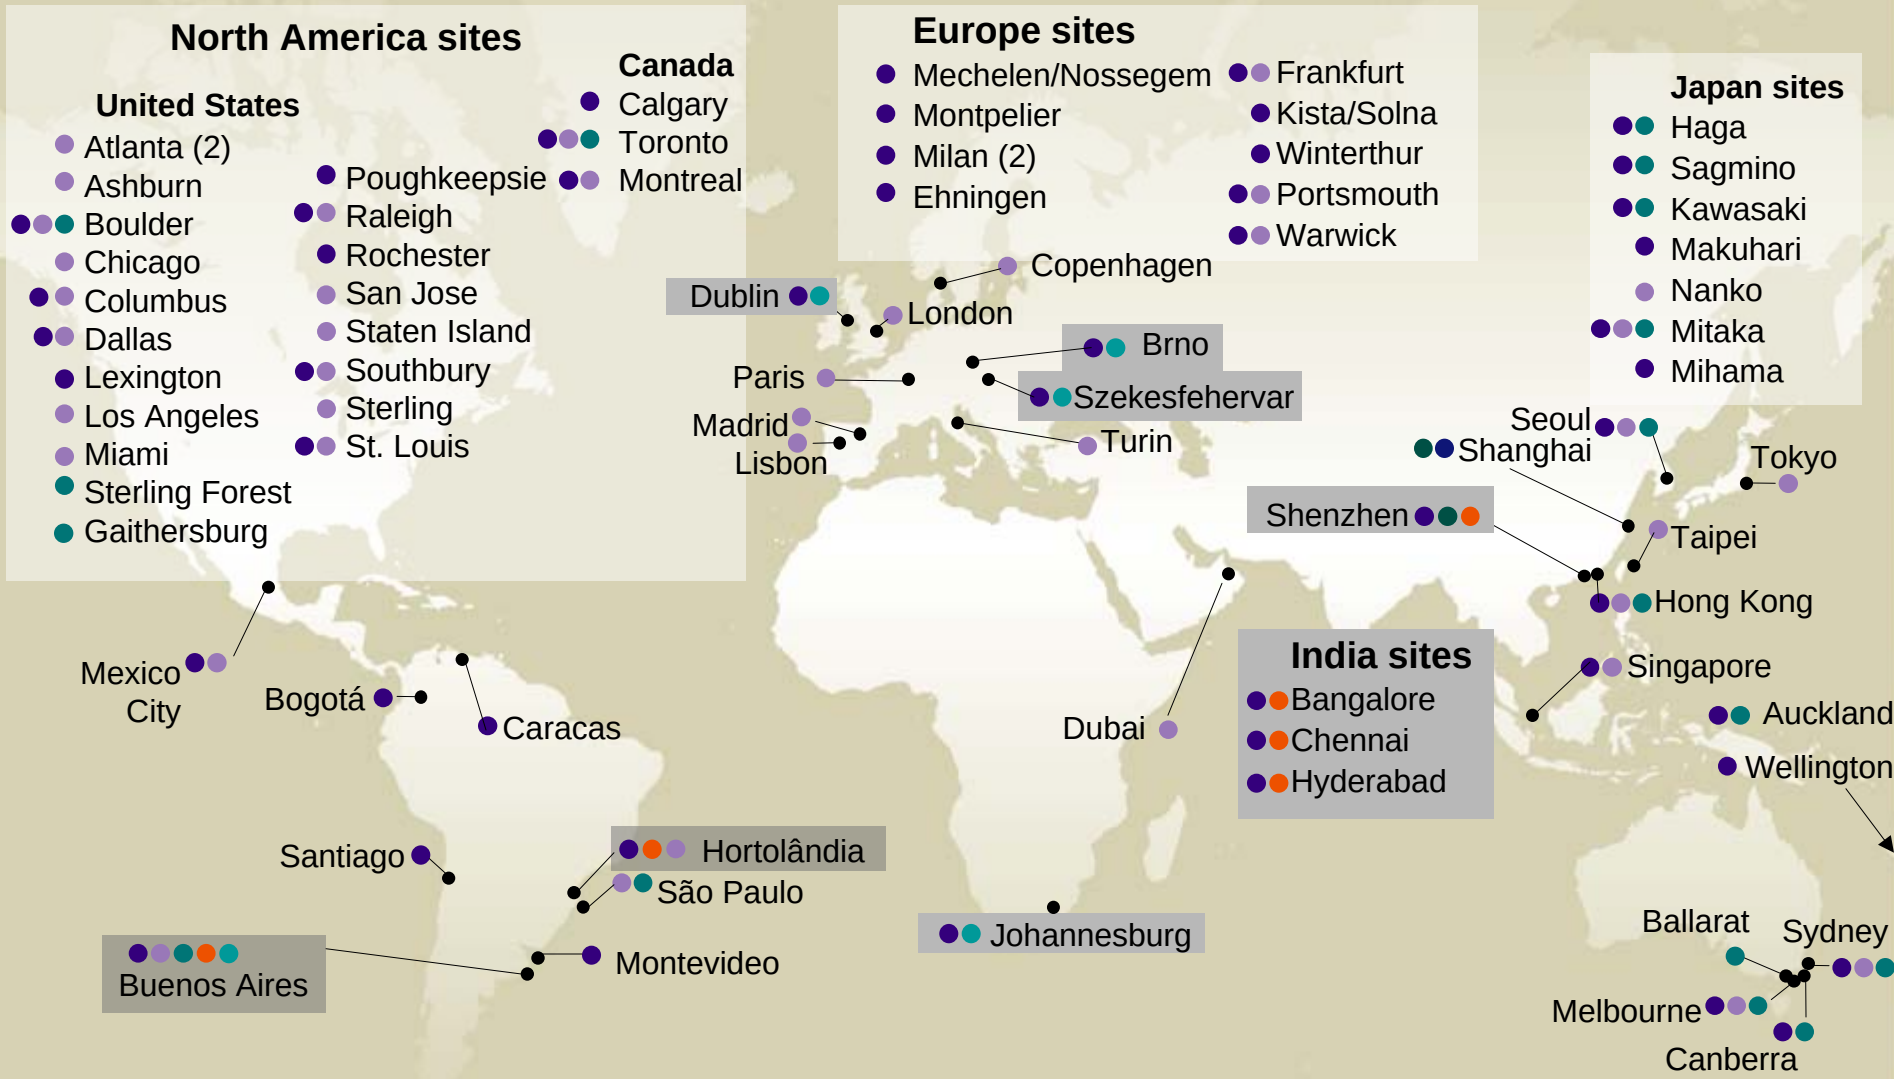

IBM Cloud Computing Centers

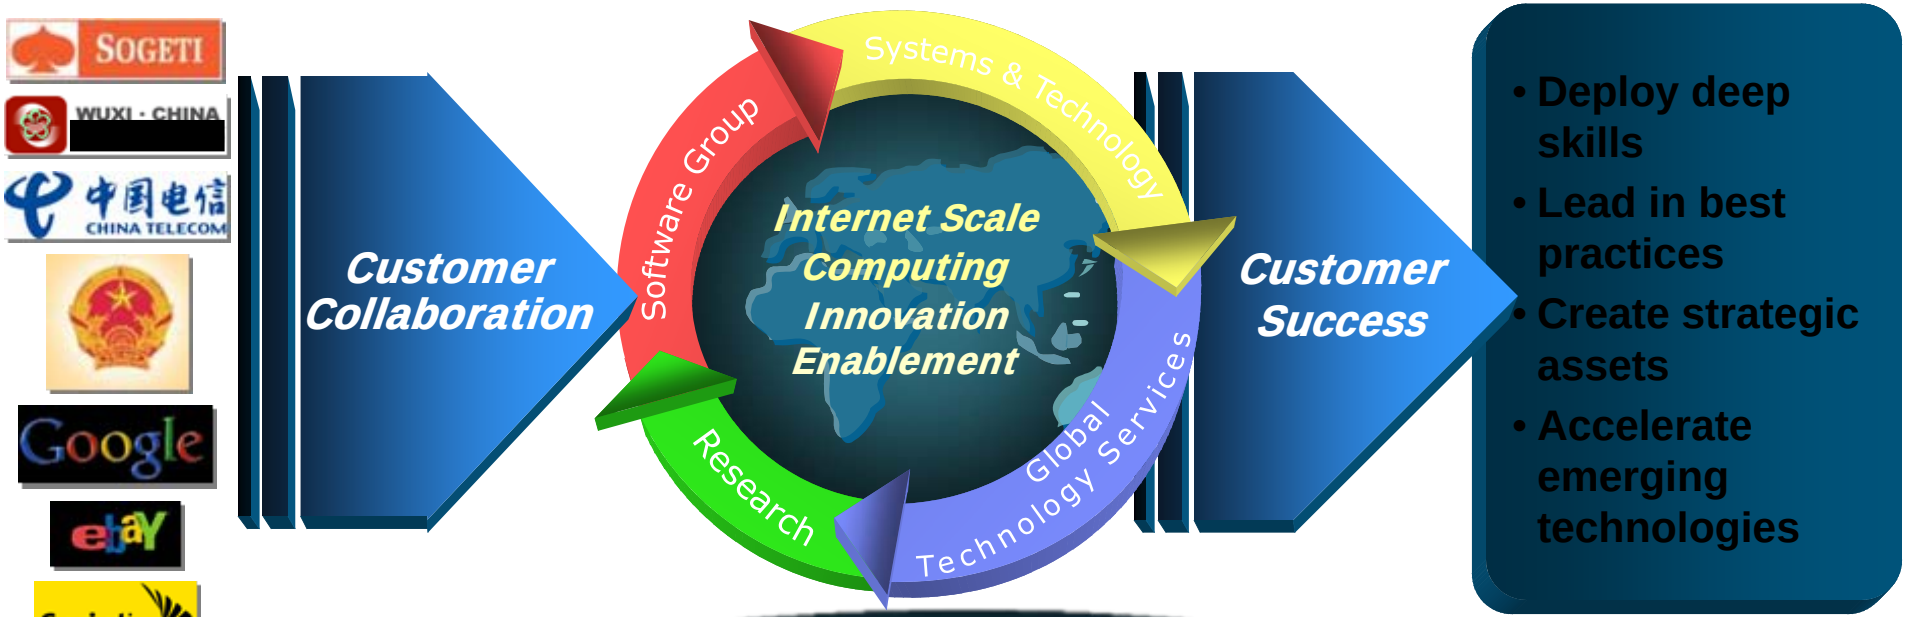

RTP Building 301 Floor Plan

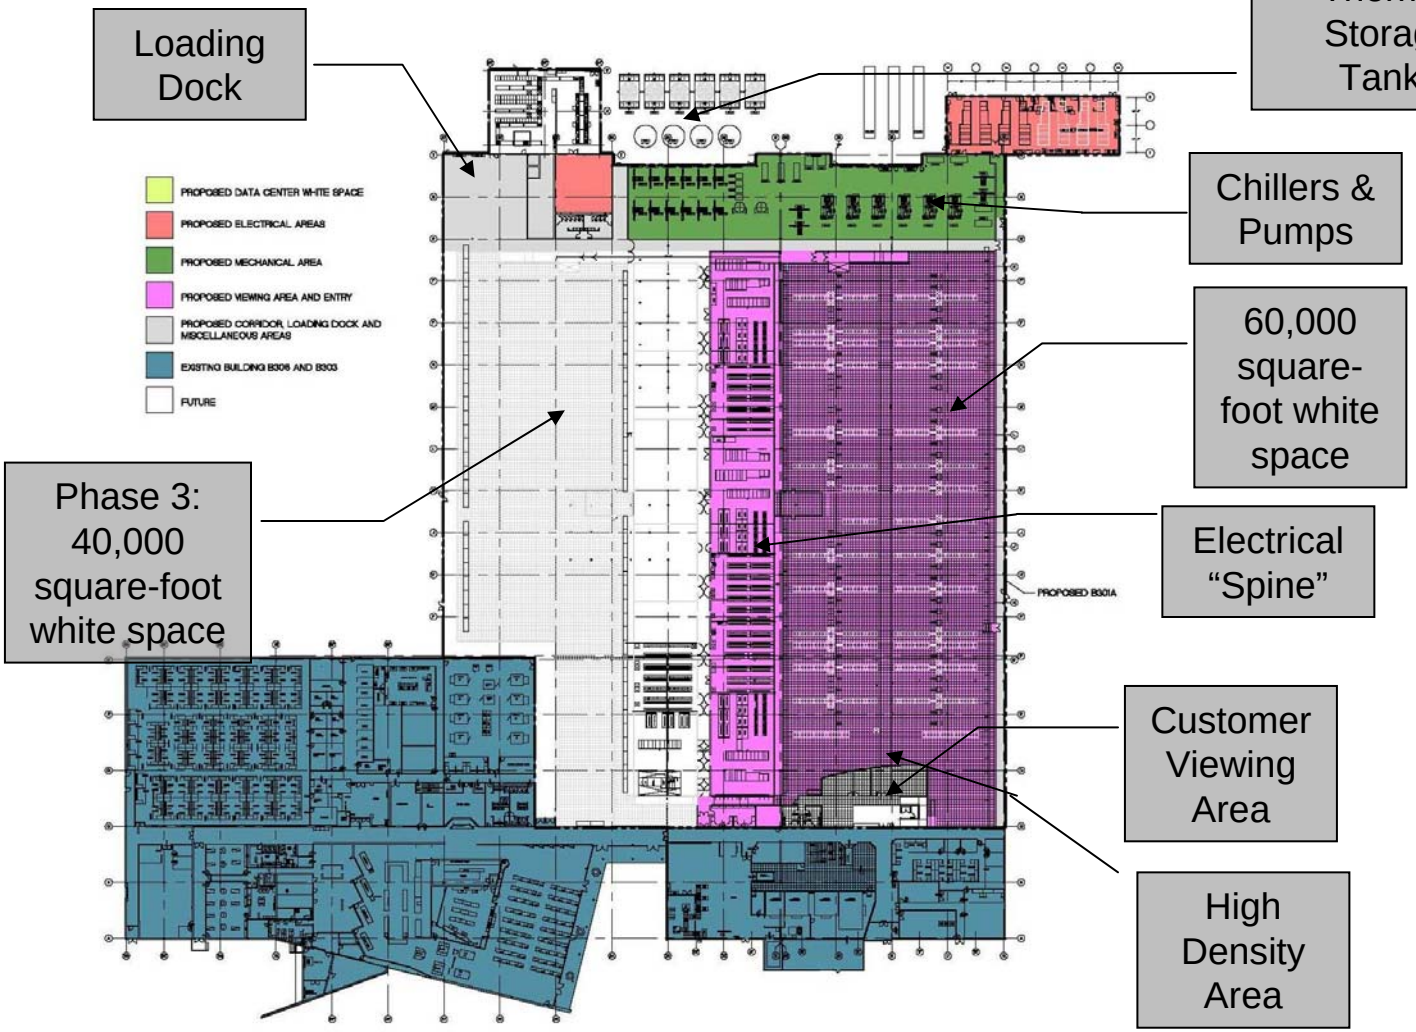

IBM Services Cloud Strategy

Enable our customers to leverage cloud computing through designing, building, and delivering

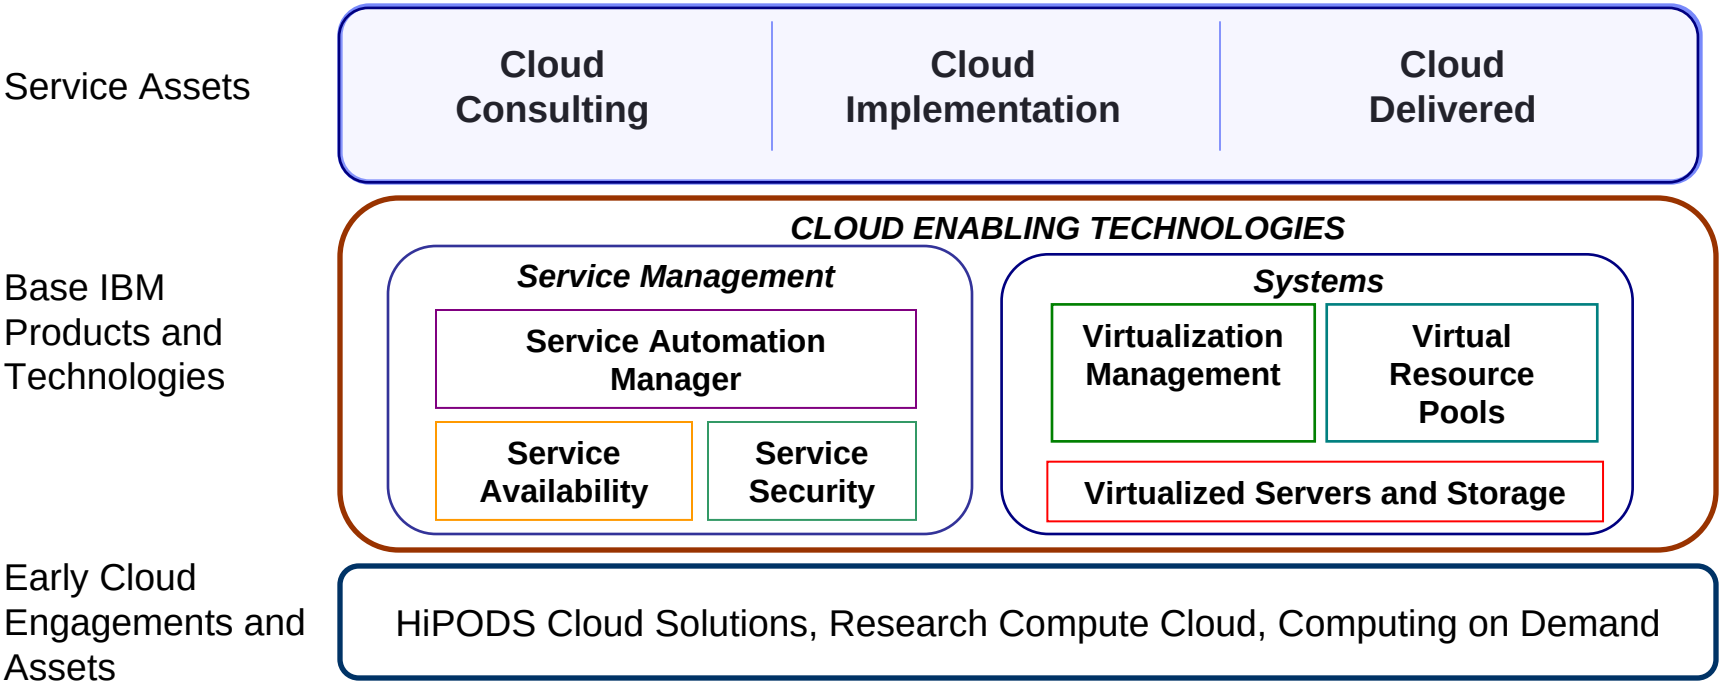

…bringing clarity and focus.

The journey to cloud…

Business Value

Centralize > **Consolidate** > **Virtualize** > **Automate** > **Optimize**





Organization Culture Governance

….requires an integrated and orchestrated approach.

IBM is Focused on Customer-Led Priorities

Cloud Consulting

- **How can cloud:**
  - **Improve responsiveness?**
  - **Save me money?**
  - **Still be secure and resilient?**

Cloud Implementation

- **How do I:**
  - **Get started?**
  - **Where?**
  - **Measure results?**

Cloud Delivered

- **What software?** collaboration, business process, applications
- **What platform?** middleware, database, systems management, tools?
- **What Infrastructure?** processing, storage, security

Cloud Enabling Products & Assets

- **What do I use to build my own cloud?**

Hardware, Middleware, Management & Billing, Tools, Services

IBM GTS Cloud offerings

Growing the portfolio to deliver industry leading cloud computing capabilities

Cloud Consulting

- Infrastructure Consulting Services for Cloud Computing
- Security and Resiliency Consulting Services for Cloud
- Resiliency Certification for Cloud Computing

Cloud Implementation

- Test and Developer Cloud Services
- Managed Security Services for Cloud Computing
- End User Cloud Services
  - Virtual Infrastructure Access
  - Self-Enablement Portal
- Scale out File Services

Cloud Delivered

- Information Protection Services
  - Remote Data protection
  - Managed Data Protection for desktops and laptops

Cloud Enabling Products & Assets

- Virtualization & Imaging technologies, Service Catalog, User Access Control
- Service Management and Automation software

Companies Waste $2.8 Billion Per Year Powering Unused PCs

"Unused PCs — computers that are powered on but not in use — are expected to emit approximately 20 million tons of CO2 this year, roughly equivalent to the impact of 4 million cars, according to report by 1E and the Alliance to Save Energy. All told, U.S. organizations will waste $2.8 billion to power 108 million unused machines this year. The notion that power used turning on PCs negates any benefits of turning them off has been discussed recently as one of five PC power myths. By turning off unused machines and practicing proper PC power management, companies stand to save more than $36 per desktop PC per year."

Developing the Cloud Strategy & Plan

Without a strategy, Cloud computing can be a threat to the CIO and IT team

- Reduced control of IT services delivered over the Internet
- Perceived cost gap between a cloud service and traditional IT

With a strategy, Cloud computing is a huge opportunity for the CIO

- Lower costs, more responsive IT, optimized delivery
- Greater range of services and capabilities
- Greater visibility in billing / chargeback to LOBs
- Better control of the users' systems, desktops, and services access

5 Steps to Cloud

1. IT Roadmap
2. Architecture
3. Workload Assessment
4. Enterprise / Cloud Mix
5. Implementation

Key questions to ask when exploring cloud computing

- Will cloud computing help create and deliver innovate business and consumer services to achieve greater competitive differentiation?
- Can cloud computing help to more quickly achieve goals for IT optimization, cost savings and faster time to market?
- Is a competitive advantage gained by using cloud computing?

What we've learned from…..

- There is nearly universal interest in Cloud computing, 20- 40% use Cloud computing today & nearly all are interested in using it in the future.
- Cloud inhibitors the same across the board – security, resiliency, and economics.
- Cloud has its sceptics – pay attention to stakeholder management
- To maximize benefits - Don't just automate, optimize
- This does achieve cost savings – really. You can cost justify the investment
- This is a paradigm shift in how the business thinks about acquiring IT resources
- Start with a pilot

… market insights through primary research – 5 countries, 650+ respondents, customers and experience.

Discussion

IBM's new Infrastructure Strategy & Planning for cloud computing service can help you develop a strategy, architecture, and roadmap

Features:

- Business and IT executive workshop to identify if cloud computing can drive business value.
- ROI value case and diagnostic assessment of the current environment to determine strengths, gaps and readiness.
- Strategy, plan, architecture, and roadmap to successfully implement the selected cloud delivery model.

Customer Benefits:

- **Lower Cost** – Identify opportunities to reduce capital and operating expense across the infrastructure.
- **Improve Service** – Streamline processes and services, improve efficiency and effectiveness**.**
- **Reduce Risk** – Architect a secure and resilient model that mitigates operational exposures and protects data.

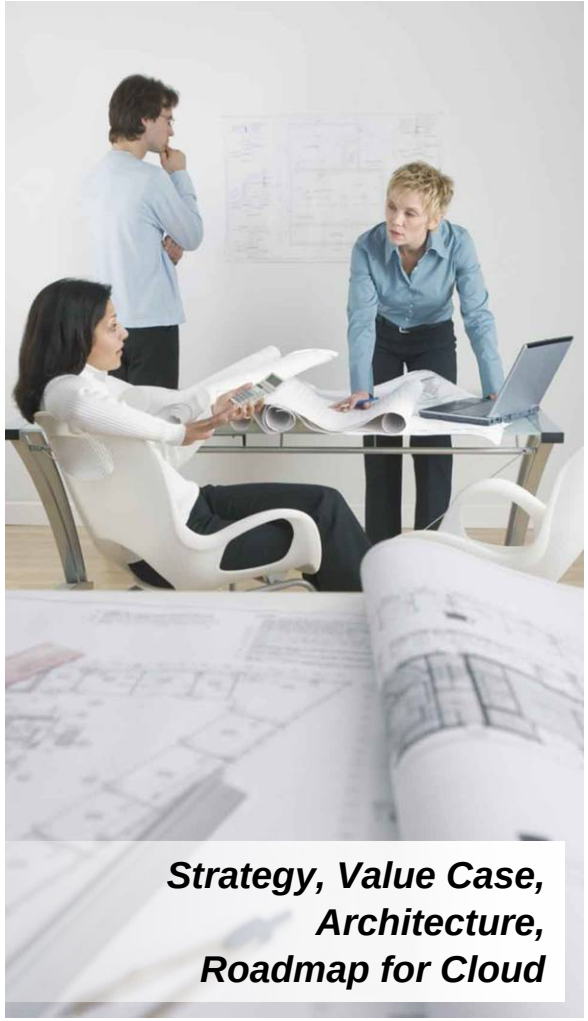

High-level cloud security concerns

Less Control

Many companies and governments are **uncomfortable** with the idea of their information located on **systems they do not control**. Providers must offer a high degree of security transparency to help put customers at ease.

Data Security

Migrating workloads to a **shared** network and compute **infrastructure** increases the potential for **unauthorized exposure**. Authentication and access technologies become increasingly important.

Reliability

High availability will be a key concern. IT departments will worry about a **loss of service** should outages occur. Mission critical applications may not run in the cloud without strong availability guarantees.

Compliance

Complying with SOX, HIPPA and other **regulations may prohibit** the use of clouds for some applications. Comprehensive auditing capabilities are essential.

Security Management

Providers must supply easy, visual controls to **manage firewall and security settings** for applications and runtime environments in the cloud.

The IBM Security Framework

Comprehensive Risk and Compliance Management

- **15,000** researchers, developers, and SMEs on security initiatives
- **3000+** security & risk management patents
- **200+** security customer references and **50+** published case studies
- **40+** years of proven success securing the zSeries environment
- **$1.5 Billion** security spend in 2008
- Managing more than **4 Billion** security events per day for clients

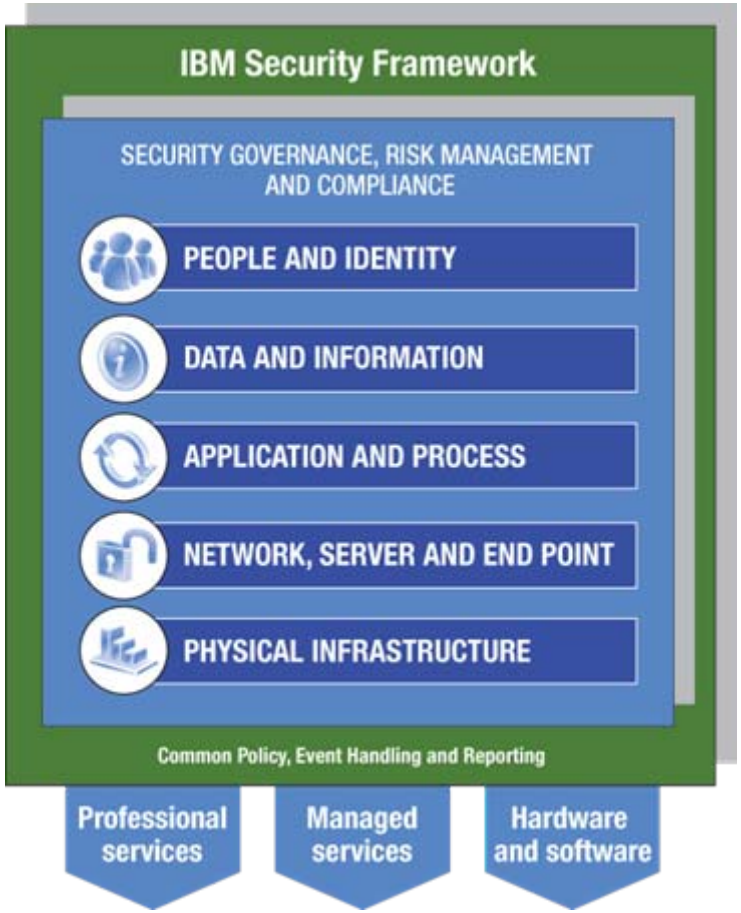

Gartner's security risks of cloud computing

…map directly to the IBM Security Framework.

- Privileged User Access
- Data Segregation
- Data Recovery
- Investigative Support
- Regulatory Compliance
- Data Location
- Disaster Recovery

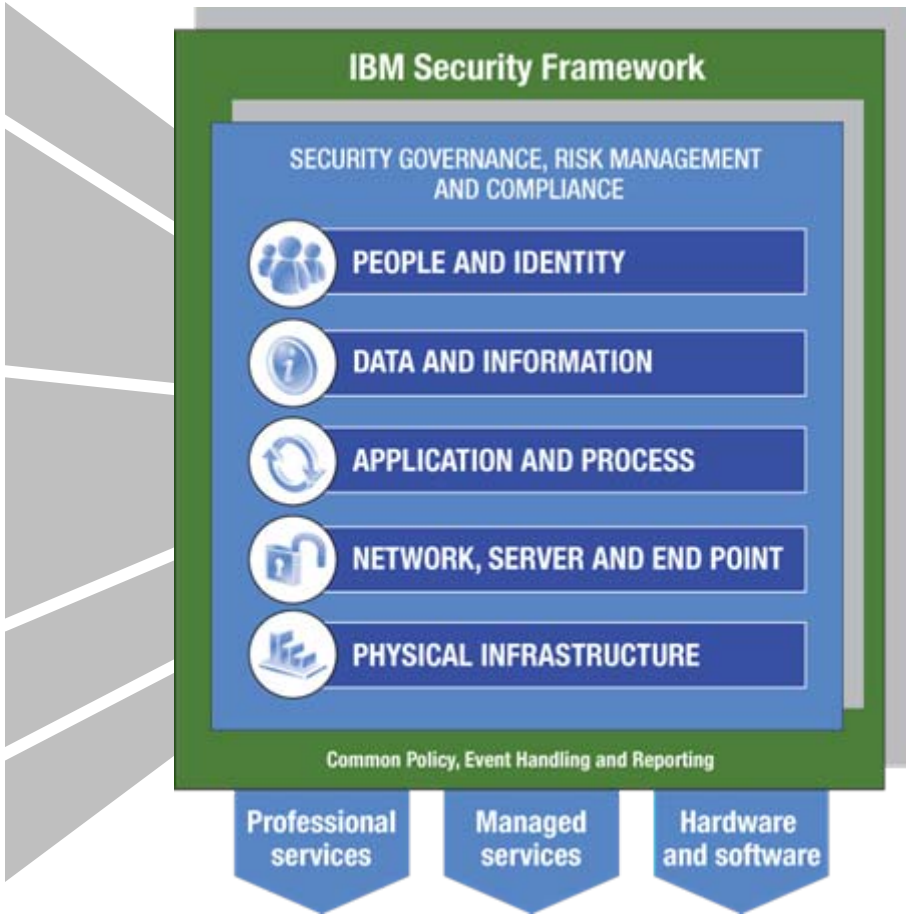

Gartner: Assessing the Security Risks of Cloud Computing, June 2008

IBM Security Consulting Services

Building secure cloud environments with IBM security expertise

- **Strategy and Planning - Helps deliver a comprehensive, detailed assessment of your business and IT security risks**
- **Vulnerability assessment and security architecture**
  - Automated scans to identify vulnerabilities
  - Documentation of gaps in key security controls
  - Policy design and definitions for secure controls
- **Security Roadmap - Prioritized recommendations against business goals for security best practice improvements that help mitigate business risks**

CLIENT BENEFITS

- **Identify and mitigate security exposures with a comprehensive assessment of your security strengths and weaknesses and streamline on-going management**
- **Enable business-aligned security controls to help avoid fine, pass audits, reduce litigation, and manage regulatory compliance**
- **Help reduce costs and complexity of security management and ease staffing pressures**
- **Leverage existing IT infrastructure to protect current IT investments and benefit from existing technology**
- **Simplify protection of your valuable, business-critical and/or confidential data**
- **Build adaptable security infrastructures and implement security best practices**

Security Implementation Services for Cloud Computing

Business challenge: *Security Optimization*

- Reduce operational expenses in managing their security posture without requiring additional physical infrastructure expenses.
- Implement a flexible security foundation to accommodate change in today's highly dynamic environments.
- Augment their IT staff with security expertise to ensure 24 x 7 protection.

Solution: *Provides Security Visibility and Control*

- Lower the total cost of ownership related to enterprise security.
- Seamless platform for visibility and control over vulnerabilities and security risks across enterprise.
- Scalable solutions offering analysis and reporting of across heterogeneous and virtual devices.
- Archive solutions for forensics and compliance reporting.
- Integrated with IBM X-Force security intelligence and backed by security assurance guarantees.

IBM ISS Security Offerings For Cloud

- Security Event and Log Management.
- Vulnerability Management Services.
- X-Force Threat Analysis Service.
- Email and Web Security Services.

Hbc Great things for Canada™

Customer Example:

Hudson's Bay Company (HBC) is required to comply with the credit card industry's data security standard (PCI-DSS) to ensure regulatory compliance and protect their brand with their current infrastructure and staff.

HBC depends on IBM Internet Security Systems to monitor the security event logs of IT assets and report through the centralized, on-line MSS portal - helping the organization meet regulatory requirements without additional infrastructure expenses.

BCRS Cloud Computing and Resiliency Validation Services

Part 1: Technical Cloud Architecture *Evaluation*

- Leveraging IBM's defined Cloud Architecture Standards, IBM BCRS consultants will evaluate a Cloud's architecture against the IBM Cloud Architecture Standards for Resiliency, Availability, Continuity/Recoverability, Scalability, and Security.
- They will identify, quantify, and prioritize gaps and risks, and then provide the ongoing design assistance and management expertise to establish a more resilient infrastructure.

Part 2: Service Delivery Resilience *Verification*

- IBM BCRS consultants will assess the ability of the Cloud environment and supporting organization to deliver IBM defined levels of service to customers.
- The assessment and testing will span the six layers of the resilience framework:
  (1) Strategy
  (2) Organization;
  (3) Processes,
  (4) Technology
  (5) Applications and Data
  (6) Facilities
  to validate the Cloud's Business Continuity and resilience capabilities.

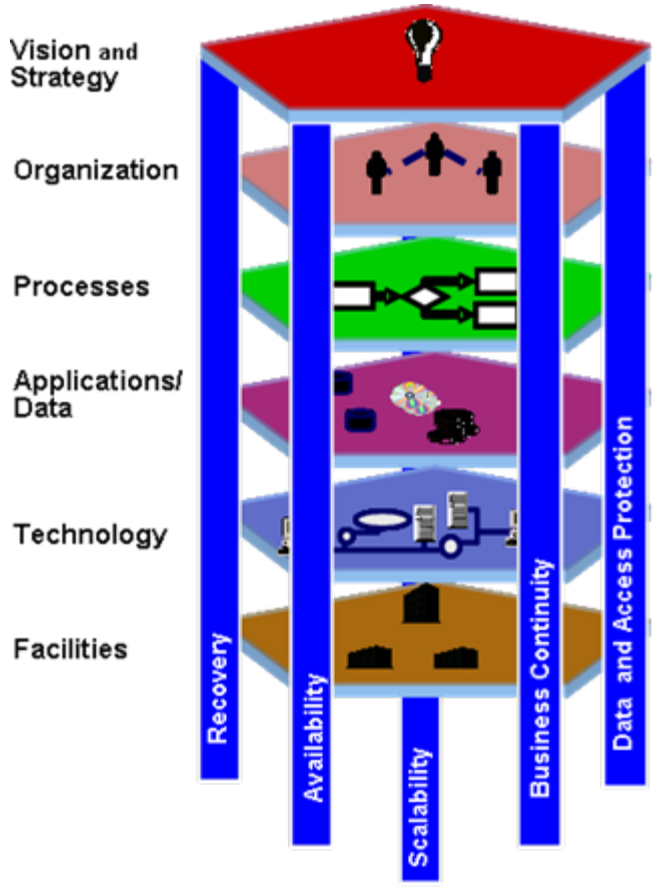

Framework for Resilient Architecture

BCRS Resiliency Validation Service Approach

Use the Framework for Resilient Architecture in conjunction with the Blue Cloud Architecture Summary to ***validate*** resiliency of cloud service providers. Clients who go through the program will be able to use "Resiliency Proven" logo

IBM Resiliency Consulting Services – Assessment and Planning - Resilient Cloud Validation

RCV Ensures Resiliency

- BCRS has over 40 years of experience of making clients infrastructures resilient and is now applying that knowledge to Cloud.
- RCV validates the resiliency of any company delivering applications or services through public or private cloud environments, and its ability to sustain those resilience capabilities over time.
- Two step process includes validation of documented architecture and hands on testing to validate resiliency

Customer Example:

A large defense contractor client was interested in the Resilient Cloud Validation program for several reasons. Their primary goal was to combat the very real fear in the industry that clouds should only run applications that you can actually do without for a couple of days.

Since their end-user clients are all DOD, they needed to be able to prove to them that the proposed Cloud services could provide the same level of reliability that traditional data center centric service offerings provide.

The Resilient Cloud Validation program demonstrates to their clients and to the press that they are running a true cloud service that can be trusted.

IBM Managed Resiliency Services - Information Protection Services – onsite and remote data protection

IPS Provides Reliable Data Protection

- Automatically backed up via your existing network through the cloud to our security-rich, offsite data centers or onsite to your own data center.
- Fully managed solution can reduce backup costs by 20-40%.
- Skilled IBM storage specialists worldwide who provide 24x7 monitoring and management.
- Quickly implement a best practices–based data protection strategy.

Customer Example:

One of Houston's largest and fastest-growing human services agencies

Serves over 200,000 citizens in Texas.

Depends on IBM cloud services to back-up server and PC data

Tom Comella, chief information officer, Neighborhood Centers Inc. *"IBM cloud services were critical in our community recovery efforts following Hurricane Ike. The benefits of cloud services reach far beyond disaster recovery. Better data protection -- demonstrating that we are good stewards of information -- has become a selling point for us in willing contracts."*

Cloud delivered data protection to your remote desktops & laptops

Information Protection Services – managed data protection for desktops / laptops

- Secure, scalable, automatic data backup solution for laptop / desktops
  - Backs up data locally and to a remote vault
- Predictable low monthly fee for fully managed service
  - Leveraging IBM's Tivoli Continuous Data Protection for Files software

Available from IBM and IBM Business Partners, April 2009!

Why should you implement this solution?

- Eliminate data loss by providing automated data protection and recovery capabilities to make your mobile workforce more resilient
- Reduce TCO 40%, no capital expense

IBM Managed Resiliency Services - Information Protection Services - email management express (EMX)

EMX Ensures Business Continuity

- Comprehensive, cloud-based standby email system can be activated in less than 60 seconds to ensure email continuity, including wireless devices.
- Email is available through any web browser.
- Recovers and restores email messages to the primary email system after outage is resolved.
- Archives email using customer-defined policies
- Can be deployed within a day with no upfront costs.
- Reduces email data stores by as much as 80%

Customer Example:

A large national hotel chain with more than 50,000 email users in the United States has used our cloud-based Email Management Express service for more than five years to ensure that their employees always have access to email.

Continuous access to email is critical to this client, as they are open 24 x 7 and information is constantly being transferred.

This client has also adopted our crisis notification to ensure they can communicate with their employees in the event of an emergency.

IBM Implementation Services for cloud computing – design and implementation for test environments

Features:

- Assessment of current test environment to project savings and ROI.
- Strategy, planning, design and implementation services of the solution.
- Create self-service portal with catalog of services.
- Integrated platform combining service request management, provisioning / de-provisioning and change and configuration management.

Customer Benefits:

- **Reduce IT labor cost by 50% +**
  reduce labor for configuration, operations, management and monitoring of the test environment.
- **75% + Capital utilization improvement**
  significant license cost reduction.
- **Reduce Test Provisioning cycle times from weeks to minutes.**
- **Improve Quality**
  eliminate 30% + of all defects that come from faulty configurations.

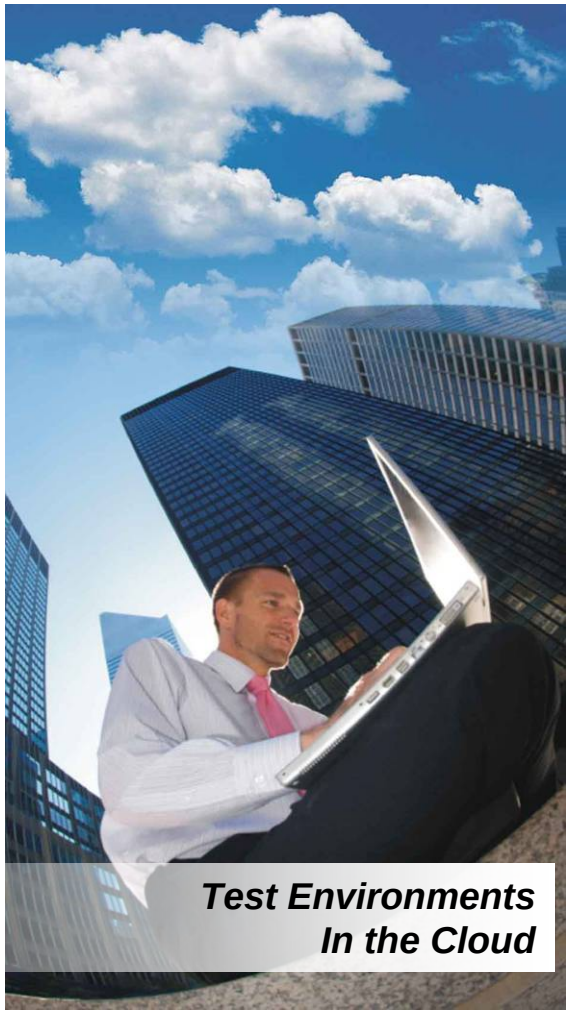

IBM Self Enablement Portal

Business challenge: *Effective Support*

- Control, manage and reduce support costs to lower the total cost of the service desk.
- Enable IT support to increase end user productivity.
- Provide 7x24 support to a mobile and geographically dispersed workforce.
- Leverage technology to continually improve the efficiency of IT support.
- Enable IT staff to focus on innovative activities that add revenue versus support that only adds costs.

Solution: *SEP Provides SaaS Style Self Service*

- Fully functioning self-help portal with web chat and remote assistance features.
- Plugs into an existing help desk with IBM managing all the execution of the portal.
- Rented solution that transforms help desks to lower cost and higher value capability.

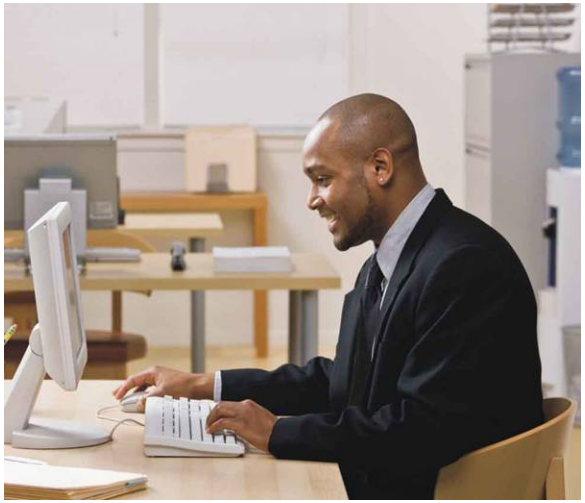

Customer Example:

- Large financial services business: $4M savings over 10-month period.
- Averages 30% chat, 40% end user issues handled without help desk call.
- 75% of portal visits avoided a call to help desk.
- Investments in the tens of millions were spent to produce these results – with SEP, we repeat over and over with a single global investment.

IBM's Virtual Infrastructure Access Services enables a virtualized IT access environment

Characterized by thin clients, higher levels of security and resilience, improved backup and recovery…and reduced cost and complexity associated with managing the "close to the box" user environment…

Integrating hardware, software and services in an innovative solution

- Using a single consistent framework connecting authorized end users
- Enabling thin clients or any other Internet-connected device to access platform-independent, hosted applications, and full client images…anytime, from any device through a Java-enabled browser
- Using open framework architecture, leveraging current top vendors and technology partners in the market today

Available as two different offerings:

- Project-based services
- Managed services

According to one major industry consultant… a successful deployment of virtual clients can render savings in excess of 40% …versus an unmanaged PC environment…and as much as 35% versus the typically managed environment.

Scale out File Services enables dedicated storage cloud computing

Business challenge: *Storage and Data Optimization*

- Speed deployment of storage applications
- Allow storage infrastructure to adapt dynamically to business requirements
- Improve cross-company collaboration and file/data sharing while controlling capital expenditures

Solution: *Provides Scalable and Integrated Storage*

- Lower the total cost of ownership related to enterprise storage, including:
  - Security authentication and encryption
  - Scalability
  - Archiving and indexing automation
  - Long-term retention

Why IBM?

- Built on unique models and technology originally developed jointly in collaboration between IBM's own CIO and IBM Research
- Leading global skills and proven storage cloud computing environments for mission-critical business applications

Customer Example:
A telecommunications firm needed a highly scalable and cost-effective environment for providing an innovative new online media and content services to their users, allowing them to share photos, video and audio including unlimited media storage and unlimited uploads.

Start small and grow rapidly - the initial client solution consisted of 50 servers, 3TB of SAN storage, networking and 80TB of Scale out File Services storage. The client expects the SOFS dedicated storage cloud component to grow to in excess of 1 Petabyte within a short period and the number of servers to potentially grow to hundreds. This telecommunications firm found that SoFS dedicated storage cloud services offer a predictable cost model for planning and growth.

Kantana Animation Studios Co., Ltd.

Solving the Data Storage Challenge

Business Challenge:

Kantana Animation Studios found that its character rendering and modeling processes were putting increasingly heavy demands on its data storage infrastructure. Needing to store and retrieve extremely large files at high speed, the studio looked for a cost-effective solution that could handle exceptional growth.

Solution:

Kantana implemented IBM Scale-out File Services (SoFS), an all-in-one data storage solution with IBM System x™ and IBM System Storage™ technologies and management services from IBM Global Technology Services. The solution allows Kantana to store large files in a single logical location, accessible by all animators. Storage capacity may be increased precisely as the business requirements dictate.

Benefits:

- Increased productivity with centralized file storage for all animators
- Enhanced, cost-effective scalability meets growing business requirements
- Reduced administrative workload and costs with IBM Scale-out File Services

Kantana Animation Studios

"The IBM Scale-out File Services solution helps our animators to co-operate by enabling rapid and easy access to shared projects."

Auchara Kijkanjanas
Managing Director of
Kantana Animation Studios Co., Ltd

Solution Components:

- IBM System x3650
- IBM System Storage DS4800
- IBM System Storage TS3310 Tape Library
- IBM General Parallel File System Services
- IBM Global Technology Services
- IBM Scale-out File Services

Katana video on SoFS:

TSP03079-USEN-00

Getting started…

Develop a strategy ▶ ***Best practices*** ▶ ***… think holistically***

Consolidate ▶ ***Reduce from many to few*** ▶ ***… start with an inventory***

Virtualize ▶ ***Assess and deploy*** ▶ ***… start now***

Manage ▶ ***Gain and maintain control*** ▶ ***… modularity and standards are key***

What questions to ask to determine if Cloud is a good fit for you?

Key Pain Points

- Lost business opportunity because IT too slow to react. Lack of agility.
- Long deployment timelines for new systems (weeks/months+).
- Many people involved in the process, high cost & complexity.
- Many steps are manual and prone to error.
- Huge up front investment for new infrastructure when I want to start small.
- Server Sprawl
- Low Utilization
- Compliance, auditing, and security patching costly.
- Don’t know what compute resources are used or how much they cost?

Key Questions to ask?

- How quickly can you react to deliver a new IT service?
- How many steps are in the provisioning process?
- What is the ratio of system admins to servers?
- Have you experienced outages due to human error ?
- How are systems sized and scaled quickly (peak usage, CUOD)?
- How many images per user?
- Am I sized for min, mean, or peak ?
- How many different configurations used?
- What level of metering and method of charging used? How do we manage license compliance ?

What is Cloud Computing?

A user experience and a business model

- Cloud computing is an emerging style of IT delivery in which applications, data, and IT resources are **rapidly provisioned** and provided as **standardized offerings** to users over the web in a **flexible pricing model.**

An infrastructure management and services delivery methodology

- Cloud computing is a way of **managing** large numbers of highly **virtualized resources** such that, from a management perspective, they resemble a single large resource. This can then be used to deliver services with **elastic scaling**.

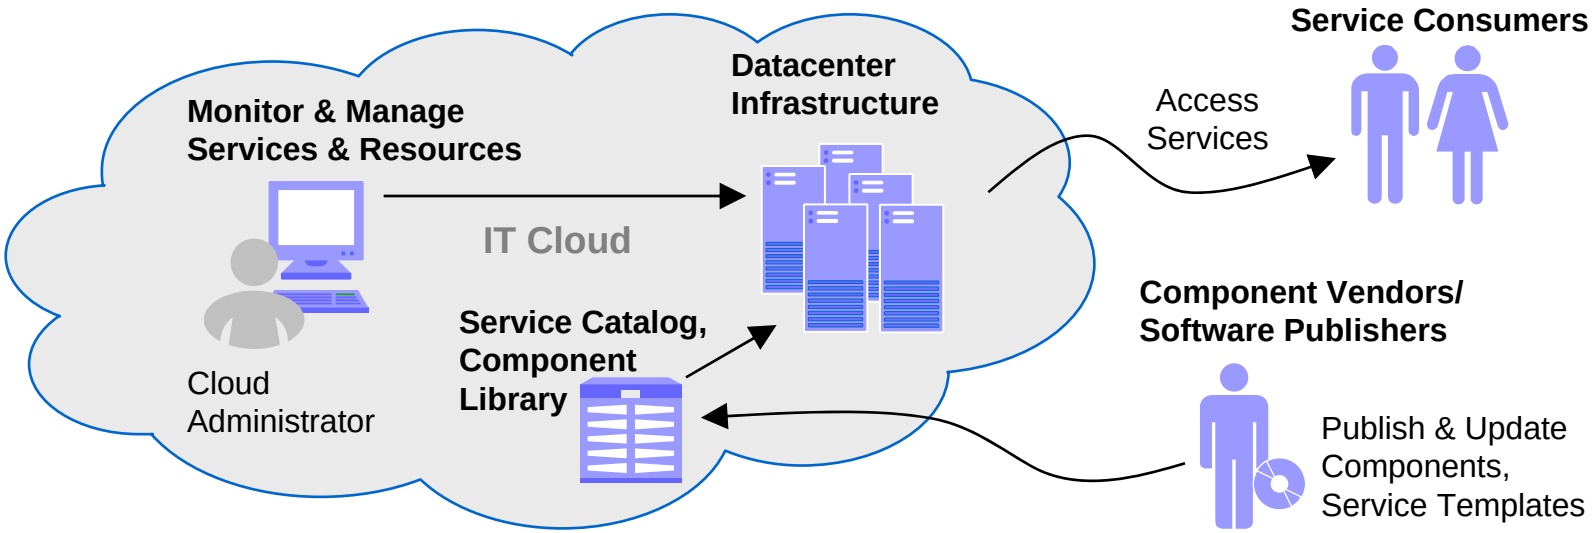

IBM Cloud Computing Gaining Momentum

2007 4Q	February	March/April	2008 May	June

Commercial Use

- IBM Blue Cloud Announced
- IBM Capacity Planning on Cloud Available
- Wuxi China Cloud Computing Center
- iTricity Deliver IT as a Service

IBM Cloud Centers

- IBM Cloud Computing Center Almaden
- IBM IDA IRELAND First Cloud Computing Center in Europe Opens
- IBM Cloud Computing Center Greater China South Africa

Collaboration and Innovation

- NORTH CAROLINA STATE UNIVERSITY NC STATE UNIVERSITY
- SOGETI Online Idea Brainstorm a "terrific success" at Dublin Cloud
- Vietnam Innovation Portal and SSME

Next Generation Internet Scale Computing Skills and Research

- IBM Google Academic Initiative
- IBM Joint research initiative with 13 European partners
- NSF IBM Google New advanced academic research
- facebook IBM Google Hadoop Workshop at IBM Dublin Cloud Computing Center

Why we need an Education Cloud in NC

March 27, 2009

Phil Emer

THE WILLIAM & IDA
FRIDAY INSTITUTE
FOR EDUCATIONAL INNOVATION

The Answer

- Outnumbered
- Cube != Classroom
- Virtual School
- ~~Stimulus~~ ARRA

By the Numbers

A **Small** District

12	Schools
334	Teachers
4,159	Students
1,896	Computers
6	Technical Staff

A **Medium** District

32	Schools
1,992	Teachers
28,881	Students
8,688	Computers
16	Technical Staff

A **Large** District

161	Schools
7,952	Teachers
127,404	Students
41,386	Computers
168	Technical Staff

THE WILLIAM & IDA
FRIDAY INSTITUTE
FOR EDUCATIONAL INNOVATION

Supporting this is one thing

THE WILLIAM & IDA
FRIDAY INSTITUTE
FOR EDUCATIONAL INNOVATION

Then there is supporting this

QuickTime™ and a
decompressor
are needed to see this picture.

21st century classrooms include projectors, audio gear, interactive boards, response systems, scientific probeware, and so-on

North Carolina Virtual Public School

- Over 25,000 high school students enrolled in NCVPS
- Expanding to include middle grades courses
- College credit courses
- Children's Internet Protection Act (CIPA) filters

NCVPS for some :-(

THE WILLIAM & IDA FRIDAY INSTITUTE FOR EDUCATIONAL INNOVATION

The Cloud increases in cost effectiveness as sharing increases

THE WILLIAM & IDA
FRIDAY INSTITUTE
FOR EDUCATIONAL INNOVATION

The Test at the End

- Technology support capabilities within districts cannot support “modernized” schools
- Schools need staff to support 21st century classrooms and teaching and learning with technology
- An education cloud provides a CIPA-friendly platform for virtual schools
- Cloud succeeds where grid fell short because having no money provides a compelling motivation for sharing

VCL Architecture – New Resource Type Add, Overview

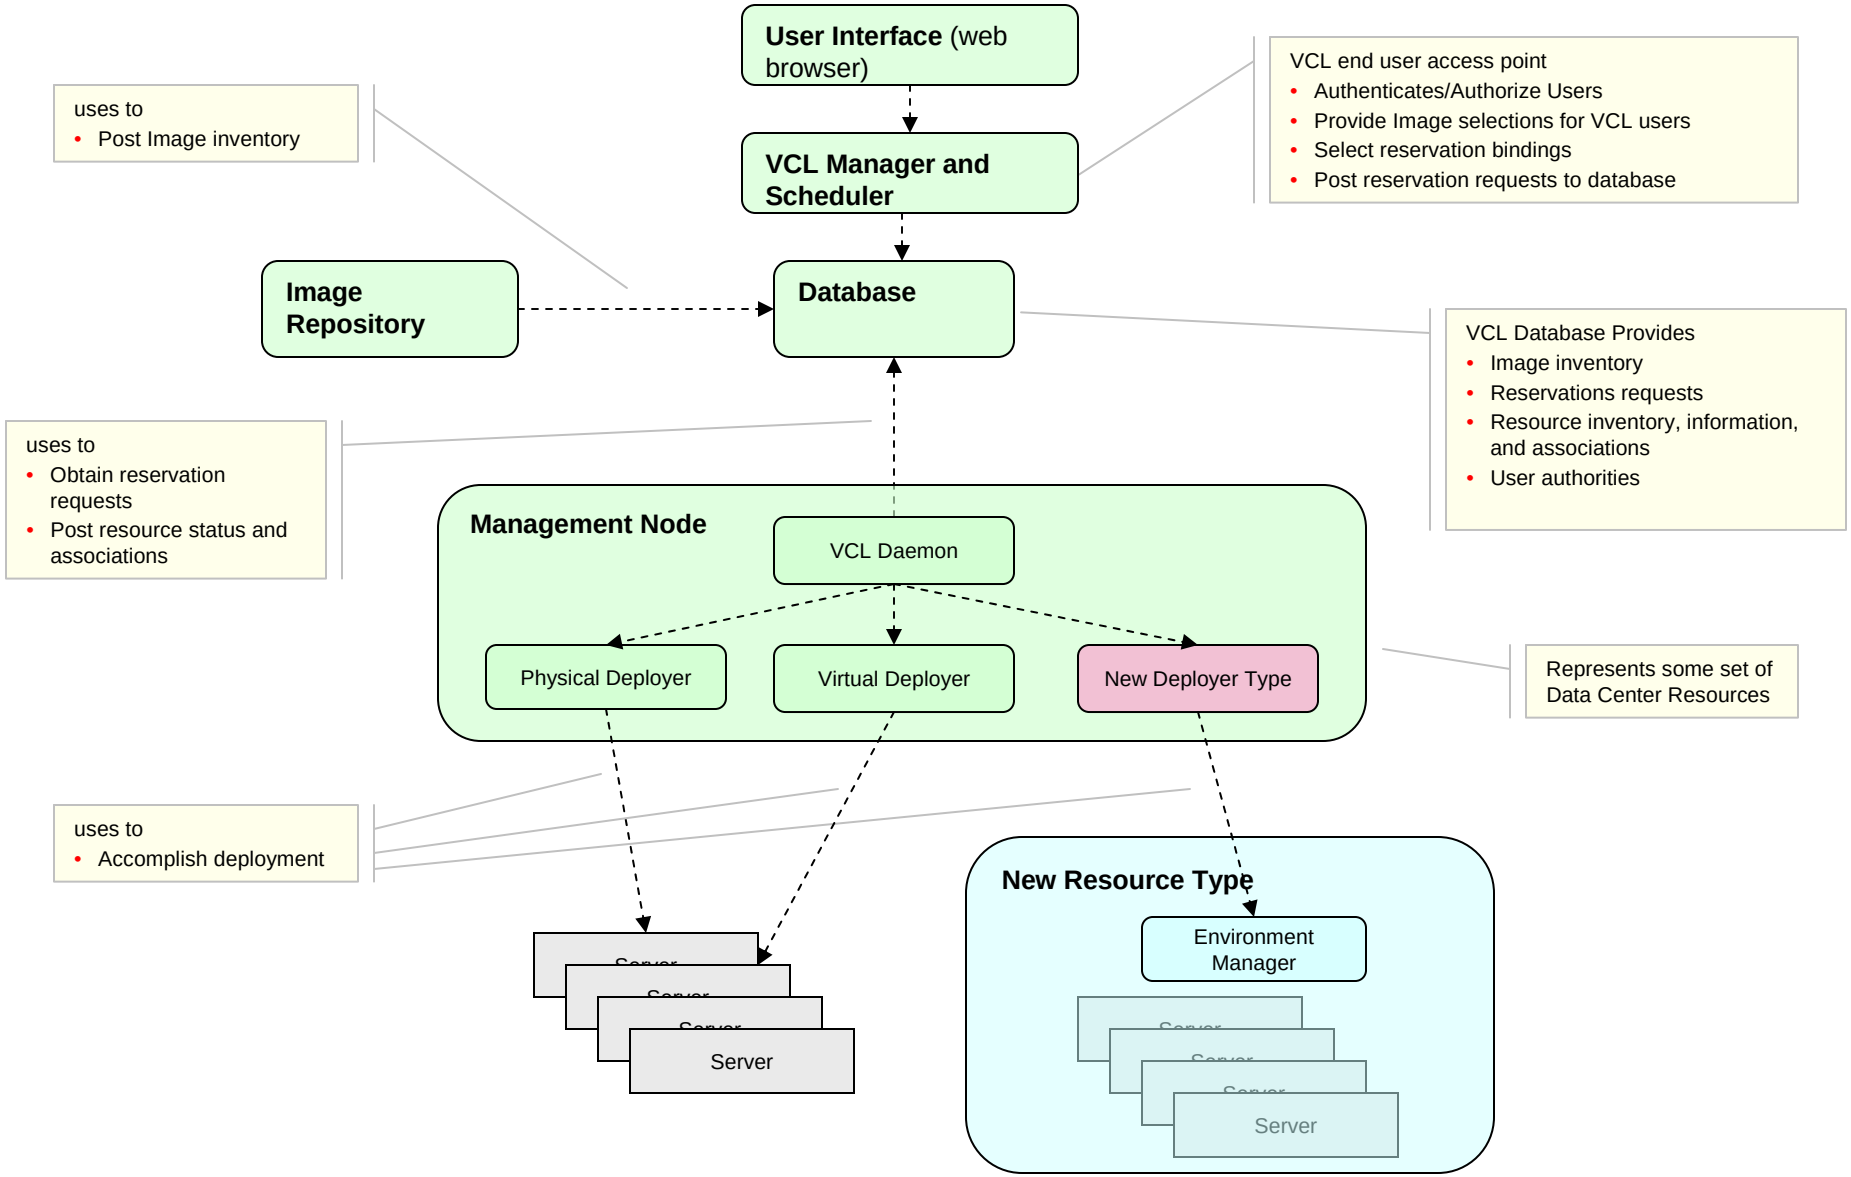

IBM Research

Cloud Computing Technologies: an IBM Research perspective

Vas Bala
vbala@us.ibm.com

IBM Research, New York

July 2008

Technology evolution

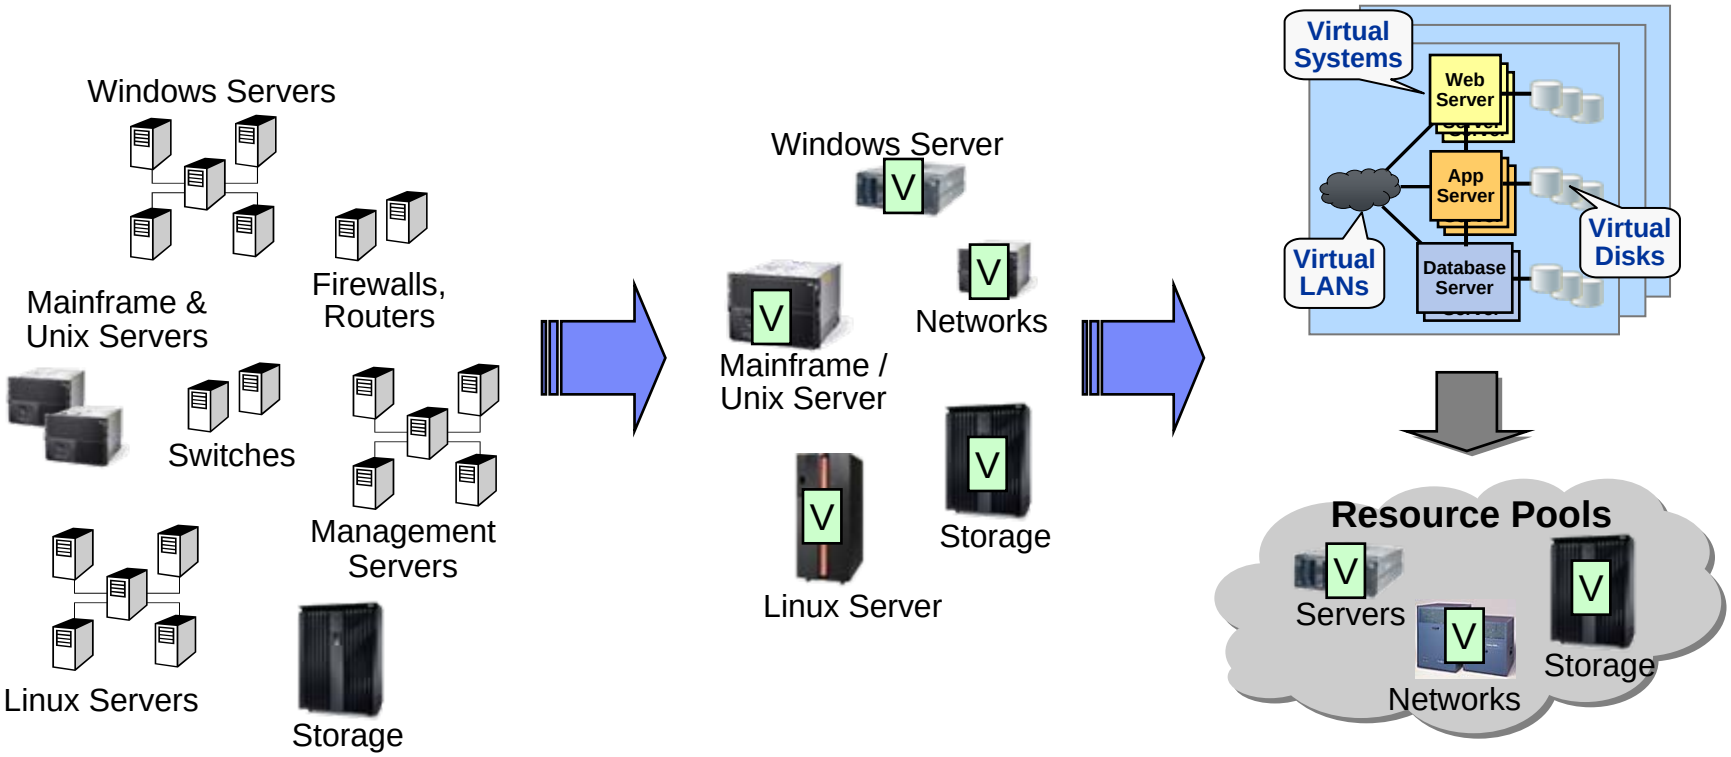

Scale-Out

- One workload per server
- Many physical constraints
- Mgmt cost prop. to # of systems

Physical Consolidation

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

Abstraction and Pooling

- Virtual resource objects
- Resource pools
- Reduced management costs

Layers of services

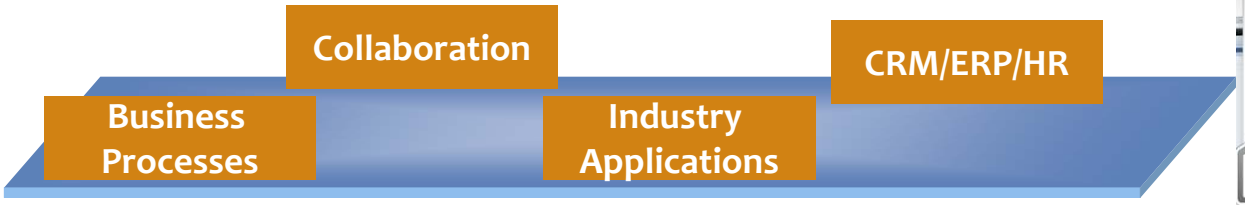

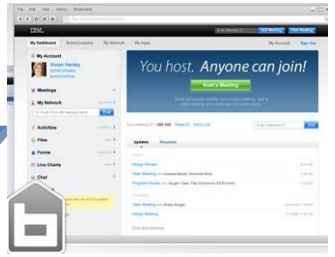

Software as a Service

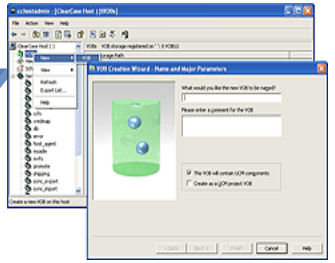

Platform as a Service

Infrastructure as a Service

IBM Research Compute Cloud (RC2): a living lab

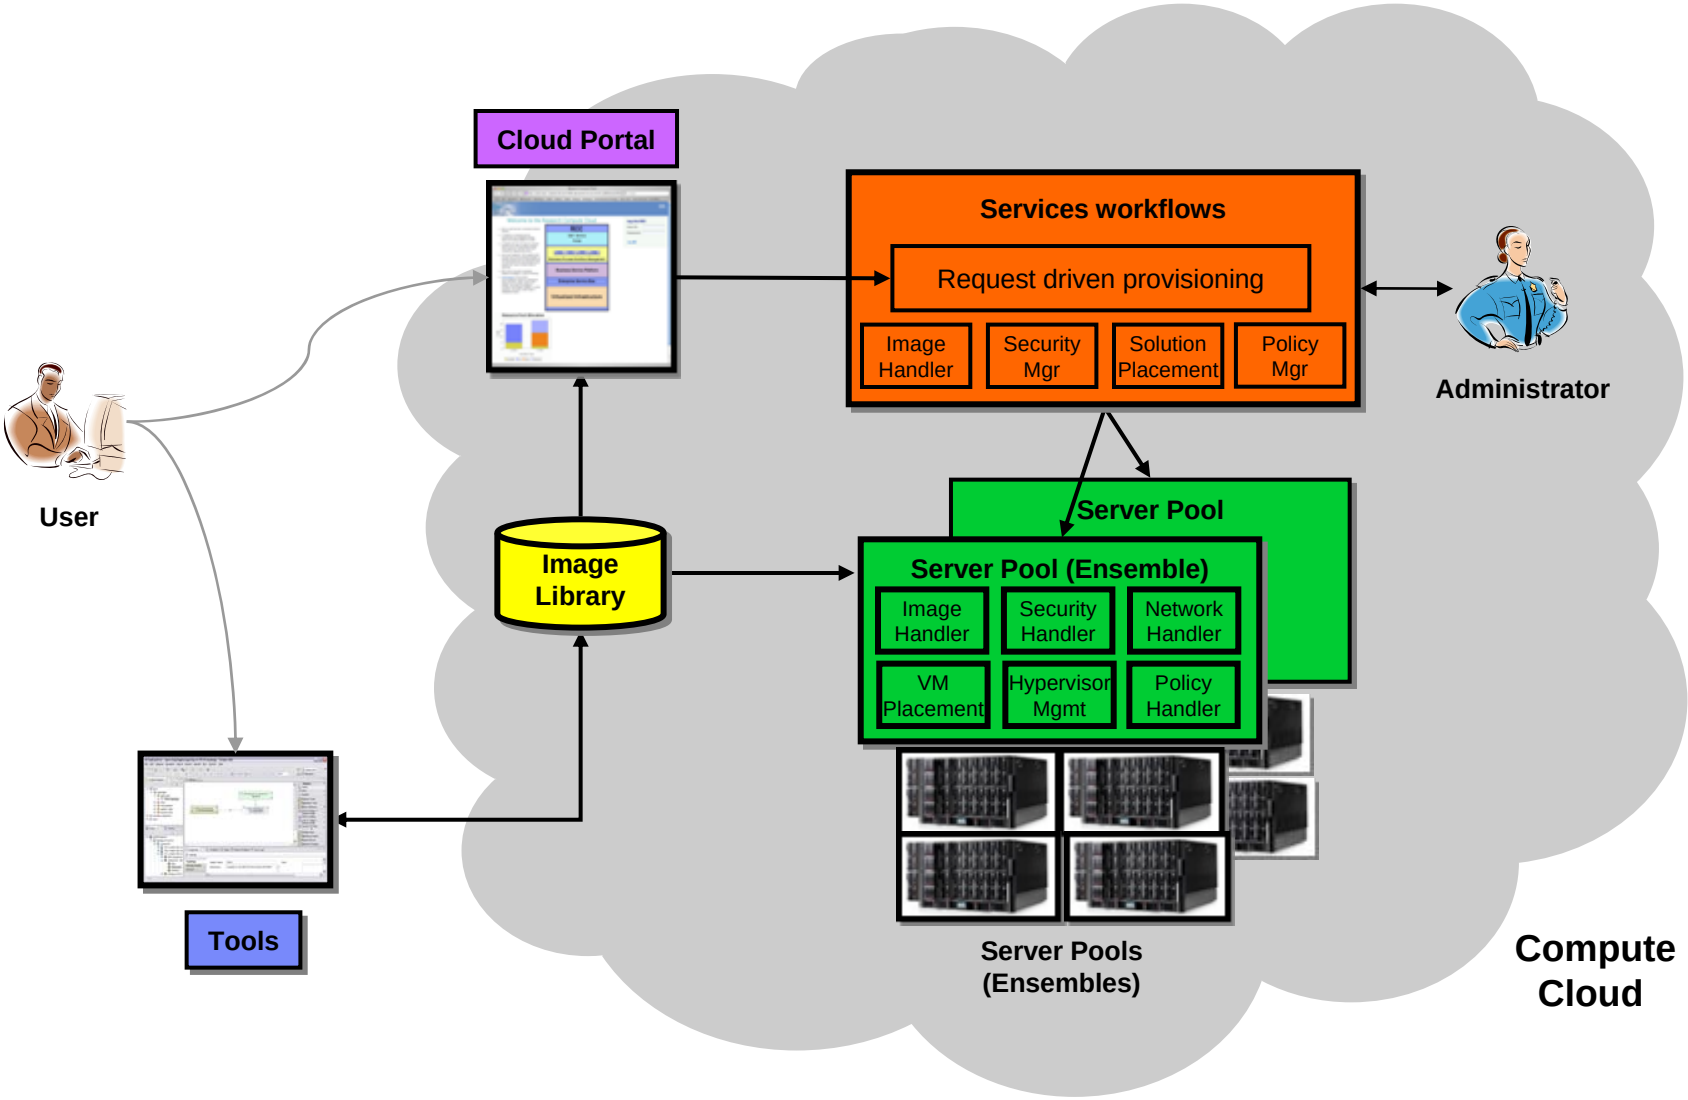

Research areas

- **Images**
  - Can VM images become durable, searchable, updateable objects?
  - Value: shift software maintenance tasks offline.
- **Tools**
  - Can complex solutions be composed using VM images as building blocks?
  - Value: shift configuration complexity offline.
- **Ensembles**
  - Can self-managing servers / server-pools be used as datacenter building blocks?
  - Value: scalability without increasing operation/mgmt costs.
- **Security**
  - Can security/isolation enforcement span hypervisor boundaries?
  - Value: security becomes a datacenter-wide concept.
- **Energy**
  - Can server utilization be optimized for power on a datacenter-wide scale?
  - Value: energy optimization become a datacenter-wide concept.

This is a subset of IBM Research activities in virtualization & cloud computing

Virtual machine sprawl, and its consequences

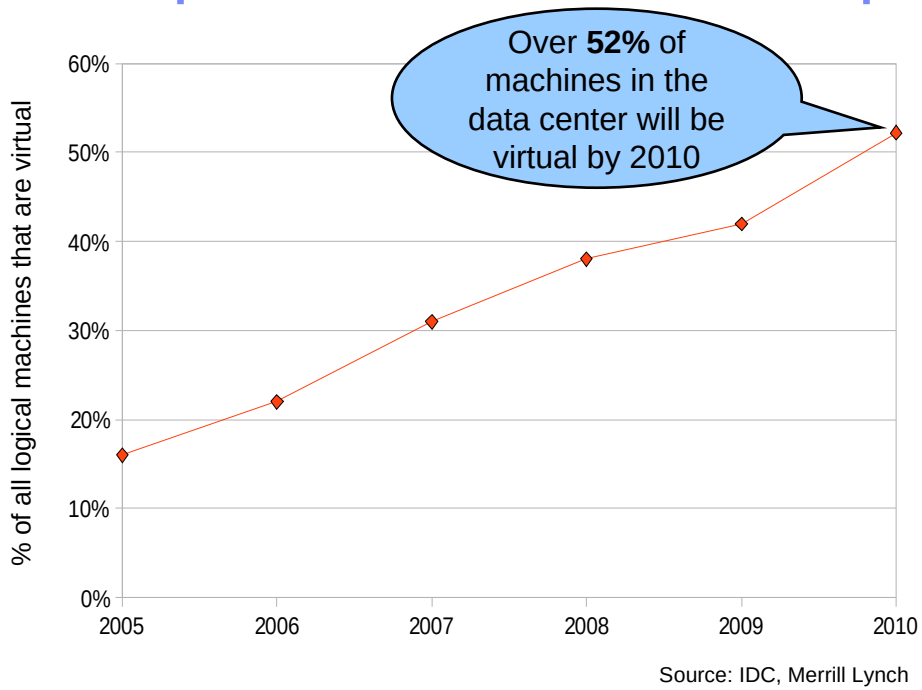

Source: IDC, Merrill Lynch

- **Which datacenter model scales better?**
  - Approach A: a few large machines/server-pools, each of which hosts lots of VMs
  - Approach B: lots of small machines, each of which hosts a few VMs.
- **How to efficiently over-commit hardware resources?**
  - Multiple degrees of freedom: capacity, lease period, avail & perf requirements, etc.
- **How to manage VM image proliferation?**
  - Should dormant images be managed the same way as running instances?

Innovations in image management

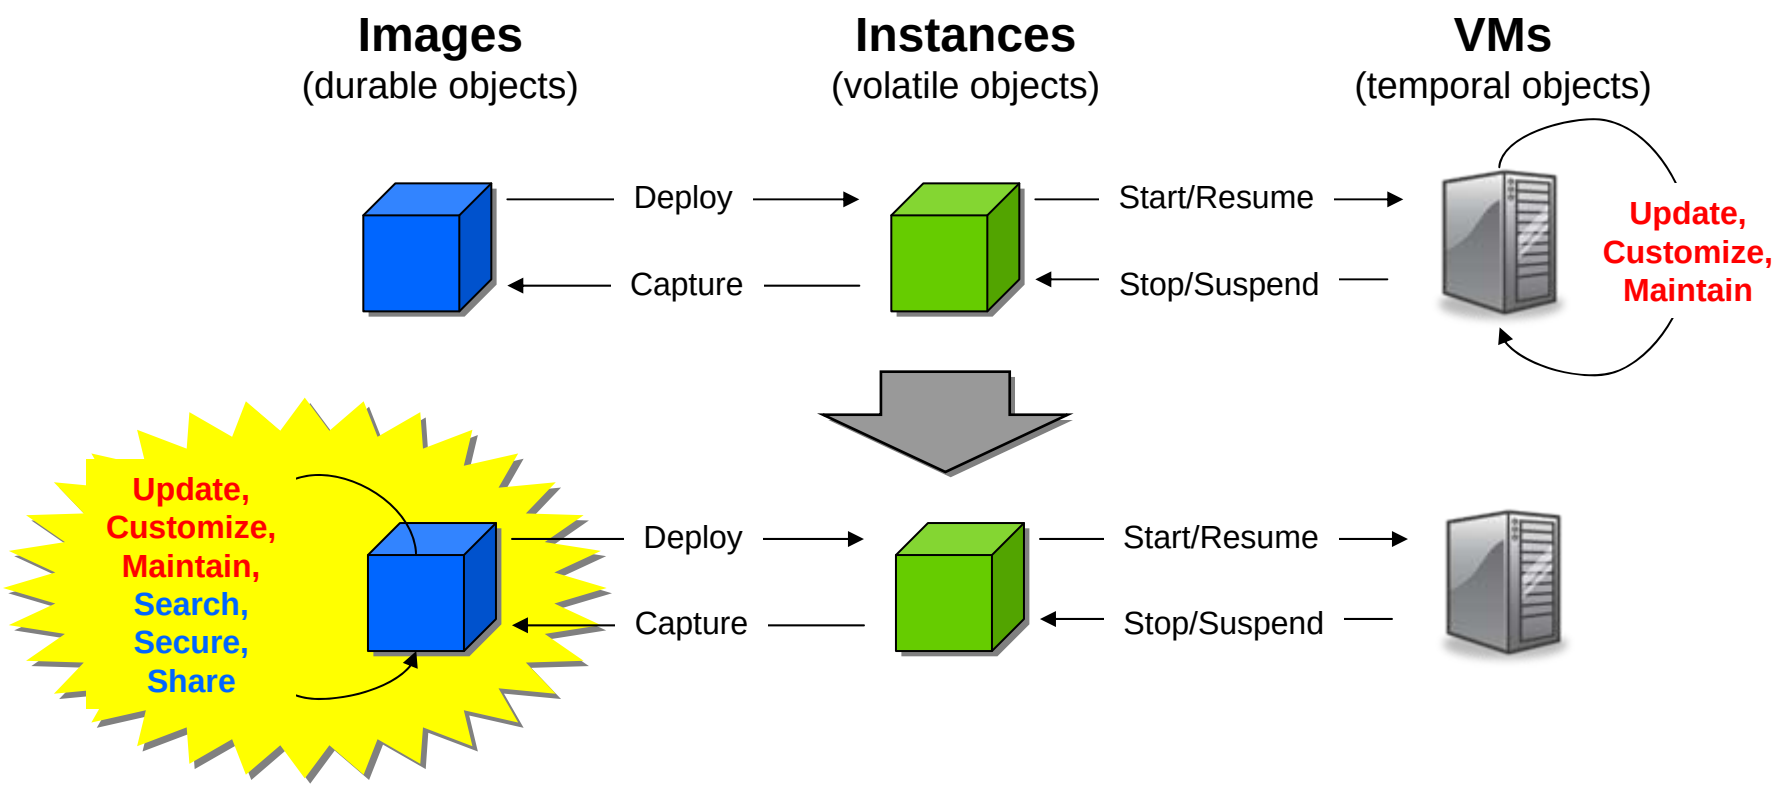

Benefits:

Scalability - don't have to start up every image in order to do maintenance on it.

Security - provenance tracking, security patching, etc without running the images.

Robustness - less noise from unintended side-effects of VM execution.

Sharing - easier to share images across a community of users.

Image sharing => faster time to value

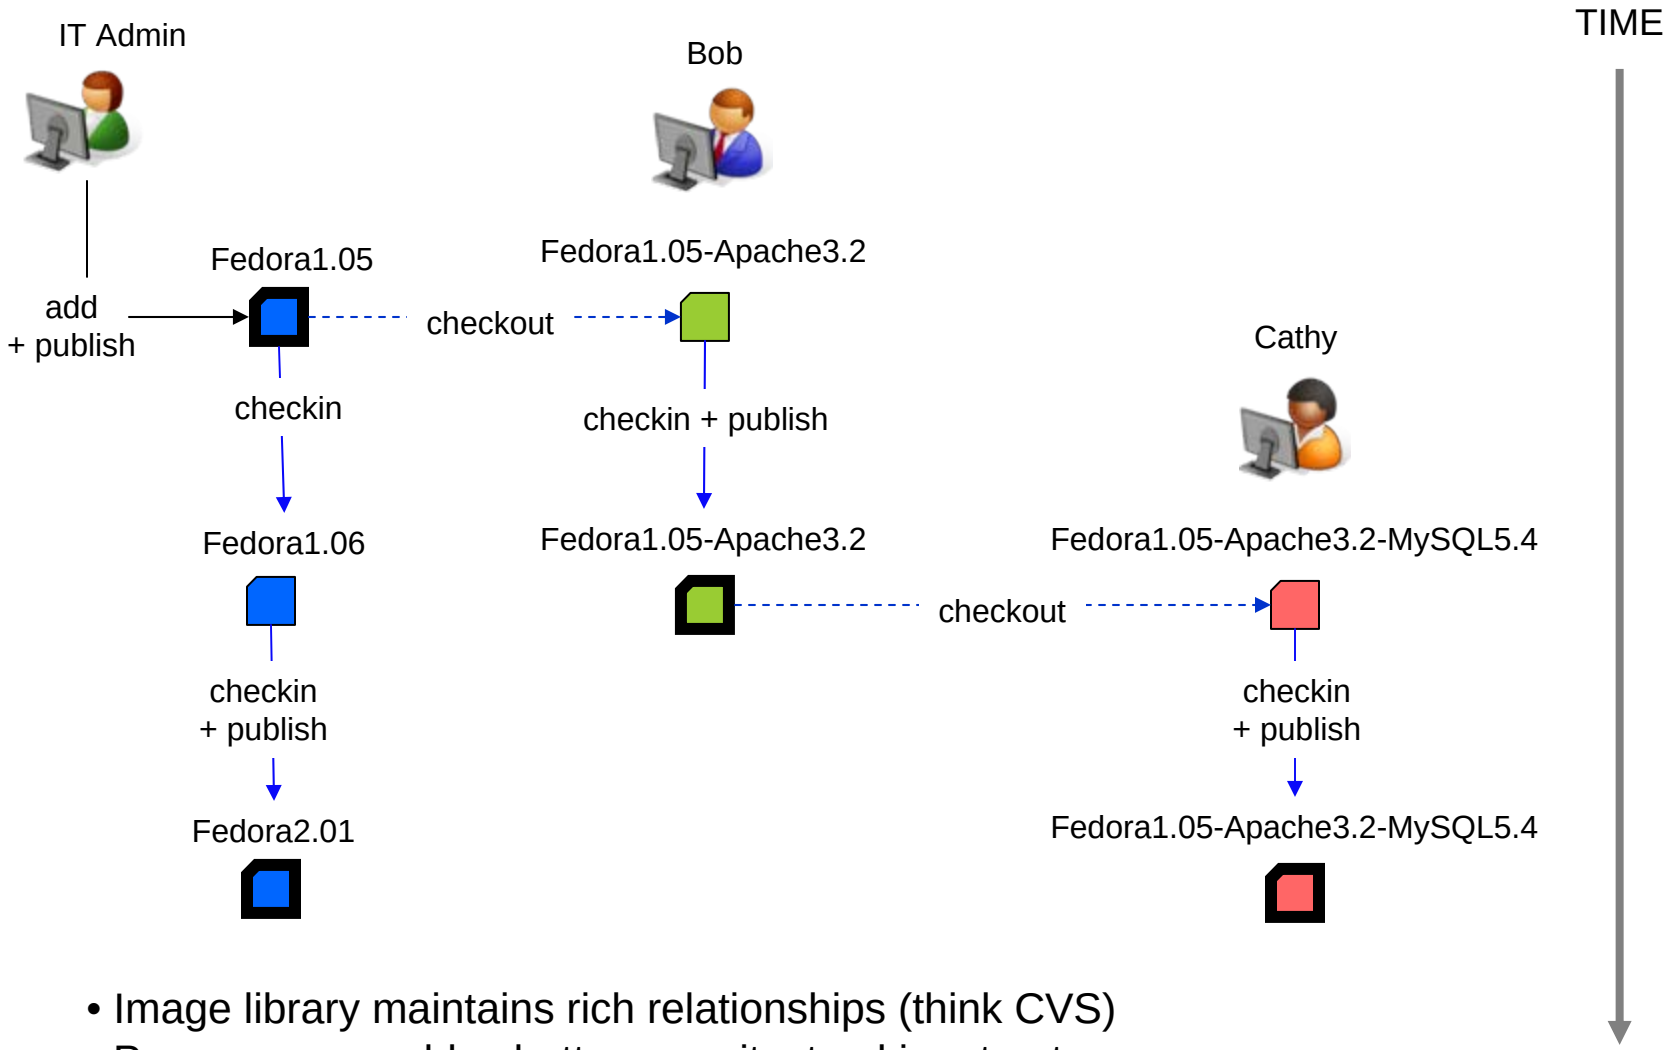

- Image library maintains rich relationships (think CVS)
- Provenance enables better security, tracking, trust

Image library: gateway to the compute cloud

Conventional image library

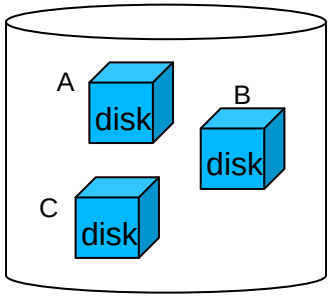

Disk block-level store

- Disk based representation
- No image relationships
- Hypervisor-dependent
- Merely a storage system for image disks

IBM Research image library

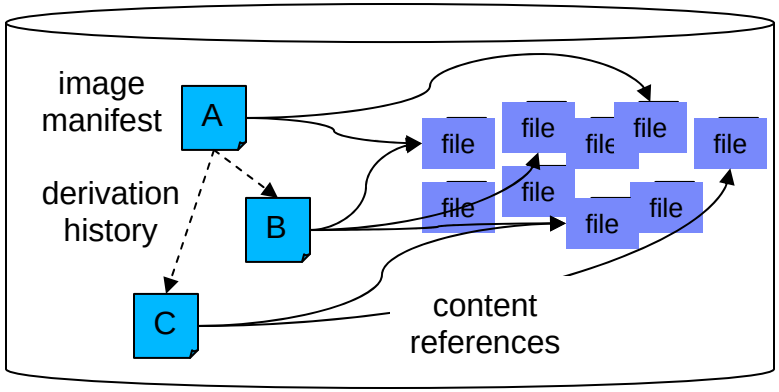

Content addressable, file-level store

- File based representation
- Image relationships (think CVS)
- Hypervisor-agnostic
- A sophisticated store with APIs to directly manipulate images without deploying them as instances or fully assembling their disks
- Conventional disk is reconstituted when an image is checked out

Open VM Image Library vision

- A logically centralized image library, hosted in the public domain (not IBM)
- Open standards for images, metadata, etc
- Pluggable interfaces to attach clouds, individual desktops
- Test bed for new research in cloud mgmt technologies

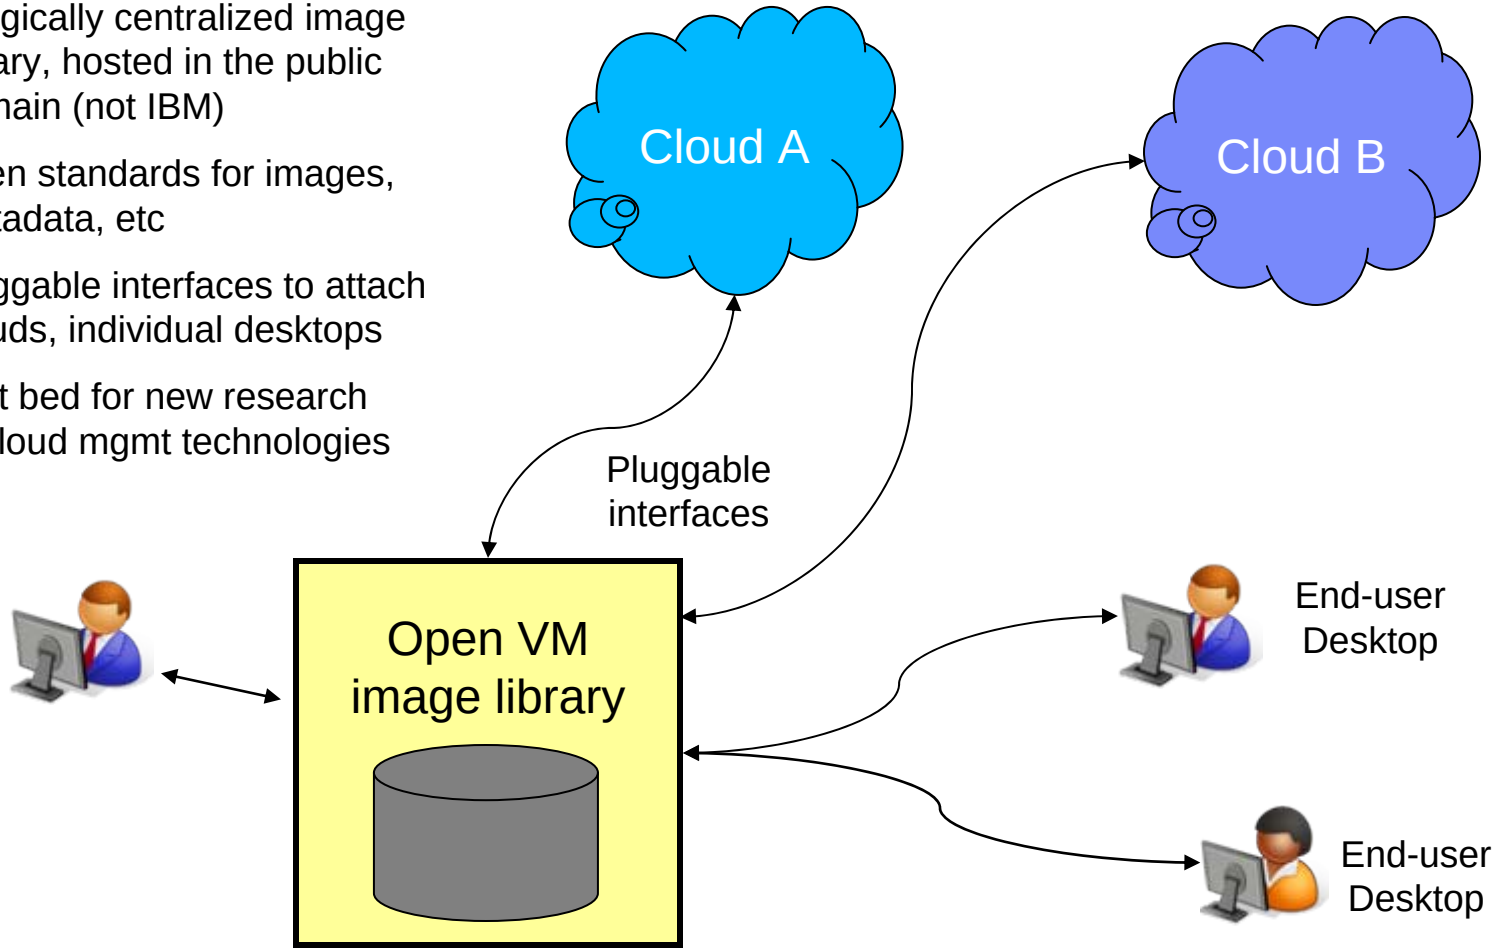

APIs for searching, securing, licensing, extending, patching, and instantiating images. Standard image formats (e.g. OVF)

Innovations in solution composition

Solution Composition Tool

Tool understands each image's metadata, enabling it to assist in the configuration

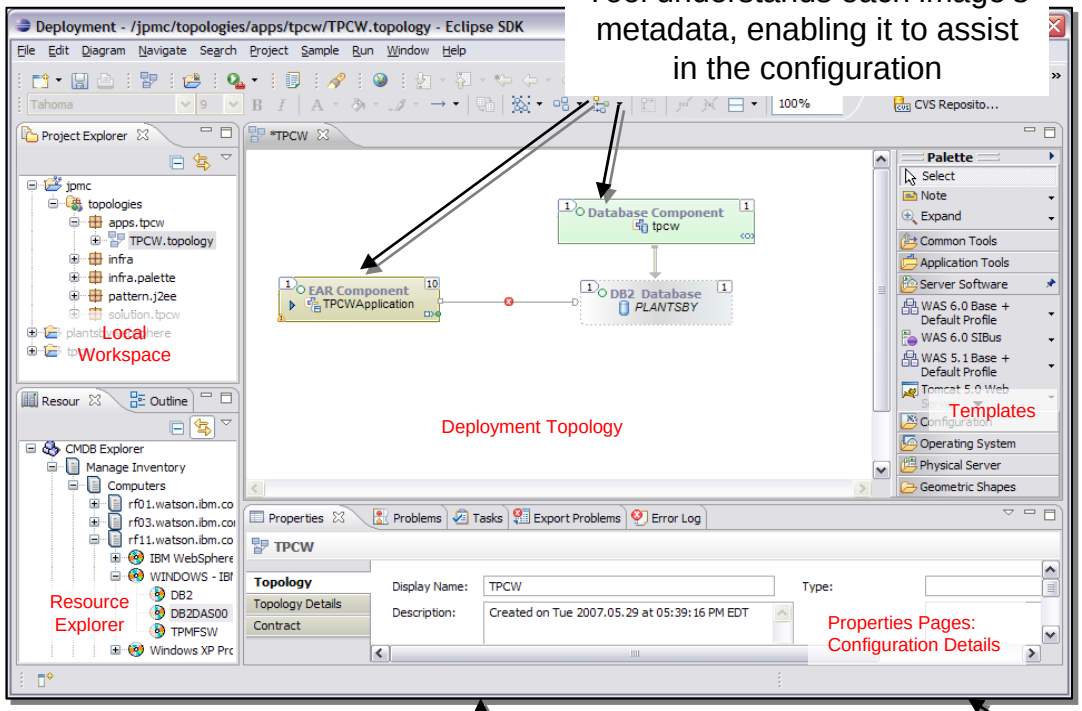

Solution = composition of images

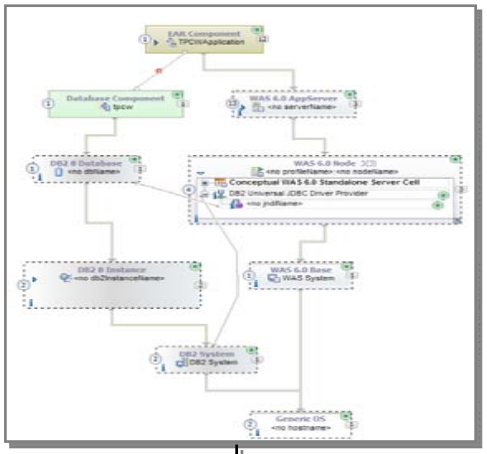

Innovations in security isolation

Isolation stops at the hypervisor boundary

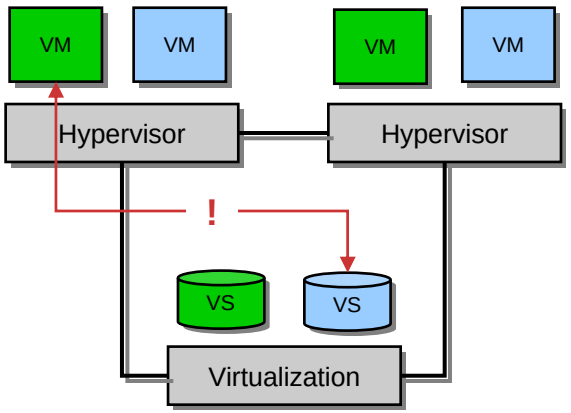

Connectivity is specified using physical device addresses

Isolation extends beyond the hypervisor boundary

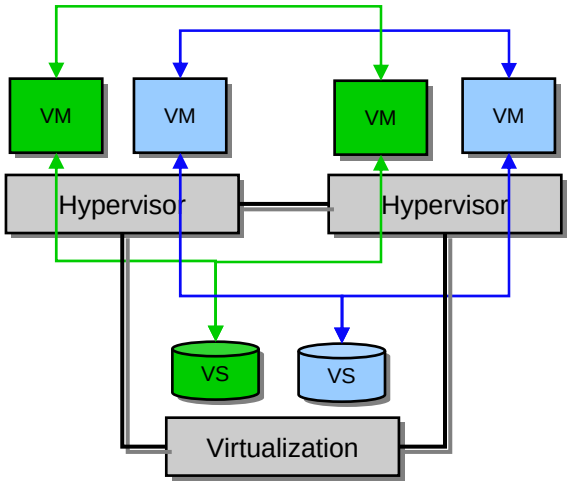

Connectivity is specified using virtual device addresses

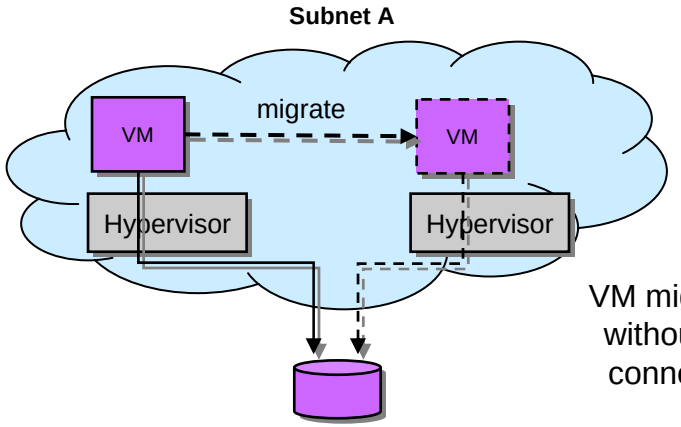

VM migration is possible without having to fixup connectivity attributes

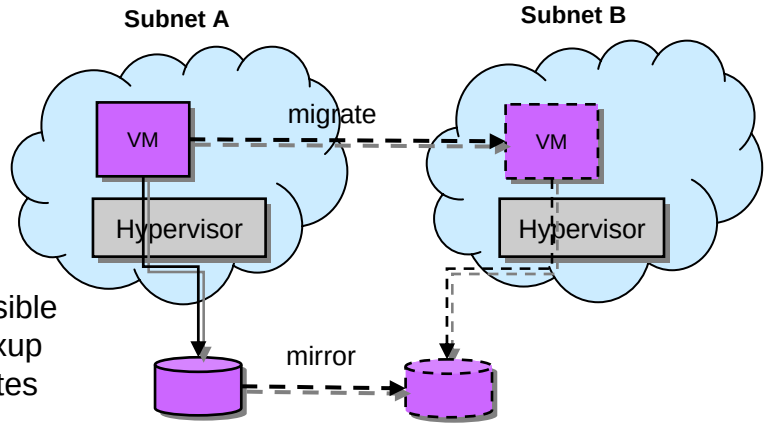

Innovations in power management

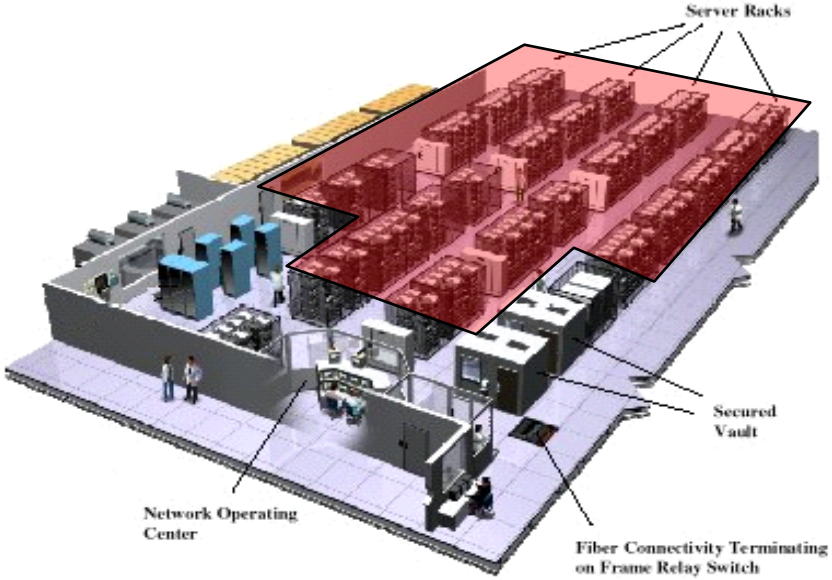

Load balance

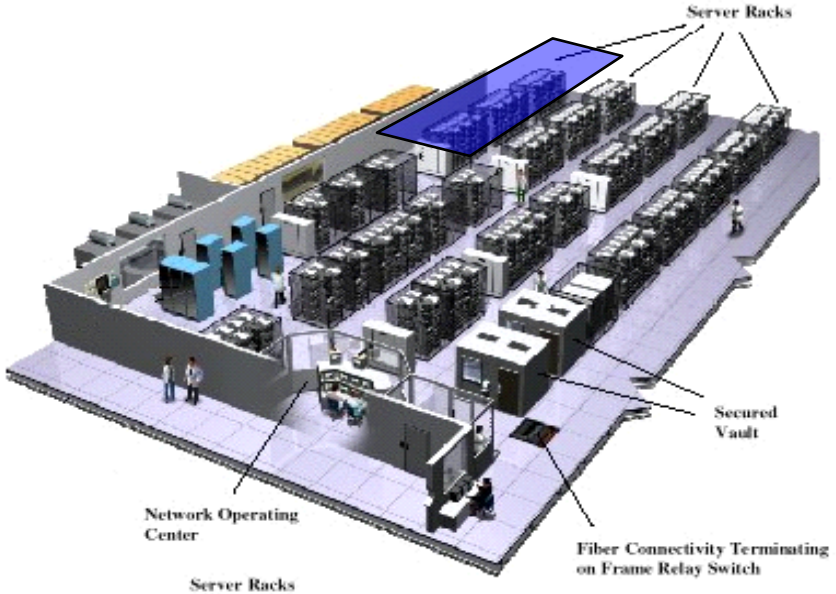

Power-aware, but cooling unaware

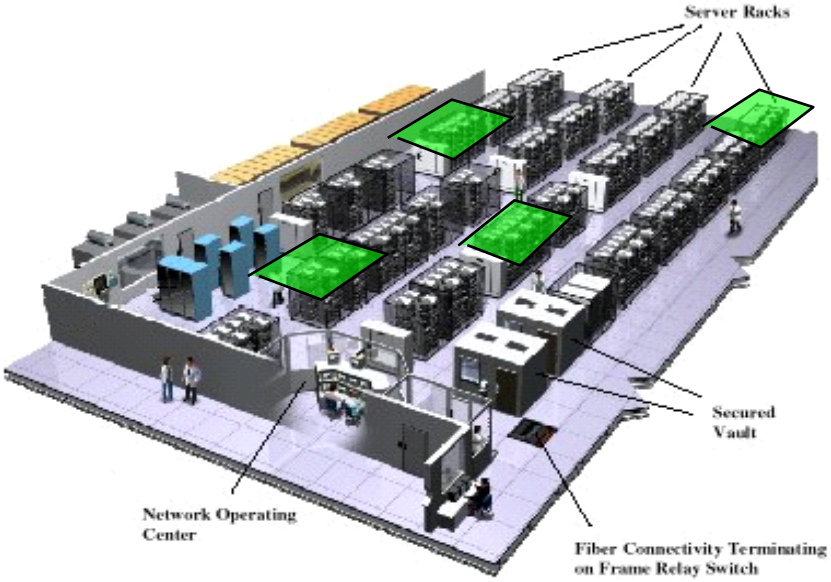

Power-aware, and cooling aware

Questions?

IBM's cloud computing centers