

Southern States Education Cloud Consortium Workshop

Goal:

To construct a multi-state community, together with SURA & TTP, as well as corporate partners like IBM, that will provide:

1. a support community for VCL production environments, addressing planning, set-up, maintenance, adding new schools to an existing community, etc.;
2. a research/development community, that will include SURA, TTP, corporate & other sponsored projects, as well as the VCL Apache open source community project;
3. pursuit of joint funding proposals to NSF & other federal programs, to industry, etc.
4. ultimately establish a governing organization and support assets (wikis/websites/dbs, tool & image repositories, best practices & troubleshooting documentation, etc.).

Additionally provide a VCL problem clinic, to help the current participants move forward with their VCL installations

Attendees:

NC: NC State, UNC-CH, NCCU, NCA&T **SC:** U of SC

DC: Howard **VA:** GMU, NSU **MD:** UMBC, MSU, USM, UMBI

GA: GSU **TN:** TSU **AL:** Tuskegee **LA:** Southern U BR

SURA, TTP/ELC, IBM

{Clemson has also joined the consortium but could not send reps}

Southern States Education Cloud Consortium Workshop

Monday, July 13th

Host: Andy Rindos, IBM

- 8:30 - 9:00 AM **Arrival at GMU, Welcome** Joy Hughes, CIO GMU
- 9:00 - 10:00 AM **"VCL Experiences and Some Suggested Best Practices: NC State,"** Mladen Vouk (NC State)
- 10:00 - 10:30 AM **"The Business Case for VCL Integration of Campus Utility Computing & HPC,"** Sam Averitt (NC State), Pat Dreher (RENCI)
- 10:45 - 11:00 AM **"Update on NSF Proposal Opportunities,"**
Art Vandenberg (GSU), Sam Averitt (NC State), Gary Crane (SURA)
- 11:00 - 11:30 AM **"The VA VCL: Plans and July 9 Event Summary"**
Sharon Pitt & John Savage (GMU)
- 11:30 - 11:50 AM **"VCL at GSU and Plans for GA,"** Art Vandenberg et al. (GSU)
- 11:50 AM - noon **Status on MD**
- 1:00 - 1:20 PM **"VCL at Morgan State U,"** Gerald Whittaker et al. (MSU)
- 1:20 - 1:40 PM **"VCL at UNC-CH,"** Ruth Marinshaw; Ken Chestnutt
- 1:40 - 2:10 PM **"VCL at NCCU & NCAT,"** Cameron Seay; Paul Campbell (NCAT)
- 2:10 - 2:40 PM **"The HBCU Cloud: Plans and Atlanta Forum Summary,"**
Ramon Harris (TTP/ELF)
- 2:40 - 2:50 PM **Status on SC,** Bob Brookshire (U of SC)
- 3:00 - 5:15 PM **"VCL Problem Clinic,"** Aaron Peeler (NC State) et al.

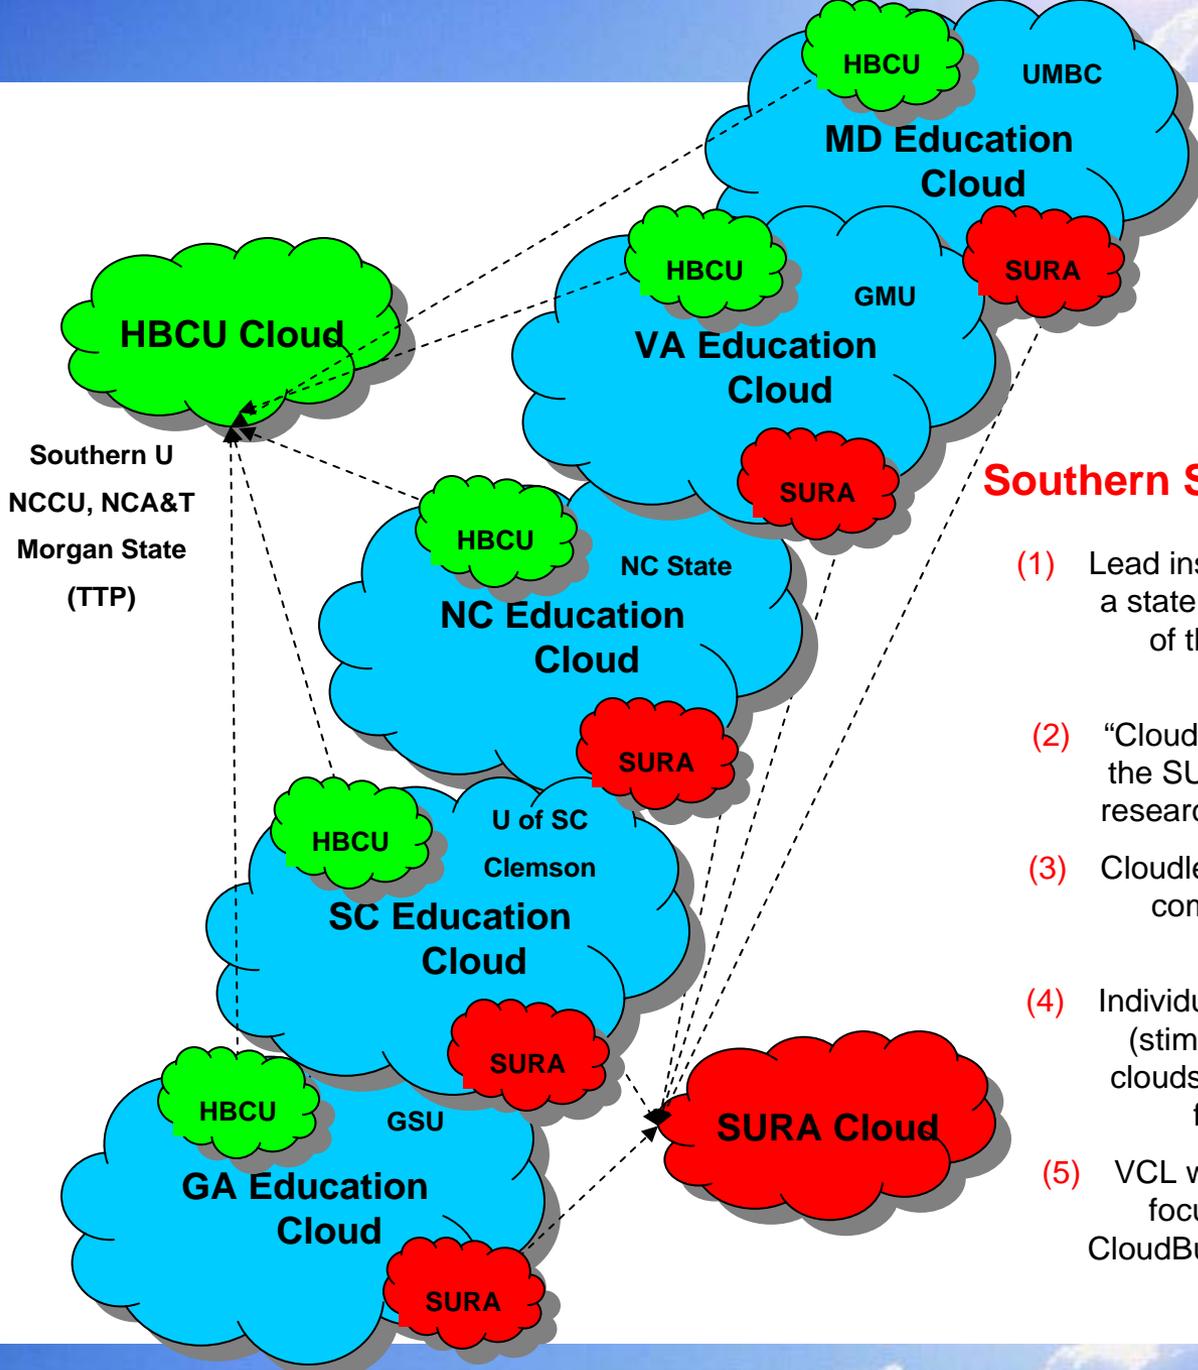
Southern States Education Cloud Consortium Workshop (cont.)

Monday, July 13th (cont.)

7:35 -10:30 PM **"NSF MRI Proposal Planning,"** All

Tuesday, July 14th

- 8:30 - 8:40 AM **Welcome** Roger Stough, VP, Res & Econ Dev, GMU
- 8:40 - 9:15 AM **"The VCL Apache Project,"** Aaron Peeler
- 9:15 - 9:45 AM **"The University Delivery Services (UDS) Center Program,"**
Jerry Haegele (IBM)
- 9:45 - 11:30 AM **"VCL Configuration Recommendations & Capacity Planning
Guidelines etc.,"** Mladen Vouk (NC State) et al.
- 11:30 AM - noon **"SURA & SURA Grid as a Model for a VCL User and Developer
Community and NSF MRI proposal,"** Gary Crane (SURA)
- 1:00 - 2:00 PM **Wrap-up/Next Steps** Andy Rindos (IBM)



Also discussions on state education clouds for FL, PA, MA, NM, CA, others

Southern State Education Cloud Consortium

- (1) Lead institutions in each state establishing VCL pilots for a state education cloud – goal is to follow the trajectory of the NC Cloud, adding other state universities, community colleges and K-12
- (2) “Cloudlets” from each cloud will be combined to create the SURA Cloud, providing resources for a multi-state research program around integration of HPC into cloud
- (3) Cloudlets from each cloud and/or member pilots will be combined to create the HBCU Cloud (led by the Technology Transfer Program or TTP)
- (4) Individual institutions will partner with IBM to seek state (stimulus and other) funds for their respective state clouds; and partner with IBM, SURA and TTP to seek federal funds for the SURA & TTP Clouds
- (5) VCL workshop for consortium took place July 13-14: focus on establishing set-up “assets”, POCs (for CloudBurst etc.), strong Apache open source community around VCL, etc.

Apache Incubator Virtual Computing Lab

- NCSU donated VCL 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/VCL>

VCL at GSU

Art Vandenberg
Director, Advanced Campus Services

Dr. Michael Russell
Director, University Academic Services

Kelly Robinson
Software Systems Engineer, Lead

About GSU

- Founded in 1913
- Enrollment of 28,238 undergraduate and graduate students (Fall '08)
- 1,046 Full Time Faculty members
- 5 labs, 242 classrooms, more than 1000 computers

Why VCL?

- Can provide more diverse offering of platforms for applications
- Off campus access to computing resources
- Ability to accommodate HPC during off peak hours
- Can maximize compute cycles and usage of the system
- Low cost

Research Computing Resources

- IBM System p5 575 with POWER5+ processors (URSA)
- Univa UD GridMP
- SURAgrid
- IBM Cluster 1350 with Intel Xeon Quad-Core processors (Octans)

OCTANS

System Specifications:

- IBM 1350 Blade Center
- 40 HS21 XM Blade Servers with 2 Quad-Core Intel Xeon Processors
- 320 CPUs
- 640 GB memory
- LSF scheduler
- xCat cluster management tools

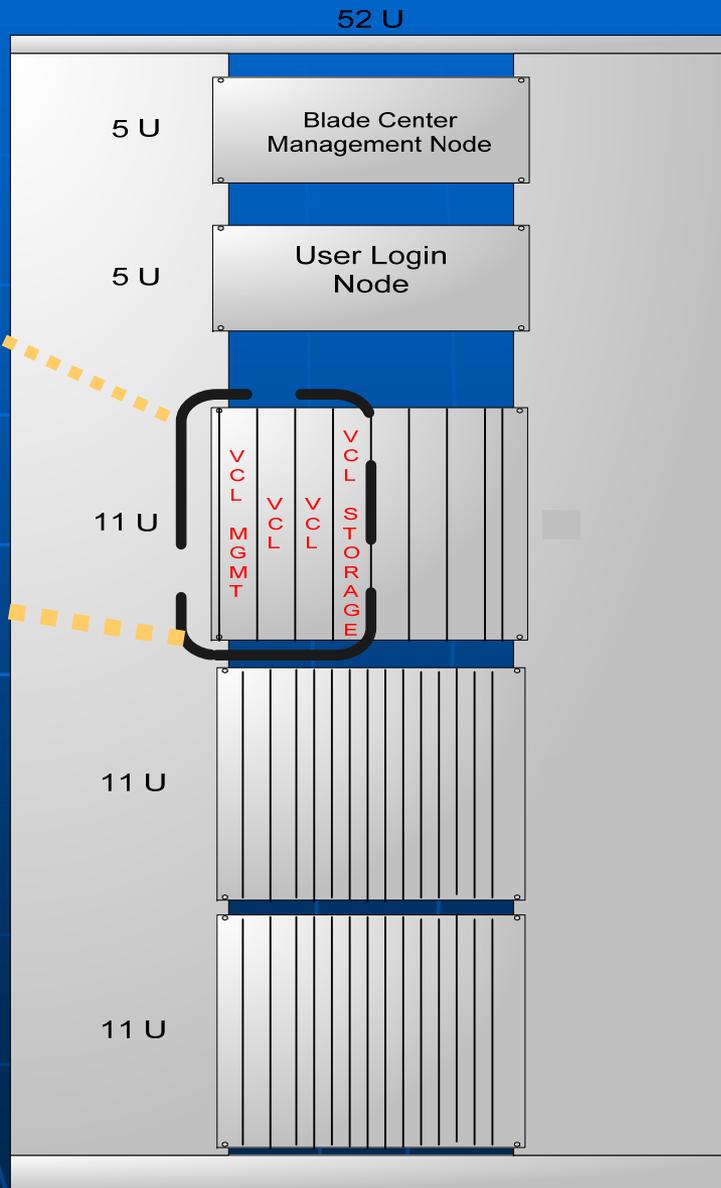
Applications:

- Gaussian
- Matlab
- CCP4

Our Approach



-VCL 2.1 Appliance
-ESXi / Vsphere client



Add an additional private and public network connection for VCL blades

Additional Considerations

- Image selection / criteria
- Security implications
- Reservation duration for HPC
- Licensing for applications used in virtual environment

The Consortium for Enterprise Systems Management

Robert G. Brookshire



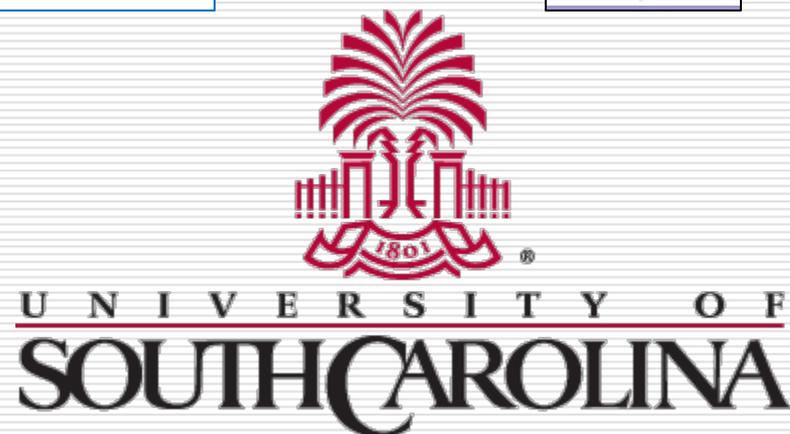
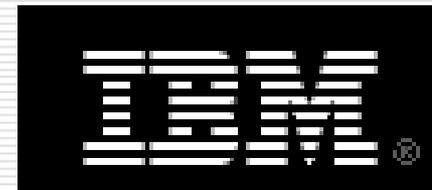
College of
HRSM
Hospitality,
Retail, & Sport
Management

university of south carolina

Mission

The Center for Enterprise Systems Management will promote education, human capital development, and applied research on the management and use of enterprise systems in business, government, and nonprofit organizations.

Consortium partners



Facilities

- ❑ Innovista Horizon II Building
- ❑ Data center with Z-series, blade servers, data storage
- ❑ Classroom space
- ❑ Seminar space
- ❑ Video conferencing system
- ❑ Office hotel space

Horizon II



The Consortium for Enterprise Systems Management will...

"Serve as a major national center for capturing, understanding, recommending, and providing proven education on the effective utilization and management of enterprise systems in business, government, and nonprofit organizations."



**NORTH
CAROLINA
CENTRAL
UNIVERSITY**
FOUNDED 1910

VCL at NCCU and the Technology Transfer Project

**Cameron Seay
School of Business
North Carolina Central University**

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:

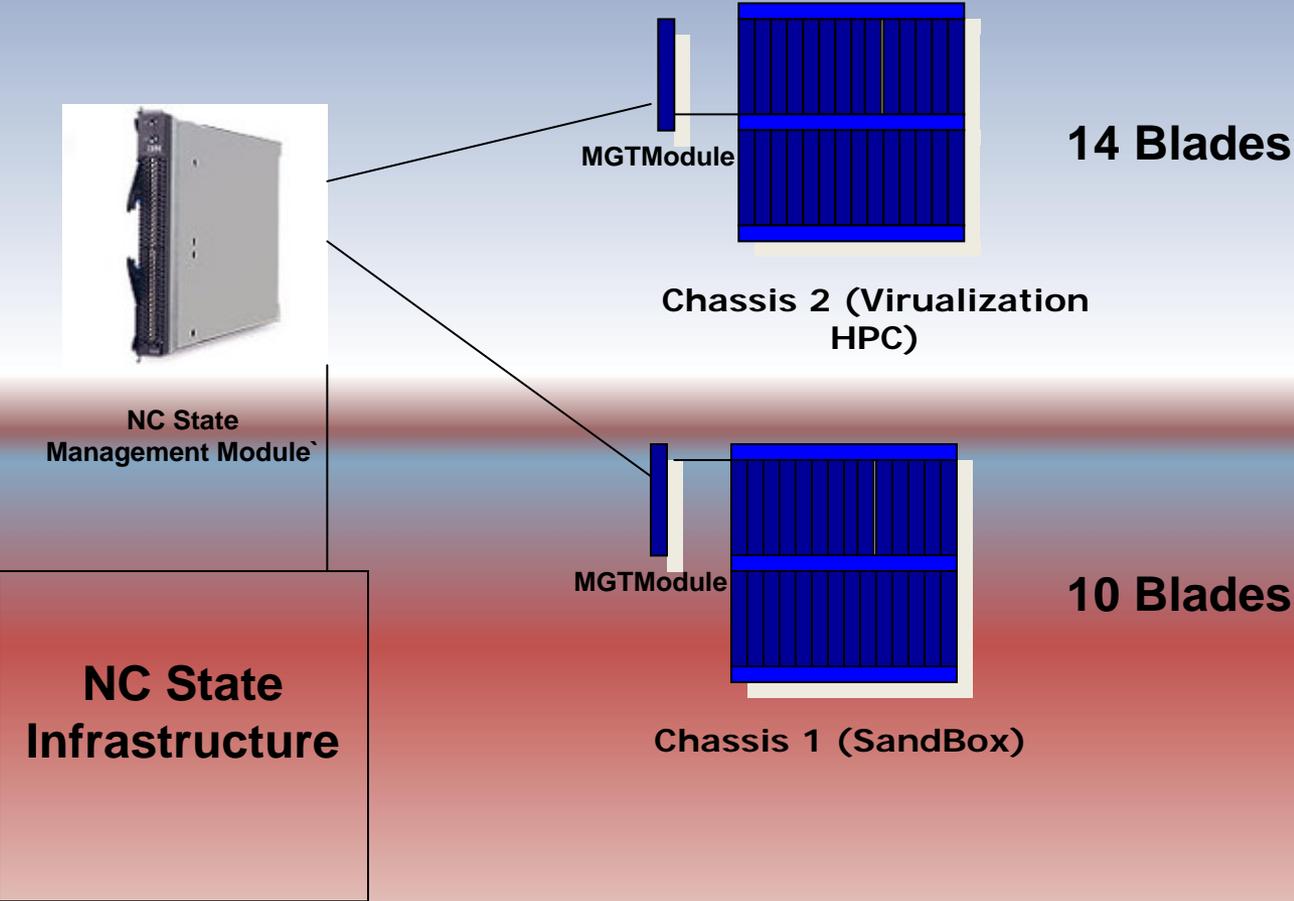
In the Fall of 2005, NCCU began to work with NC State and IBM to develop a presence of the Virtual Computer Lab (VCL) at NCCU

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

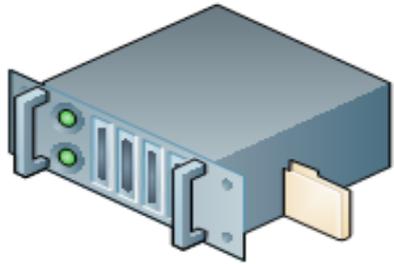
The result is a small VCL infrastructure tied directly to the VCL management logic maintained by NC State

NCCU Blade Center at Present

Housed at MCNC, RTP, NC



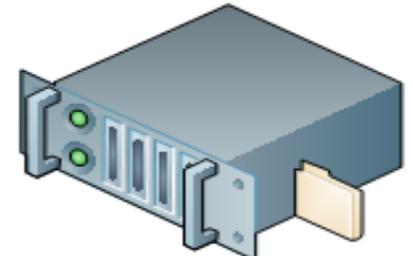
TTP Cloud



Howard
(proposed)



z10



A & T
(proposed)



Shaw U.
(proposed)



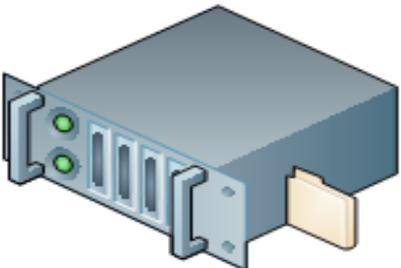
Bennett Col.
(proposed)



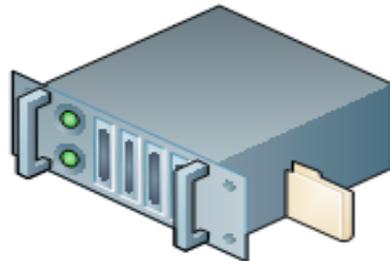
H.S.



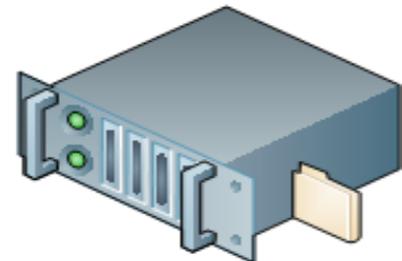
Middle School



Southern



NCCU
(MCNC)



Morgan

MSU VCL/Cloud Computing Initiative

Morgan State University Virtual Computing Initiative

July 13, 2009

July 22, 2009



The Beginning

- Held meetings to discuss the possibilities of a small project VCL
 - Faculty
 - Staff
 - Many were concerned about resources for support
 - University Information Systems Department
 - Systems department supported the concept
- Met with the Vice President of Information Systems Technology
 - Associate VP
 - Director
 - He supports the concept and assign a staff contact

Virtual Computing Laboratory

- The Virtual Computing Laboratory is a project by which the School of Engineering/to Morgan State University plan to leverage computing architecture provide a variety of services to the faculty, staff and students.
 - Using this system, students will be able to log in to the VCL over the internet.
 - Faculty and graduate students will be able to run specialized software packages.
 - Allow us to configure a limitless number of “systems” with a lower operational cost than most traditional IT setups.

Virtual Computing Laboratory

Current Status

- NOC has configured and installed the related Hardware/Software.
- Have some issues with the provisioning service
- Next, the Network Operations Center will start constructing the images to run on the VCL system.
 - Clone the School of Engineering General Computing Laboratory system.
- Solicit project / image requirements input from the VCL Team
 - Faculty
 - Staff
- Construct images and open the Lab for business.

Morgan State VCL Team

- Interested Faculty and Staff
 - Biology
 - Computer Science Department
 - School of Education
 - School of Engineering
 - Civil
 - Electrical
 - Industrial
 - Math Department
 - Research Staff

Potential VCL Projects

Arthur Neal Willoughby / Civil Engineering

- **Research Project: Computational Fluid Dynamics**
 - Highly computational.
 - Model and analyze such fluid scenarios that can not be performed experimentally.
 - CFD is predicting what will happen, quantitatively, when fluids flow, often with the complications CFD uses a computer to solve the relevant science-based mathematical equations, while using information about the circumstances in question.
- **Geographical Information Systems**
Using Satellite Tool Kit (STK) to plan and analyze mission to Mars and Moon

Potential VCL Projects

Electrical and Computer Engineering - Dr. Ladeji-Osias

- **Course #:** All departmental courses
- **MATLAB/Simulink:** Used by all departments in the School of Engineering. Mathematical and simulation software utilized in sophomore, junior, and senior courses in electrical engineering. Heavily utilized in signal processing and communications.
- **Electronics Workbench (NI MultiSim):** Used by electrical engineering students for circuits simulation especially in electronics courses.
- **ModelSim XE:** Used for the design and simulation of integrated circuits especially in computer engineering courses. This software is also used for **research** by Dr. Ladeji-Osias.

July 22, 2009

Potential VCL Projects

Dr. Jochen Mueller

Biology Department

Undergraduate and Graduate computational biology

Bioinformatics

Computational Analysis in Genomic Science Research

Dr. Vojislav Stojkovic's

Computer Science dept

- Intelligent software agents
- Computational Biological systems
- Quantum Computing applications
- DNA Computing applications
- Bioinformatics

Potential VCL Projects

Ahlam Tannouri **Mathematics Department**

- Give engineering and other students access to **MAPLE** and **MATLAB** via the internet. This will enable them to work on applied Mathematics problems through the MSU VCL, especially for courses like Differential Equations and Numerical Analysis.

Cynthia Brown-LaVeist **Research Engineer, C.I.B.A.C**

- Interested in partnering with Baltimore City/County public schools. Plan to focus for STEM topics to enrich the educational experience of the K-12 students.

Potential VCL Projects

Center of Excellence in Math and Science Education

Center focuses on and works with the Baltimore City School System in the STEM areas and Professional Development for Teachers.

Dr. (Mrs.) Anasutan N. Swamy

- Several others professors from the School of Education and the Computer Science Dept participate in CEMSE outreach programs

In the process of identifying several current and additional projects to VCL Implement in the VCL environment.

SOUTHERN STATE EDUCATION CLOUD CONSORTIUM



The IBM Application Services University Delivery Services (AS-UDS) Program: Towards VCL Support

June 13th and 14th

George Mason University, Fairfax, VA

G.E. (Jerry) Haegele, IBM Distinguished Engineer, CTO North America Application Services Delivery

Andy Rindos, Senior Technical Staff Member , Head, RTP Center for Advanced Studies (CAS) and WW
CAS Coordinator

First, consider a very short history of computing leading to the clouds...

“Computers were human.

Then they took the shape of metal boxes, filling entire rooms before becoming ever smaller and more widespread.

Now they are evaporating altogether and becoming accessible from anywhere.”

October, 2008

The
Economist

...but, there is 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

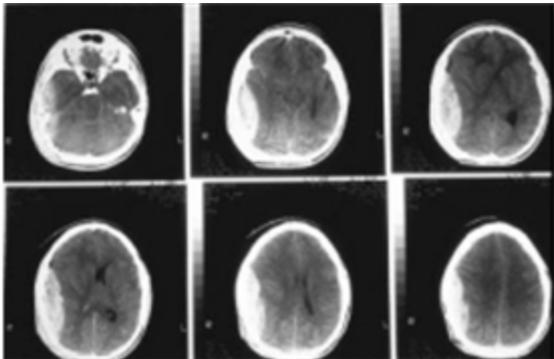
70% on average is spent on maintaining current IT infrastructures versus adding new capabilities.

85% idle

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

There may soon be an IT Gap (infrastructure, business services, data, and security).

Medical imaging



- By 2010, medical images will take up 30% of the world's storage

Financial markets



- By 2010, over half of equities trading will be algorithmic

Logistics and transportation



- By 2010, more than 30 billion radio frequency tags will be embedded in processes

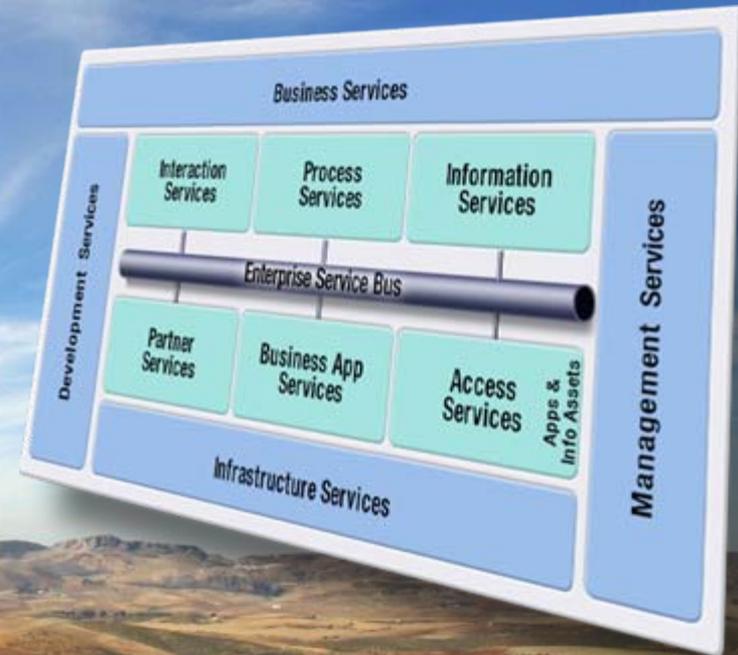
Cloud Computing provides many benefits to help fill the IT gap.

- **Business innovation** – Enables organizations to quickly and cost-effectively deploy new solutions for use inside and outside the company
- **Service delivery** – Enables the dynamic availability of IT applications and infrastructure. More rapid and enhanced service delivery reinforces efforts for customer retention, faster time to market and horizontal market expansion. **Cloud enhances SOA, information management and service management initiatives.**
- **IT optimization** – Offers massive scalability and can help lower cost of ownership, which drives higher profitability, enabling the ability to more easily reinvest in infrastructure.



A Service Oriented Architecture complements the capability of cloud computing to fill some IT gaps.

- **Both require similar capabilities:**
 - Architectural and organizational models
 - Optimization, Innovation and Value Delivery
 - Flexibility and Agility
 - Secure, reuse and sharing of 'services'
 - Separation of Concerns (Requestors, Providers, Creators, Brokers, etc.)
 - Improved Administration
- **Virtualization at all layers of the architecture**
- **SOA provides flexibility, reuse, separation of concerns, etc.**
- **Exploit a dynamic and elastic environment to enable innovation and to get optimum use from resources.**



A dynamic infrastructure provides the foundation for effective cloud services delivery, but there is more to a cloud than infrastructure.

Cloud Services Delivery

- Elastic Scaling
- Rapidly provisioned
- Flexible pricing
 - Ease of use
- Standardized offerings

Required Infrastructure Characteristics for effective Cloud Delivery

- Open standards-based, service-oriented
- Advanced virtualization and automated mgmt
- Common components and processes
- Advanced security and resiliency
- Easy to use service catalog

What IT services 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

We must move past today's challenges to seize tomorrow's opportunities.

HOW CAN WE ADDRESS ...

HIGHER SERVICE EXPECTATIONS

Internet-savvy consumers and employees expect 24x7 access to quality services.

RISING COST PRESSURES

Staggering levels of complexity and inefficiency drive up cost and stifle innovation.

NEW RISKS & THREATS

The connected, collaborative world is also a more vulnerable world.

WHILE ALSO LAYING A FOUNDATION FOR ...

BREAKTHROUGH PRODUCTIVITY

Almost any person, object, or service can become digitally aware and connected creating new possibilities for change.

ACCELERATED VALUE CREATION

More adaptive capabilities like cloud computing create new opportunities.

INCREASED VELOCITY

The faster pace of business and society demands a more responsive, agile infrastructure.

IBM is focused on several Cloud Computing client priorities and university cloud initiatives.

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 Delivery

- **What software?**
collaboration, business services, 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?

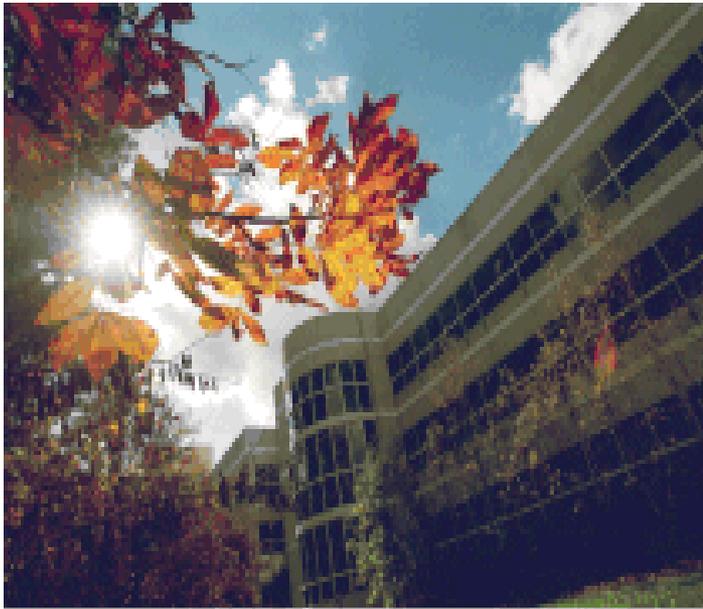
Hardware, Middleware, Management & Billing, Tools, Services?

How can universities team with IBM University Relations Center for Advanced Studies and IBM Application Services University Delivery Services?

It's time to start thinking **differently** about infrastructure, applications (business services), data, and security. The world is converging.

IBM needs universities to team with to provide key education, skills, and research so we all can work a lot smarter and more efficiently.

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



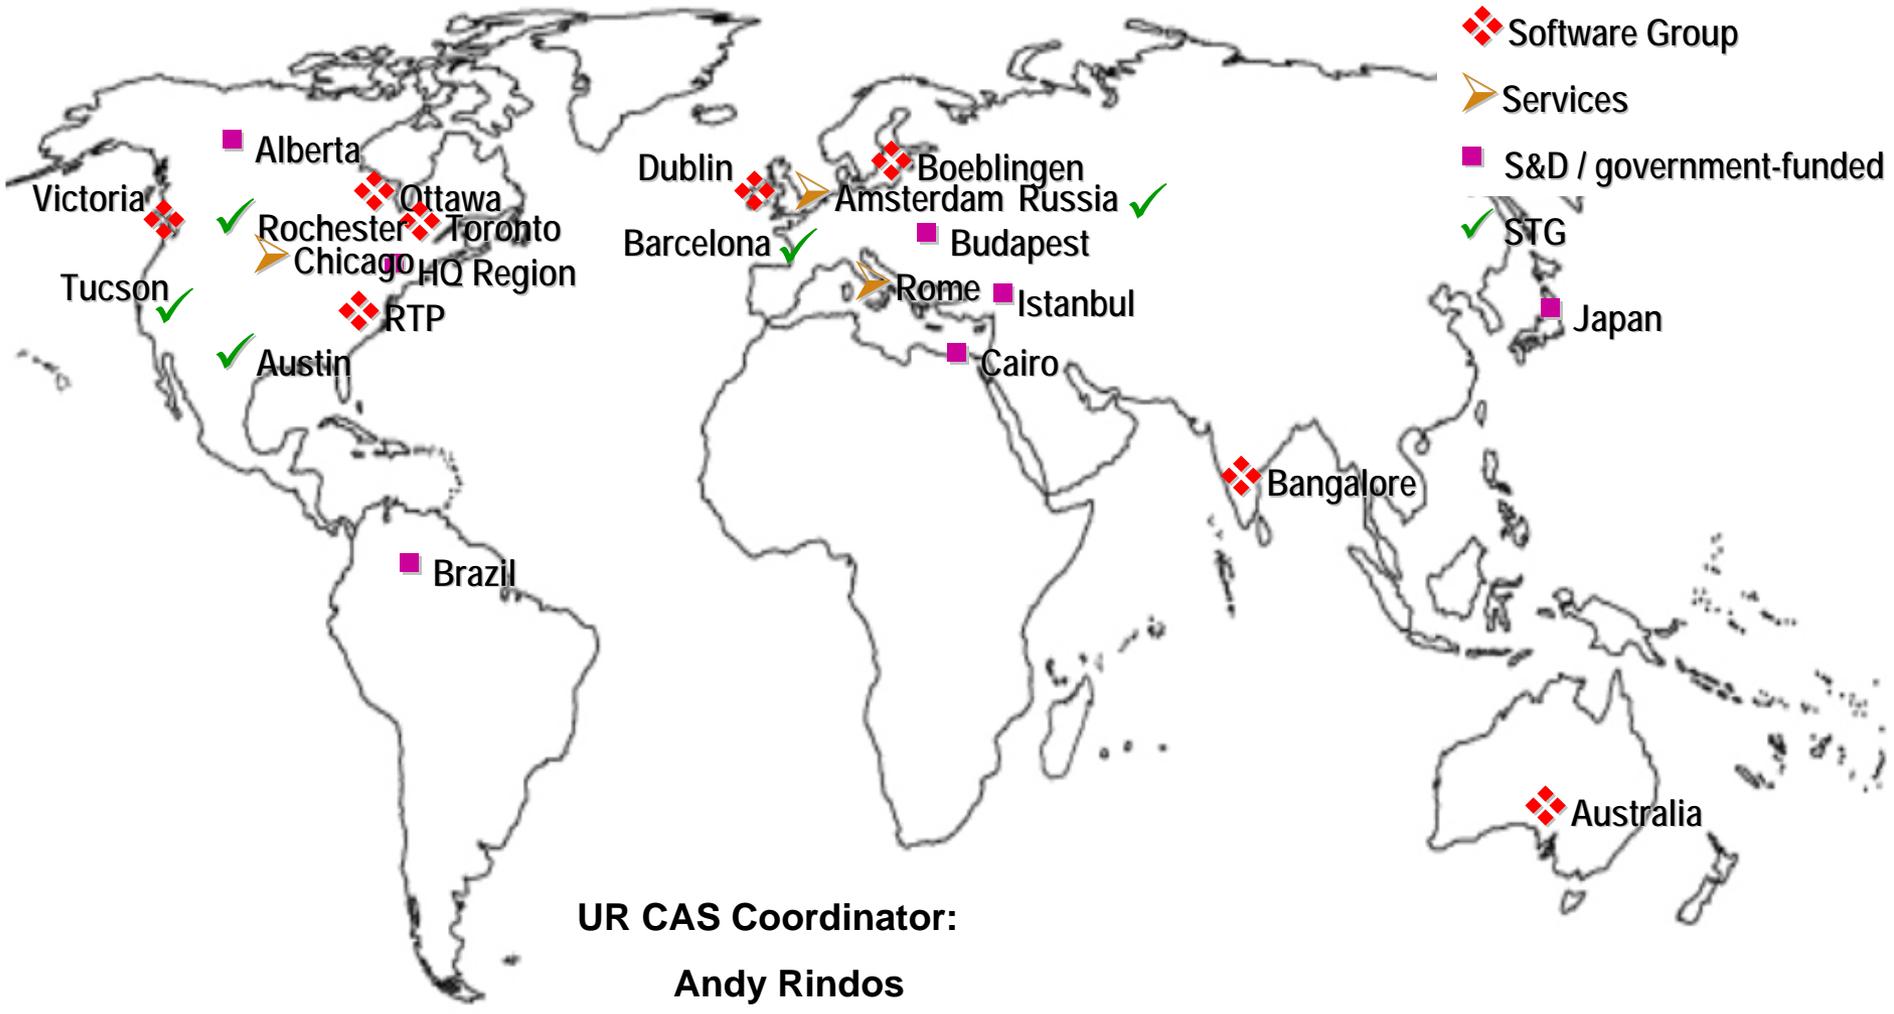
Universities and IBM have a long collaborative history. Cloud Computing will further enhance our relationships.

- IBM and university teaming is nothing new. The IBM T.J. Watson Research Center was formed with Columbia University during WW II and was located on its campus for many years afterward.
- According to many analysts, IBM is a leader in Cloud Computing.
- IBM can bring its focus on business and technology through CAS and AS-UDS to add to a university emphasis on education and research to create an innovative end-to-end value environment.
- IBM may provide services that a university may find difficult to invest in.
- The IBM AS-UDS helps raise the level of innovation and teamwork for the university and students with IBM real-world industry experiences.
- We want to create a true value-add Win/Win for both a universities and IBM.
- IBM Application Services University Delivery Service (AS-UDS) can deploy university students to support VCL production environments.

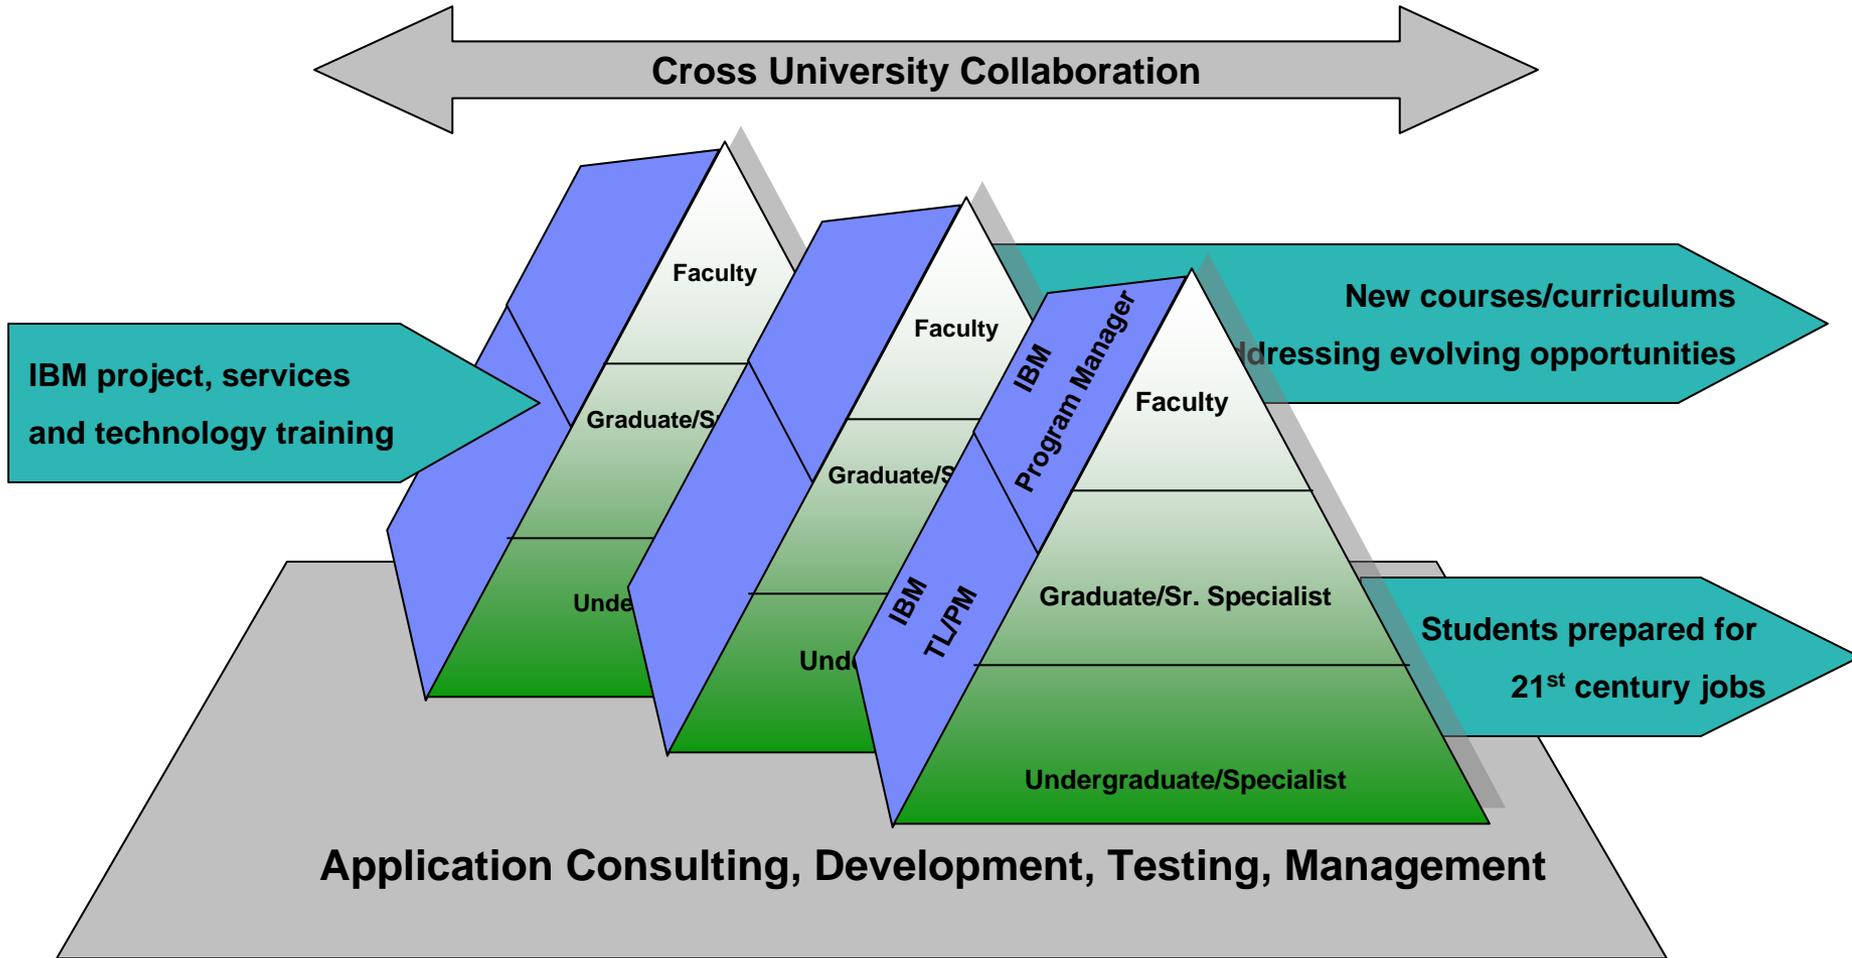
University and IBM teaming provides many benefits and creates a Win/Win scenario.

- IBM University Relations programs provide many opportunities to team with IBM development and IBM Research.
- For example, 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 is eager to team with universities in the development of curricula that provides graduates with highly marketable 21st century skills.

The IBM Centers for Advanced Studies includes 23 centers worldwide.



The Talent Solution: AS University Delivery Services Model provides an environment for University and IBM collaboration



Marketplace Synergy: The AS-UDS capability creates benefits for universities, clients, IBM, and the marketplace.

Clients

- Lower cost domestic delivery services
- Access to best and brightest students
- Tailored services
- Influence the skills students develop
- Pipeline to new, “ready-to-roll” employees

Students

- Enhanced market value
- Targeted course work and Access to latest IBM technology
- IBM mentoring and Employment during school year and summers
- Pay scale based on course work completion
- Inside track to full time IBM employment
- Real client project and delivery experience

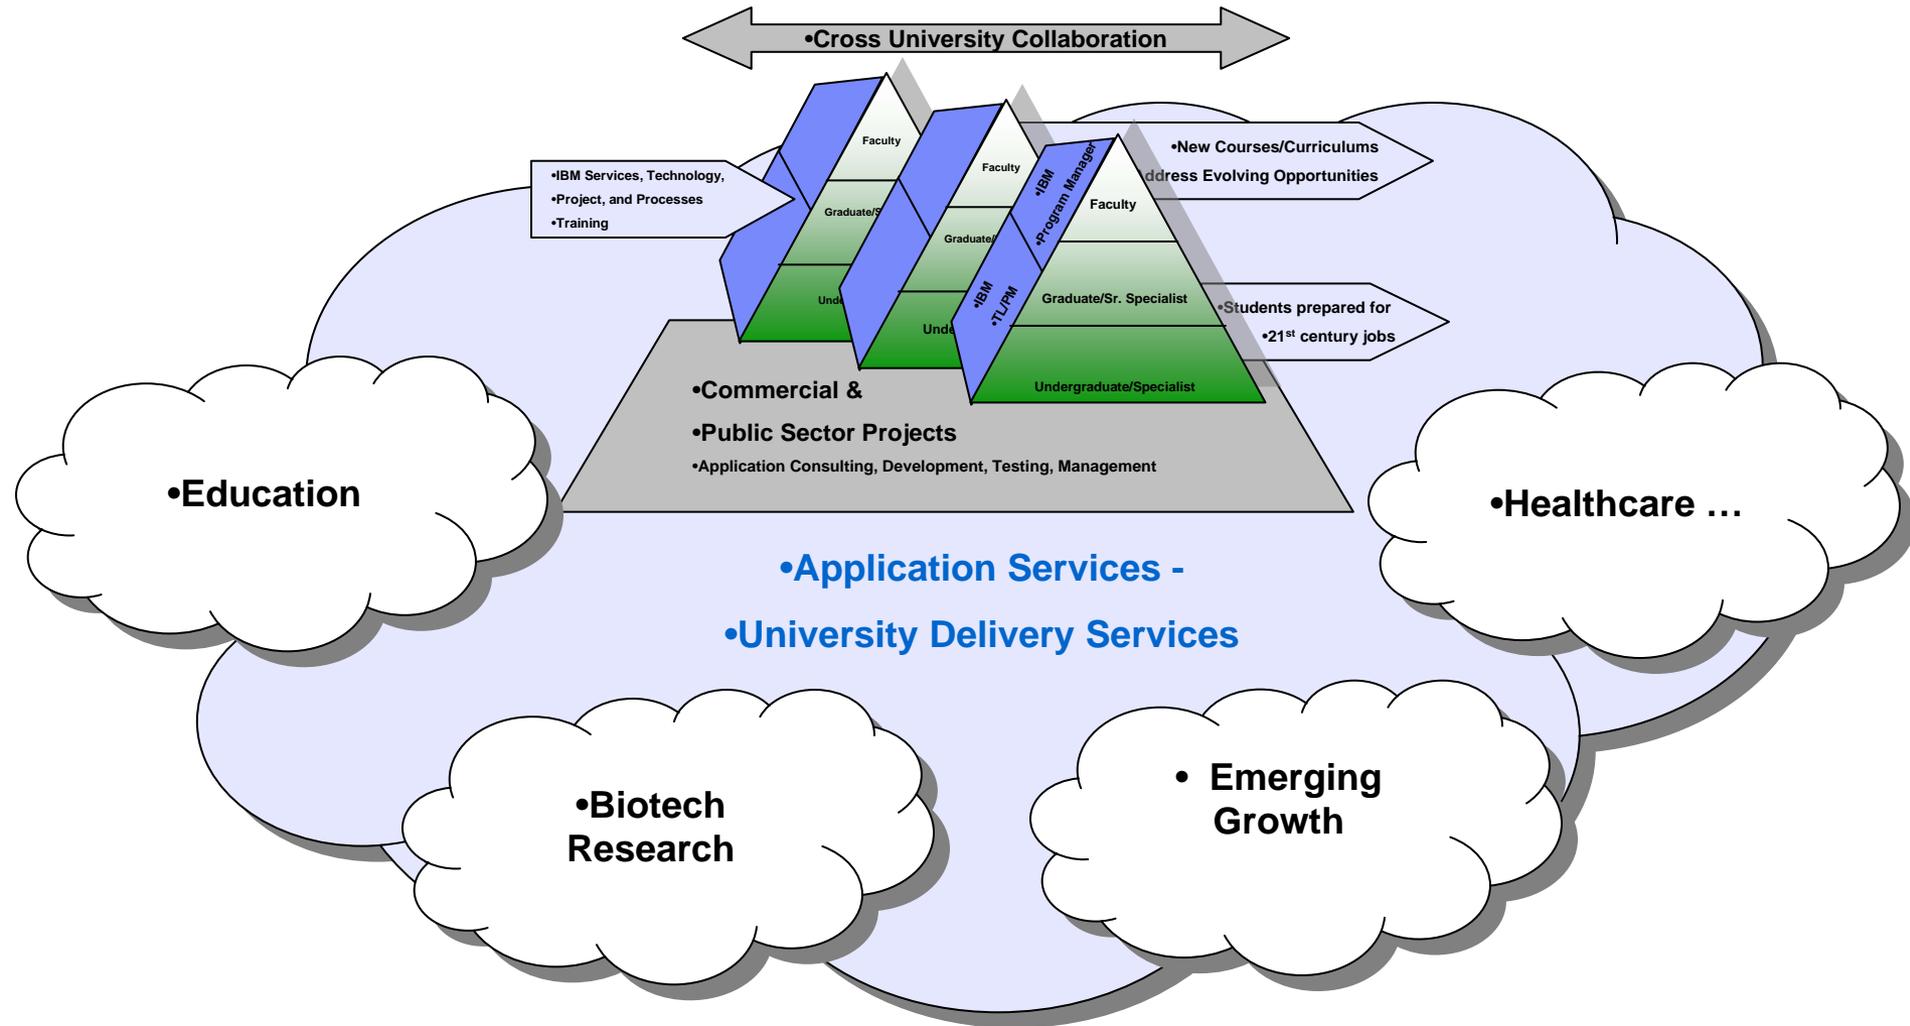
Universities

- Curriculum to match business needs
- Professors engaged in client projects & research
- IBM technology expertise – commercial & emergent
- Recruit best students/more applicants
- IBM on campus presence

IBM & Marketplace

- Curriculum input = marketplace needs
- IBM technology infused into curricula
- Provide visiting lecturers
- Training ground for high-potential employees
- Access to best students
- Onshore centralized delivery capability
- Cost effective/flexible delivery model
- Pipeline to new IBM employees

The IBM AS-UDS enables effective delivery of cloud computing-based applications and services



The VCL Apache Incubator Project and The Virtual Computing Initiative (VCI) are two examples of teaming.

- **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

The SMARTER CLASSROOM through Cloud Computing incorporates the Virtual Computing Lab (VCL) and VCI

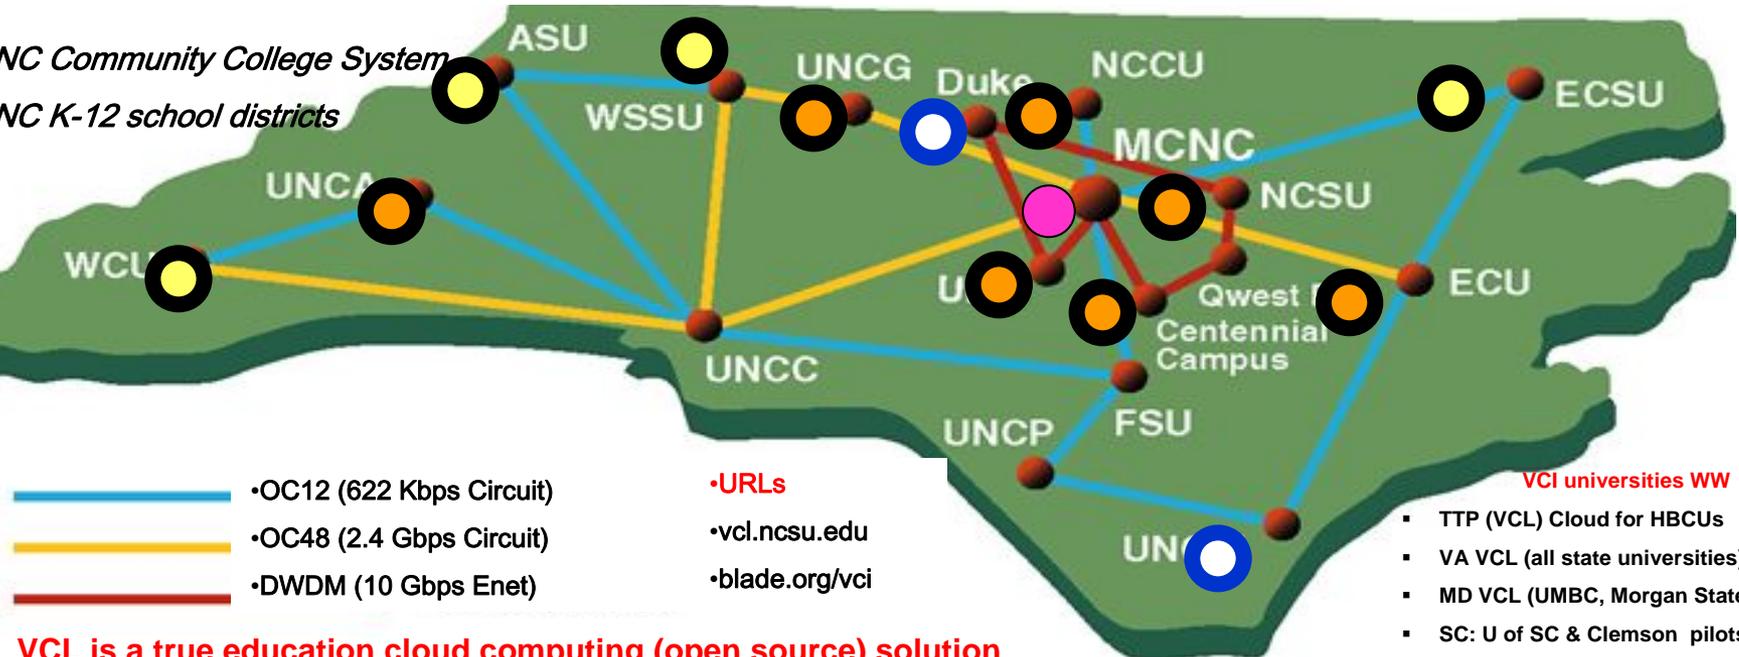


East Carolina University



... And Growing Across North Carolina and the world: VCI announced October 2006 (and sponsored by blade.org); WSTI sponsored the VCL Apache project in October 2008 ('07 Computerworld Honors Program Laureate)

- + NC Community College System
- + NC K-12 school districts



VCL is a true education cloud computing (open source) solution developed by NC State & in production for almost 5 years.

IBM Software Group (SWG) provides 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**
 - **WebSphere CloudBurst Appliance**, a DataPower-based box, will manage and provision WebSphere-based clustered images into the VCL cloud (for use in NC State educational programs in SOA, etc.). Currently working to place VCL appliance into box.
 - 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 Services 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 is also part of the overall collaboration.

- **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) Desktop Cloud (formerly VIA or 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 data center in RTP.**



Thank you!

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

Or buy us coffee and talk to us here at the consortium,

Or contact us at:

Jerry -- geh@us.ibm.com

Andy -- rindos@us.ibm.com

Legal Disclaimer

Notice: This document contains information on products in the design phase of development. The information here is subject to change without notice. Do not finalize a design with this information. Contact your local Intel sales office or your distributor to obtain the latest specification before placing your product order.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY RELATING TO SALE AND/OR USE OF INTEL PRODUCTS, INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT, OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, life saving, or life sustaining applications. Intel may make changes to specifications, product descriptions, and plans at any time, without notice.

The panel discussion is of the opinions of each respective individual, and not of the company they represent.

All products, dates, and figures are preliminary for planning purposes and are subject to change without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

The Intel products discussed herein may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's website at <http://www.intel.com>.

Intel® Itanium®, Xeon®, Pentium®, Intel SpeedStep® and Intel NetBurst® are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. Copyright © 2007, Intel Corporation. All rights reserved.

*Other names and brands may be claimed as the property of others.

NORTH CAROLINA A&T STATE UNIVERSITY

VCL 2009 REPORT



VCL REPORT

- ❑ Autodesk Applications 2008 & 2009 Pilot
- ❑ VCL login using our LDAP
- ❑ Collaboration with School of Technology & College of Engineering
- ❑ Collaboration with NCCU
- ❑ Campaign: Informing other departments on campus
- ❑ Training Faculty
- ❑ Campaign: Recruiting other HBCUs to VCL



Thank You & Have a Nice Day

Executive Leadership Foundation's
Technology Transfer Project:
History, Current State, & Outlook



VCL Technical Workshop
June 13, 2009

Technology Transfer Project - History

- **TTP Phase 1: Sunset in 2006**

- Focused on the enterprise level infrastructure and organization of academic institutions
- Developed strategic planning model

- **TTP Phase 2: 2005 - 2007**

- Exclusive IBM partnership with Executive Leadership Council
- Focused on enabling HBCUs to teach on demand and open IT concepts; and their applicability to science and business
- Leveraged technology-enabled instructional models and improved school access to web-based learning technology
- Concentrated on faculty and student enablement activities through existing IBM programs

- **TTP Phase 3: 2008 - 2010**

- Integrate additional ELC industry partner corporations for technical and financial support
- Increase the number of participating schools, professors, and students
- Assist HBCUs with ensuring that graduates are competitive and have skills required by industry

TTP Value Proposition

- **Council Partner Corporations**
 - A pipeline of diverse students with specific knowledge and skills
 - Preferential recruitment opportunities
 - Product and curriculum branding opportunities
 - Opportunities for research and development partnerships
 - Opportunities to partner with diverse institutions of higher education
- **HBCUs**
 - Access to premier software, technology, and curriculum
 - Effectiveness of overall programs increased through collaboration with industry
 - Exposure to emerging technologies and their use within the various industries
 - Faculty internship opportunities
 - Ability to ensure an industry-relevant curriculum
- **Students**
 - Ability to compete for valuable co-op/internship and full-time opportunities
 - Practical application experience provided within and outside of the classroom
 - Opportunities to develop a competitive edge
 - Enhanced understanding of the corporate and industry environments
 - Hands-on experience with leading-edge technology and subject-matter experts
 - Potential mentoring relationships

Technology Transfer Project Participating HBCU Institutions

- Bowie State University
- Florida A&M University
- Hampton University
- Howard University
- Morehouse College
- Morgan State University
- Norfolk State University
- North Carolina A&T University
- North Carolina Central University
- Southern University
- Savannah State University
- Tennessee State University
- Tuskegee University

TTP Phase 3 – Current State

Commitment of Support from ELC Partner Corporations

- **Chubb Insurance Companies** = Focus on Developing a Mainframe Pipeline (NCCU & NCAT)
- **Cisco Systems** = Cisco Academies, Recruiting (TSU, Howard, NCA&T, NCCU)
- **Deloitte Group** = Enterprise ICT Engagements (Business Intelligence – TSU)
- **IBM** = Academic Initiative, Recruiting, Professor Internships, VCL/Cloud Computing
- **Sun Microsystems** = Sun Academic Initiative, TTP Java Boot Camp
6 Sun Campus Ambassadors,
OSUM, Alice, Greenfoot, BlueJ
- **Symantec** = Information Assurance (Florida A&M University)
- **Verizon Foundation** = Thinkfinity.org (Virginia HBCU's Schools of Education)

Technology Transfer Project Phase 3 – Outlook

Goals

- Facilitate and drive the involvement of additional ELC and non ELC corporations
- Increase the number of HBCUs actively participating in the TTP
- Improve the awareness and use of corporate educational programs and platforms to stay abreast of industry requirements and future trends
- Increase the utilization of the Collaborative Learning Environment
- Facilitate and focus on improving the collaboration between TTP partner intuitions
- Partner, where appropriate, corporations with specific TTP participating educational institutions that have shared requirements and objectives
- Initiate a process where the activities of the TTP can be self maintaining
- Develop an HBCU Cloud Computing Environment

Technology Transfer Project VCL/Cloud Computing Initiative

Draft Vision

The TTP HBCU Cloud Computing Initiative (HBCU-CCI) will be viewed as a best practice in collaboration and sharing of academic, educational, and research resources for the development of students and faculty in higher education.

Draft Mission

The mission of the TTP HBCU Cloud Computing Initiative is to work collaboratively with the TTP, IBM, HBCUs, other institutions of higher education, as well as other entities to develop a Cloud Computing infrastructure that will leverage the intellectual, infrastructure, human, and financial resources of participants to enhance and maximize the use of information and communication technology in the preparation of students for industry and post graduate education.

Technology Transfer Project VCL/Cloud Computing Initiative

Goals

- Encourage faculty at TTP HBCUs to participate in the Apache Open Source initiative related to VCL
- Encourage TTP HBCUs to set up VCL Cloud Computing pilot projects
- Leverage the experience and work of existing organizations (North Carolina State University, North Carolina Central University, IBM, Apache.org) that have or are in the process of developing Cloud Computing environments
- Develop a private HBCU Cloud comprised of current TTP institutions
- Develop a governance and management structure for the development and implementation of the HBCU Cloud
- Develop a funding strategy that will support the initial development and implementation of the HBCU Cloud as well as a strategy for financial self maintenance
- Create content for educational programs to run on VCL, ie System Z, Service Oriented Architecture

Technology Transfer Project VCL/Cloud Computing Initiative

Lead Institutions

- North Carolina Central University
 - Dr. Cameron Seay
 - Operating for 18 months
- Morgan State University
 - Mr. Gerald Whitaker
 - Final Stages of implementation
- Southern University, Baton Rouge
 - Ms. Deanna Roquemore
 - IBM Faculty Award for VCL/Cloud Computing
 - Initial stages of implementation

Technology Transfer Project VCL/Cloud Computing Initiative

Highly Interested Institutions

- **Bowie State University**
 - Mr. Al Vabuena, CIO
- **Hampton University**
 - Dr. Jean Muhammad, Chair, Computer Science
- **Howard University**
 - Dr. Legand Burge, Chair, Computer Science
- **Norfolk State University**
 - Ms. Margaret Massey, CIO
- **North Carolina A&T State University**
 - Dr. Benjamin Uwakweh, Dean, School of Technology
- **Tennessee State University**
 - Dr. Dennis Gendron, CIO
- **Tuskegee University**
 - Al Bryan, CIO

Next Steps

- Continue to collaborate with existing organizations pursuing Cloud Computing
- Develop a governance structure for the HBCU Cloud
- Seek funding from government and non government organizations

Questions

7/13/09

VCL Technical Workshop
George Mason University

12



UNC
INFORMATION
TECHNOLOGY SERVICES

Update: VCL at UNC Chapel Hill

Ruth Marinshaw and Ken Chestnutt
UNC Chapel Hill
July 13, 2009



UNC
INFORMATION
TECHNOLOGY SERVICES

Piloting VCL Technologies in Campus Research Computing

Ruth Marinshaw
UNC Chapel Hill
July 13, 2009



How We Got Here

- UNC Chapel Hill's original VCL involvement thanks to Paul Jones (iBiblio), UNC Computer Science, Andy Rindos (IBM). Proximity to NCSU helpful as well.
- SUR award to get going
- Additional funding and expanded campus participation in 2009 (iBiblio, Research Computing (ITS), Information & Library Science, School of Public Health, RENCI, Computer Science)



Central Research Computing Resources

- Linux clusters
 - Multiple vendors, ~5000 cores, RHEL 5
 - >40 centrally-maintained applications
 - Multiple compilers
 - Multiple parallel environments
 - ~1300 accounts
- Shared memory systems (SGI Altix, IBM P575s)
 - SUSE , AIX

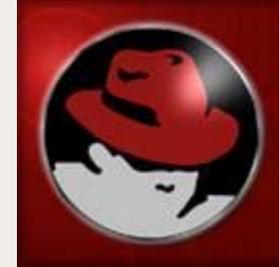


Central Research Computing Resources

- Lots of storage resources
- Lots of compute resources



BUT...



- Homogeneity of user environment within each system
- No Windows systems





Research Computing Goals for VCL

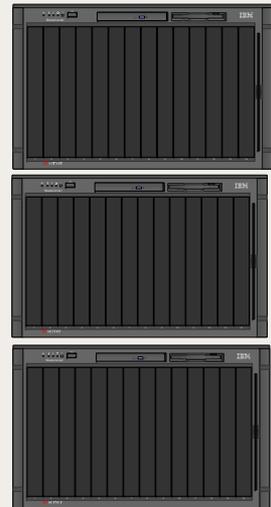
- Windows application environment
 - Social Sciences
 - Biosciences
 - Student research & learning
- Linux sandbox
- Test/dev environment for production services





■ Hardware

- 1 IBM BladeCenter H chassis - 14 blades (Primarily for Research Computing use)
- 2 IBM BladeCenter E chassis - 20 blades (IBM SUR grant partner use)
- Each blade:
 - ◆ Xeon E5405 @2GHz dual quad-core
 - ◆ 16Gb memory
 - ◆ 73Gb or 146Gb disk
 - ◆ 2 1Gbps Ethernet

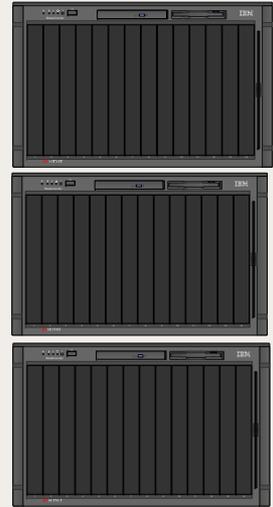


■ Images stored in GPFS file system



VCL Current Status

- Partnership with College of Arts & Sciences to build and pilot Windows applications
 - ArcGIS, SAS, Stata, Mplus, Matlab, Maple, Kurzweil 3000, NMIM, etc. In use in some departments & schools
- iBiblio assistance to faculty with linux builds
- Working on Ubuntu for some researchers; RHEL for research project computational modeling training; hope to pilot with some courses in fall
- Other IBM SUR partners to do their own thing
- Research Computing maintains hardware, base images, standard apps; provides admin





VCL to do's/wish list

- Documentation
- Windows XP x64
- Vista
- Linux virtualization
- Ubuntu
- Evaluation for lab environments
- Integration with Blackboard & Sakai
- Policies/governance

The End of my part ... now Ken's turn
Questions? ruth@unc.edu



UNC
INFORMATION
TECHNOLOGY SERVICES

The University System of Maryland institutions are very interested in developing a plan for a centralized VCL infrastructure for USM institutions, and possibly other Maryland higher education institutions, to leverage for supporting student access to software resources, faculty research computing, and potentially disaster recovery opportunities. UMBC is presently the lead institution and will be setting up a small VCL infrastructure in the fall where campuses can test out using different images on VCL and get a true sense of the potential.

The USM institutions are very interested in the Virginia VCL model and will be following that closely. We anticipate making further plans for expansion in the spring as our campuses get a better understanding of the opportunities and limitations of VCL and we work through a financial model for how this will work.



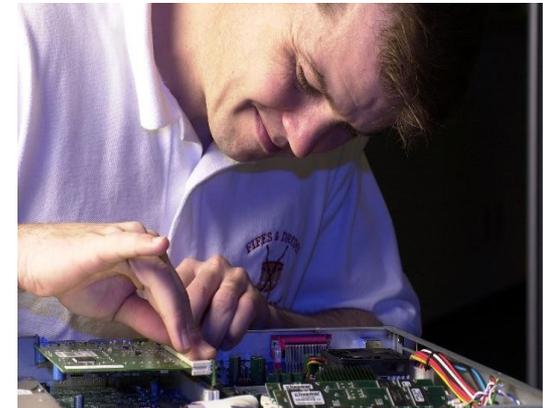
Supporting Transformative Research Through Regional Cyberinfrastructure (CI)

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



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

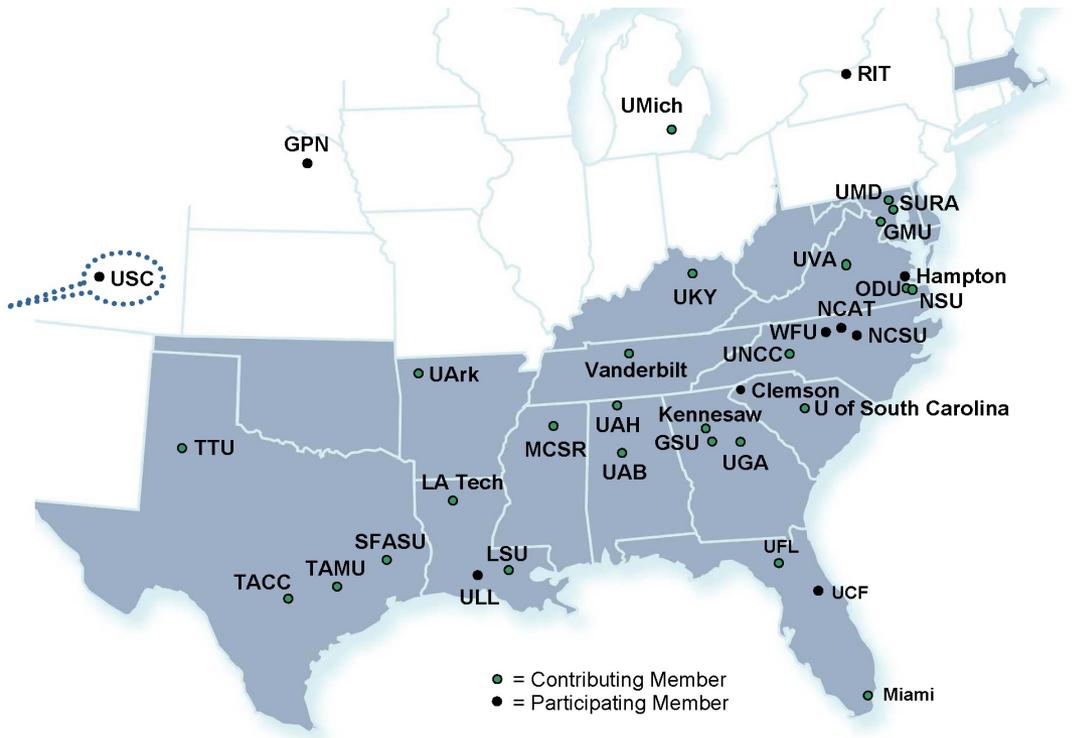
Relations – to formulate and sustain internal and external relations strategy and support for SURA’s scientific and research programs



History of community development in IT

- SURAnet
- Regional Infrastructure Initiative
 - RON development (SoX, MAX)
 - AT&T Dark Fiber Program
- ViDe – Video over IP Development Community
 - Annual ViDe workshops
 - IP Video Cookbook (www.vide.net/cookbook)
- Atlantic Wave - www.atlanticwave.net/
 - an international research network peering fabric interconnecting: US, Canada, Europe, and South America.
- SURAGrid – www.suragrid.org
 - Originated from SURA's participation in NSF Middleware Initiative (NMI) – managed the NMI Test Bed Program
 - An active community of over 30 member institutions developing and deploying a shared, distributed HPC environment

SURAGRID Lowering Barriers for Deploying and Utilizing CyberInfrastructure



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

SURAgriD Governance Structure

- An established and documented governance and decision making process (www.sura.org/programs/sura_grid_gov.html)
 - Elected SURAgriD Governance Committee
 - 9 elected members with staggered 2 year terms
 - An annual election process
- Many active working groups focused on evolving and growing SURAgriD

SURA Corporate Partnerships



Microsoft



- 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

VCL MRI-R2 Proposal Discussion

- ❑ Focus is on MRI-R2 Development Proposal (NSF 09-561) due Aug. 10, 2009
- ❑ Concept conceives of 4 VCL development partners and several test bed, verification sites
 - NCSU, GSU, GMU, UMBC – development sites
 - NSU, U South Carolina, (*need several more*) – Test Bed sites
- ❑ SURA will act as administrative lead
 - Resolves institutional 3 proposal limit
- ❑ NCSU as technical lead
- ❑ Biggest challenge is clearly positioning VCL as a tool (instrument) in support of science domain(s)
 - Meeting planned with NSF OCI Program Director tomorrow
- ❑ IBM VLC support model – Dev. system – Latest Blade Centers
- ❑ ELF/TTP – Outreach partner?
- ❑ Others: Need to identify participants/proposal team now!

Interested? Contact

Sam Averitt, sam_averitt@ncsu.edu

Gary Crane, gcrane@sura.org



The Business Case for VCL Integration of Campus Utility Computing and HPC

Patrick Dreher
RENCI

Sam Averitt
NCSU

VCL Workshop
George Mason University
July 13, 2009



VCL Workshop
Fairfax/V2/Jul-09



Outline

- **Computing value proposition for the campus user and the organization**
- **Brief overview of VCL at NCSU**
- **VCL operational data and statistics**
- **Financial analysis of the data**
- **Observations and remarks**



Computing Value Proposition For Campus User And The Organization

User Requirements

- **Build, save, modify, run custom virtual CEs**
- **Ability for user to install and configure all operating environments and apps**
- **Root privileges (*as required/authorized*)**
- **Secure dropped-session recovery**
- **Time and place independent access**
- **Full functionality via consumer devices and platforms**

University Requirements

- Sustainable funding model
- Comprehensive performance and usage data
- Auditable and open data records
- Software tracking and license compliance
- Enterprise/user defined and managed access policies and priorities
- Transparently interoperable with enterprise identity management system.
- Security level configurable as part of CE load
- Direct high speed access to campus and public networks

Value Proposition for Educational Cloud Computing Architecture

- Deliver services and support to a wide range of users from the novice to the most sophisticated expert researcher
- Deliver a wide-range of course materials and academic support tools to instructors, teachers, professors, and other educators and university staff as part of the academic mission of the institution
- Deliver research level computational systems and services in support of the research mission of the university

Characteristics Within This Value Proposition

- **Benchmarks**
 - **Reliable and fault tolerant**
 - **Flexible to changing user requirements**
 - **Sustainable**
 - **Extensible architecture**
 - **Modular and reusable**
 - **Scalable**
 - **Accessible on-demand or batch**
 - **Secure**
 - **Ability to audit data, process, output**
 - **Efficient and cost effective to operate, service and maintain**

Goal for Cloud Computing in Educational Environment

Develop business models and technical models that can deliver these capabilities within these value propositions and benchmarks and be cost effective



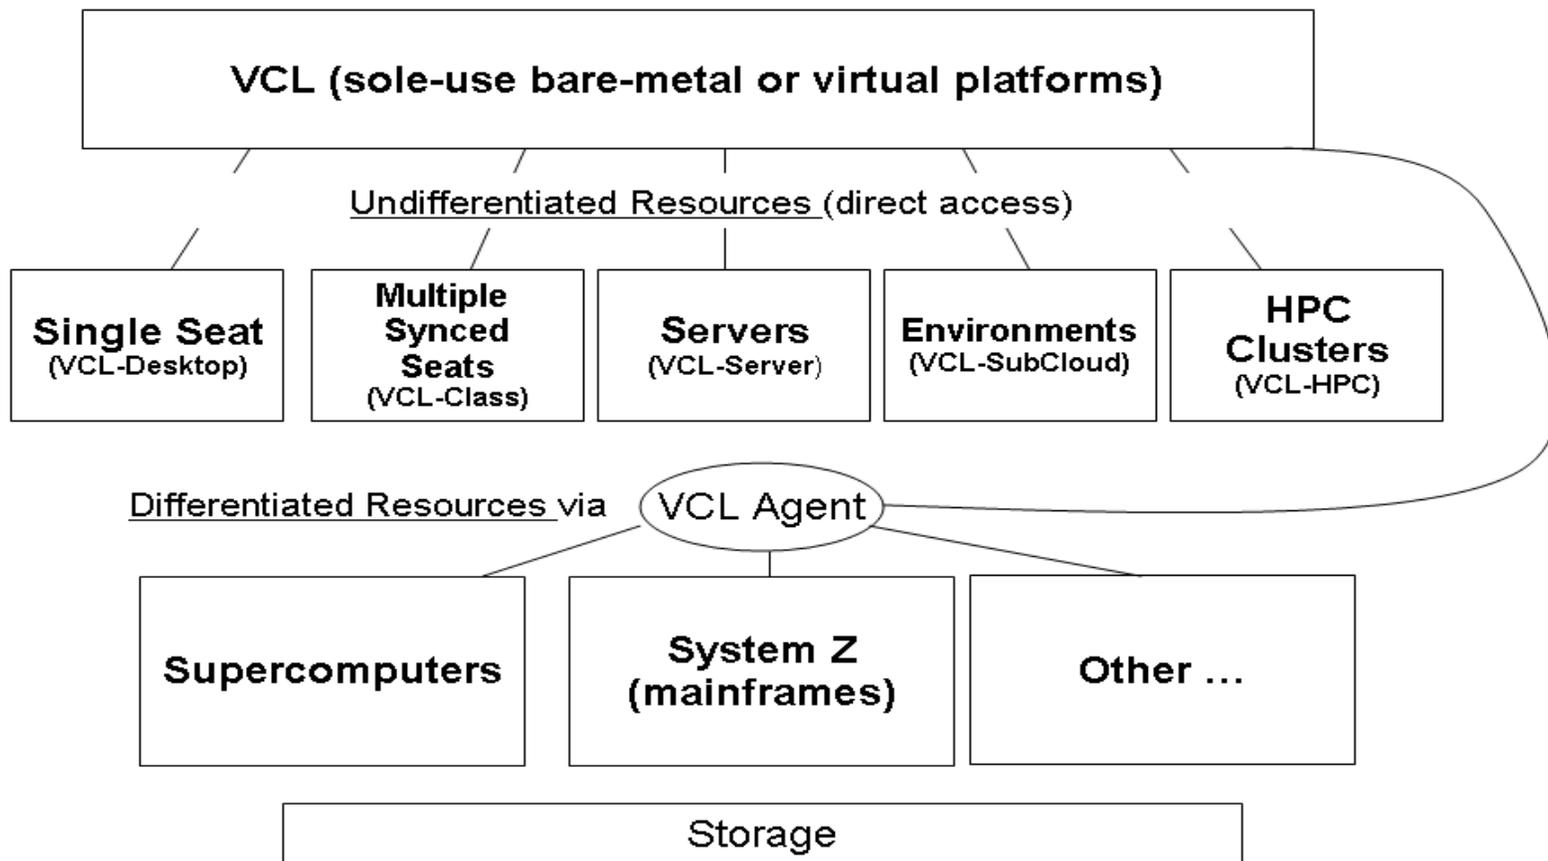
Brief Overview of VCL at NCSU



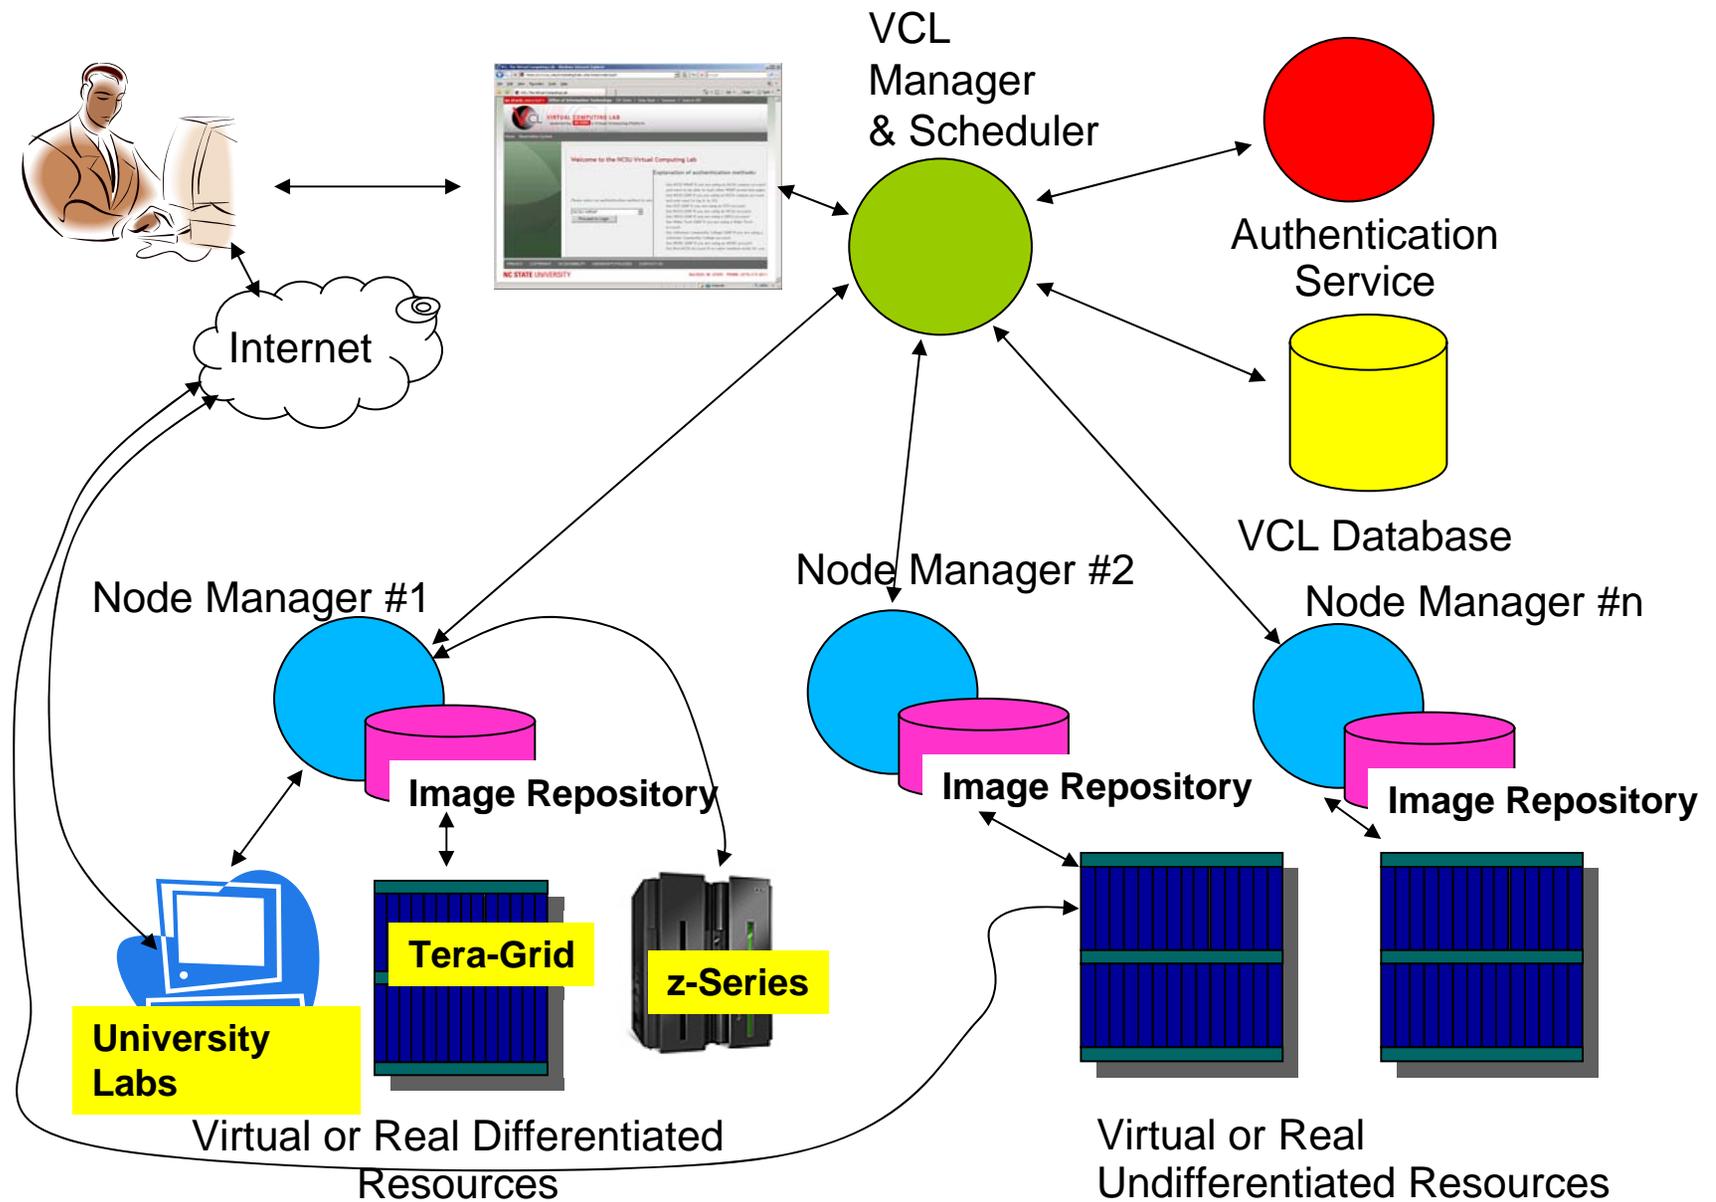
VCL Workshop
Fairfax/V2/Jul-09



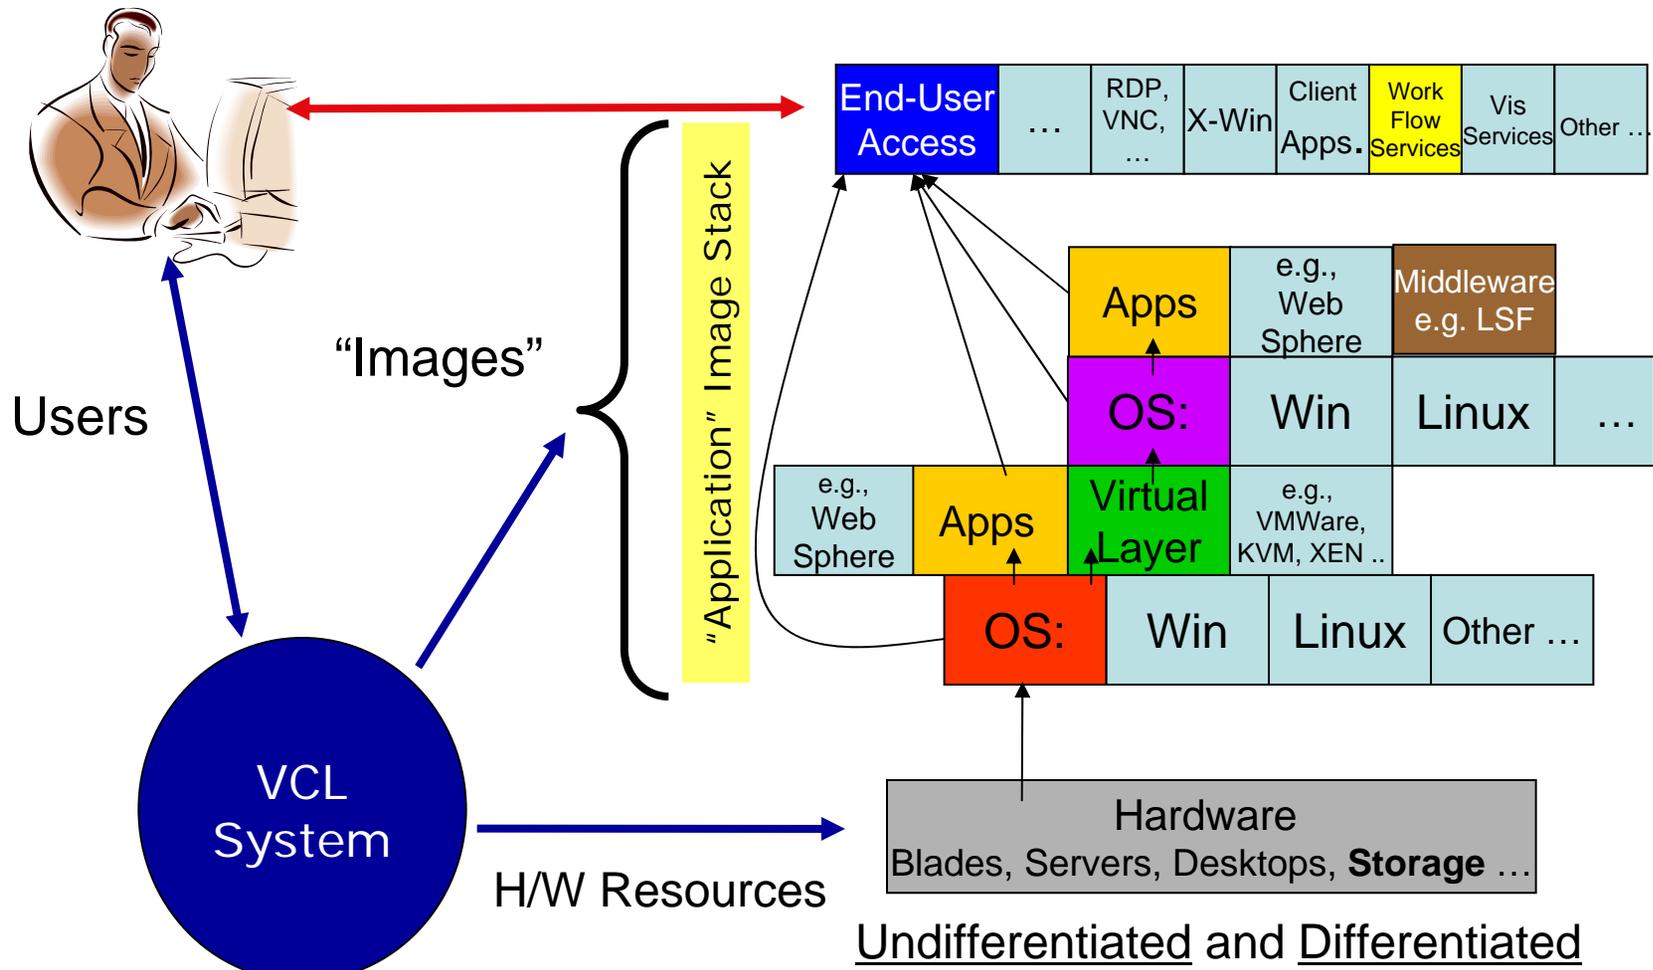
VCL Service Categories



VCL Architecture



Users, Bare-metal and Virtual "Images" and Stacks



xCAT	VCL code	PHP
Apache	MySQL	Etc.

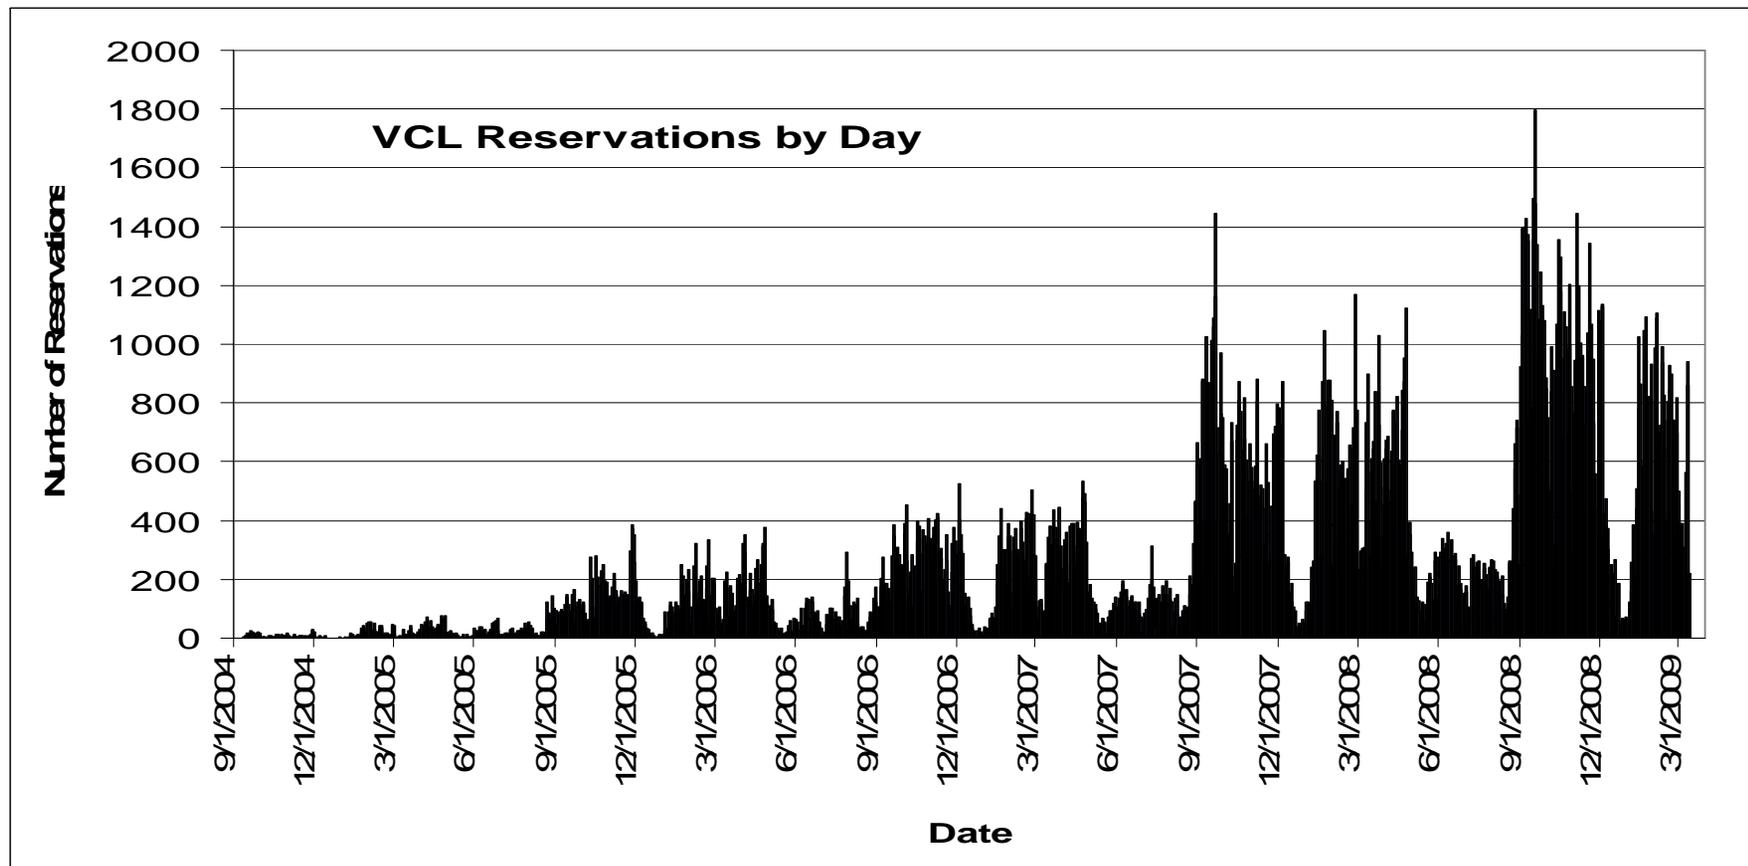
Undifferentiated and Differentiated
Local or Distributed

Differentiators: Reliability, Componentization, Scalability, Economy



VCL Operational Data and Statistics

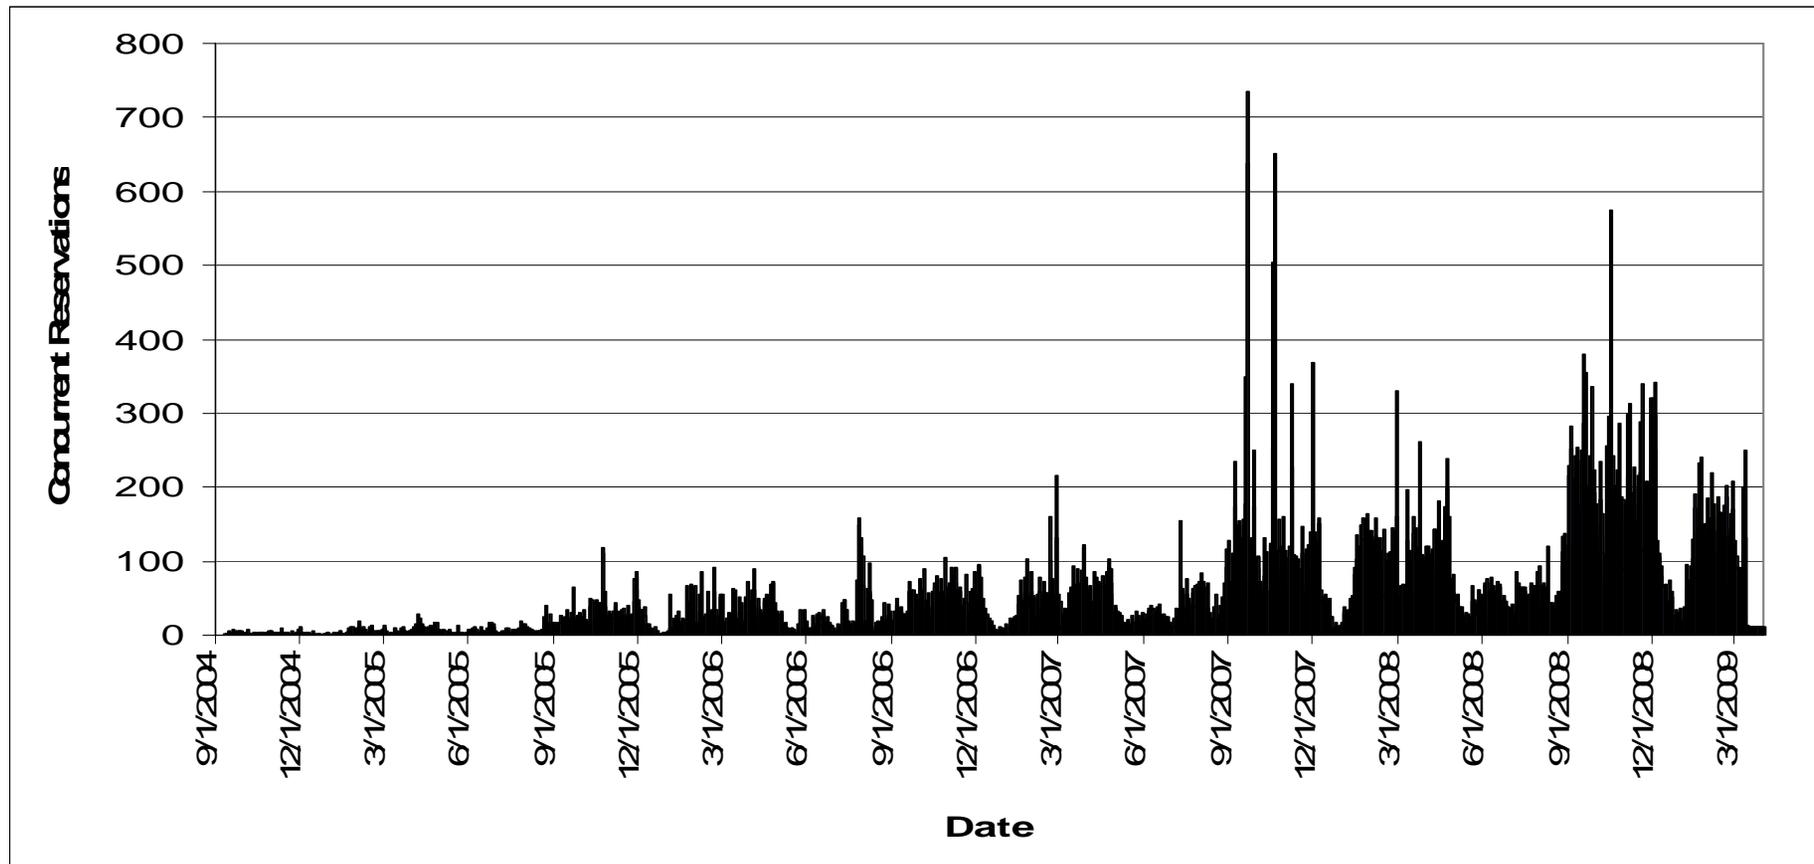
VCL Reservations* at NCSU Sept. 2004 Through Feb. 2009



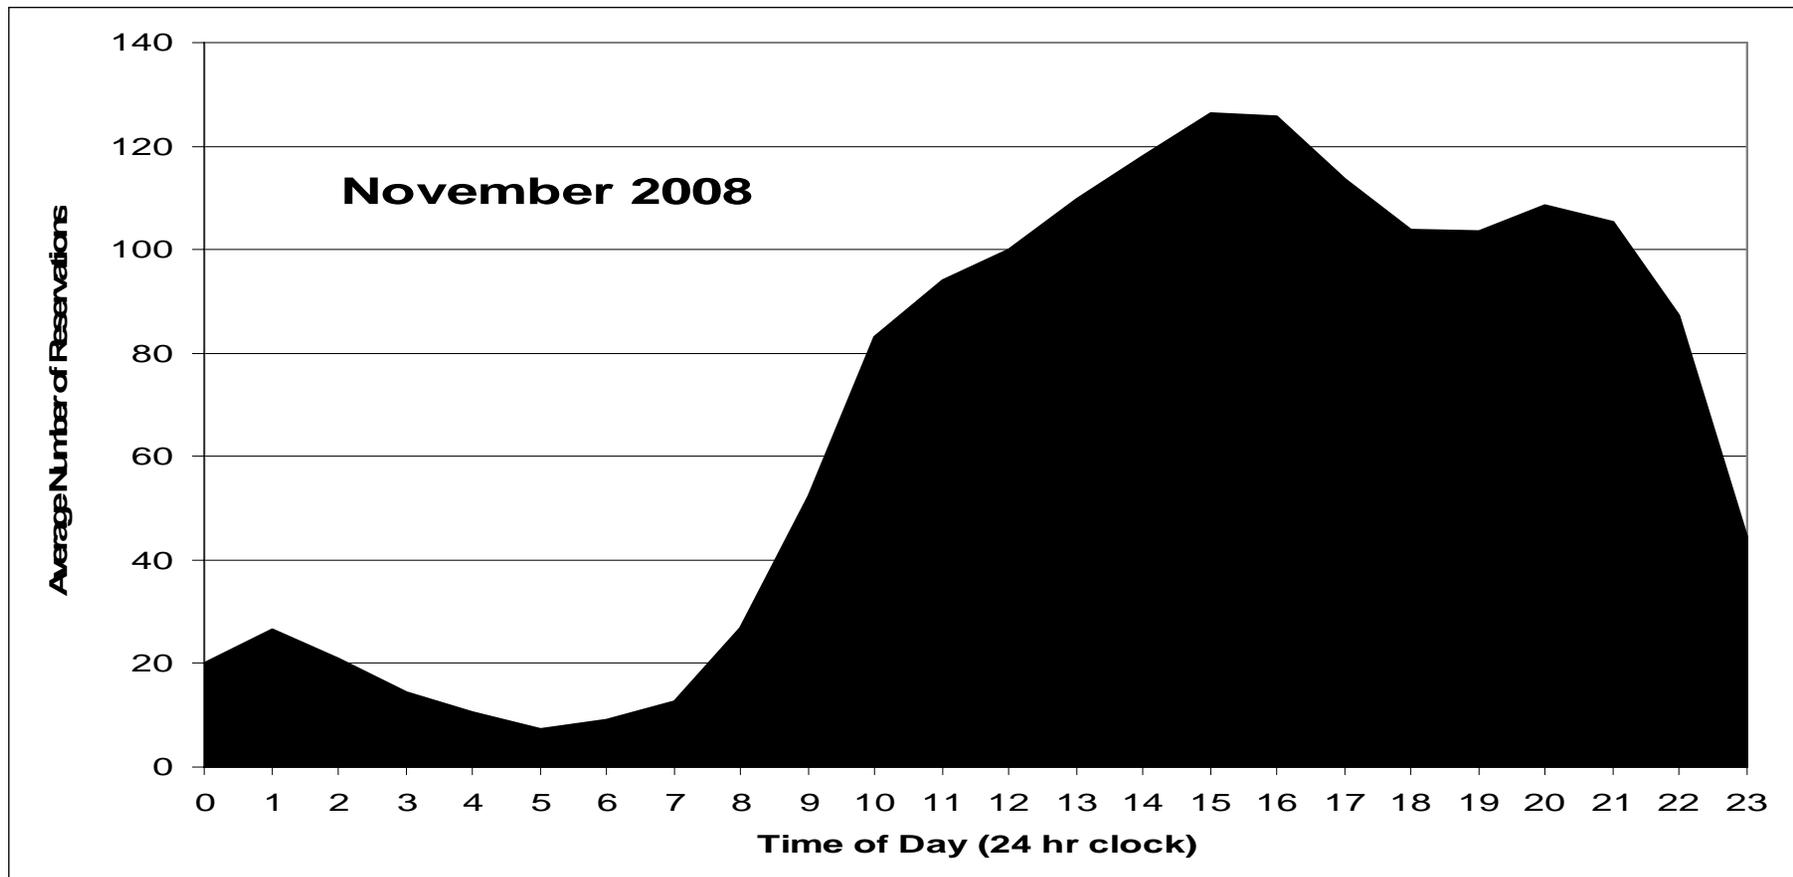
* 352,488 total reservation transactions

VCL Workshop
Fairfax/V2/Jul-09

NCSU VCL Concurrent Reservations Sept. 2004 Through Feb. 2009



Average Daily Number Active NCSU Reservations For November 2008



Observations

- Periodic fluctuations both in total and concurrent number of VCL utility computing reservation requests
- Periodic fluctuations with the average number of VCL utility computing reservations in 24 hour period
- To deliver the value proposition -- tune the number of blades active in VCL so that there are a minimum number of on-demand reservations that fail or cannot be satisfied



Financial Analysis of the Data

Identify Costs Elements To Maintain This Utility Computing Service

- Multi-year life cycle implies a heterogeneous mix of hardware
- Depreciation schedules
 - Computers
 - Network
 - Facility infrastructure
- Personnel costs
 - Inelastic – one-time scalable activities
 - Elastic – e.g. volume dependent support costs
- Operational costs (supplies, maintenance, etc)

Non HPC VCL Hardware for Guaranteed Utility Computing

- If virtualization factor is 1 (one)
- Measure (non-HPC statistics)
 - Number of blades required to be concurrently active in the NCSU cloud to guarantee VCL access anytime in last 12 month period? (as many as 359)
 - Number of reservations (161,850)
 - Total VCL utility computing hours (358,154)
- Load factor - $\text{hours}/\text{max \# of concurrent blades}$ (997)

Parameters for Operational Cost/Service Hour

- **Guaranteed reservation**
 - Dedicated new blade (3 yr depreciation)
 - Older blade with 5 yr depreciation
 - Mix of blade hardware
- **non-HPC reservations can allow option for blades to be virtualized**
- **Classroom and light duty non-HPC reservations can handle ~ 10 or more virtual sessions per blade**

Operational Cost/Service-Hr For Non-HPC VCL Reservations*

Depreciation	HPC blade	Mixed blade	Older blade
3 year	\$1.60	\$1.25	\$1.03
10x Virtual	\$.16	\$.13	\$.10
Mix	\$1.41	\$1.06	\$.84
10x Virtual	\$.14	\$.11	\$.08
5 year	\$1.27	\$.92	\$.70
10x Virtual	\$.13	\$.09	\$.07

*Guaranteed reservation dedicated blades

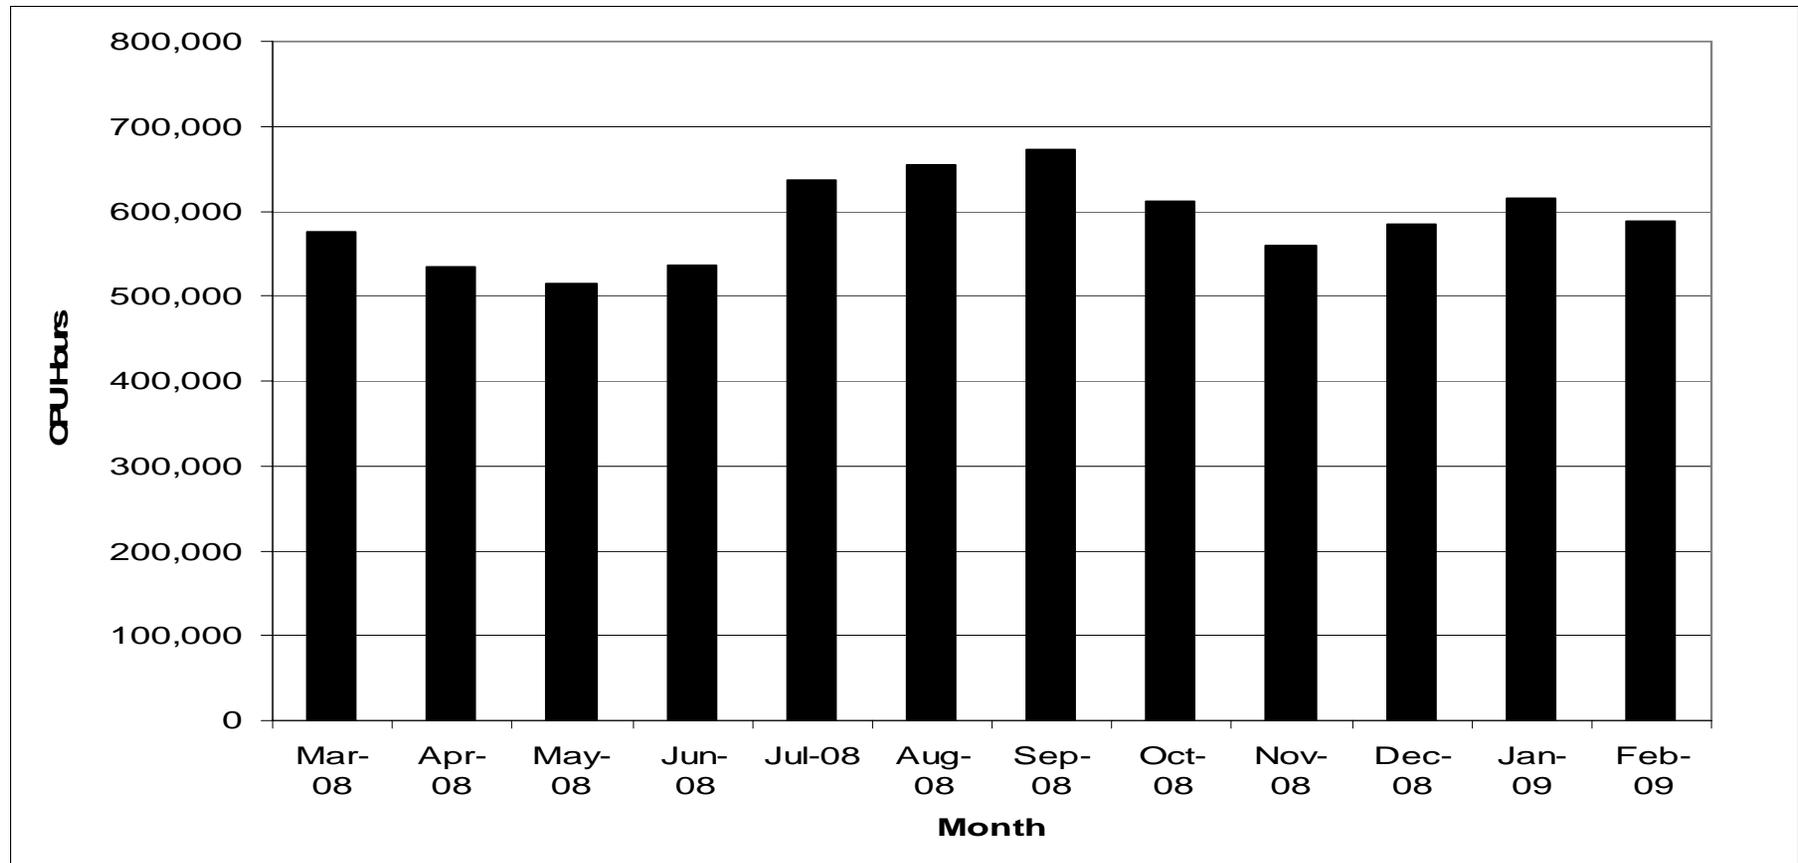
Multi-Use Cloud Utilization

- **Guaranteed reservation policy requires ~ 359 blades dedicated to non-HPC use**
- **There are 8,760 hours in a year**
- **10% operational overhead implies there are 7,884 useable hours/year/blade**
- **Guaranteed availability of ~ 359 blades is only needed at peak times during the year**
- **Tuning the cloud with load factor of ~ 1000 to guarantee peak reservation requests results in only 10 to 12% average cloud utilization.**
- **Only covers operational costs - very expensive with underutilized/wasted assets!**

NCSU VCL Cloud Characteristics

- **Non-HPC usage patterns**
 - Typically 1000 to 1200 blades are active in the NCSU cloud
 - Several hundred are testing and research blades
 - Several hundred are powered down
 - For example, on July 11, 2009 (a summer Saturday)
 - 528 in HPC mode (100% blade utilization)
 - 457 in non-HPC mode (16% utilization)
 - **Range of utilization of non-HPC blades**
 - ~ 100% (school day, around midday)
 - ~ 15% during the day
 - As low as 5% during the night.
- **Does campus HPC usage follow utility computing usage patterns?**

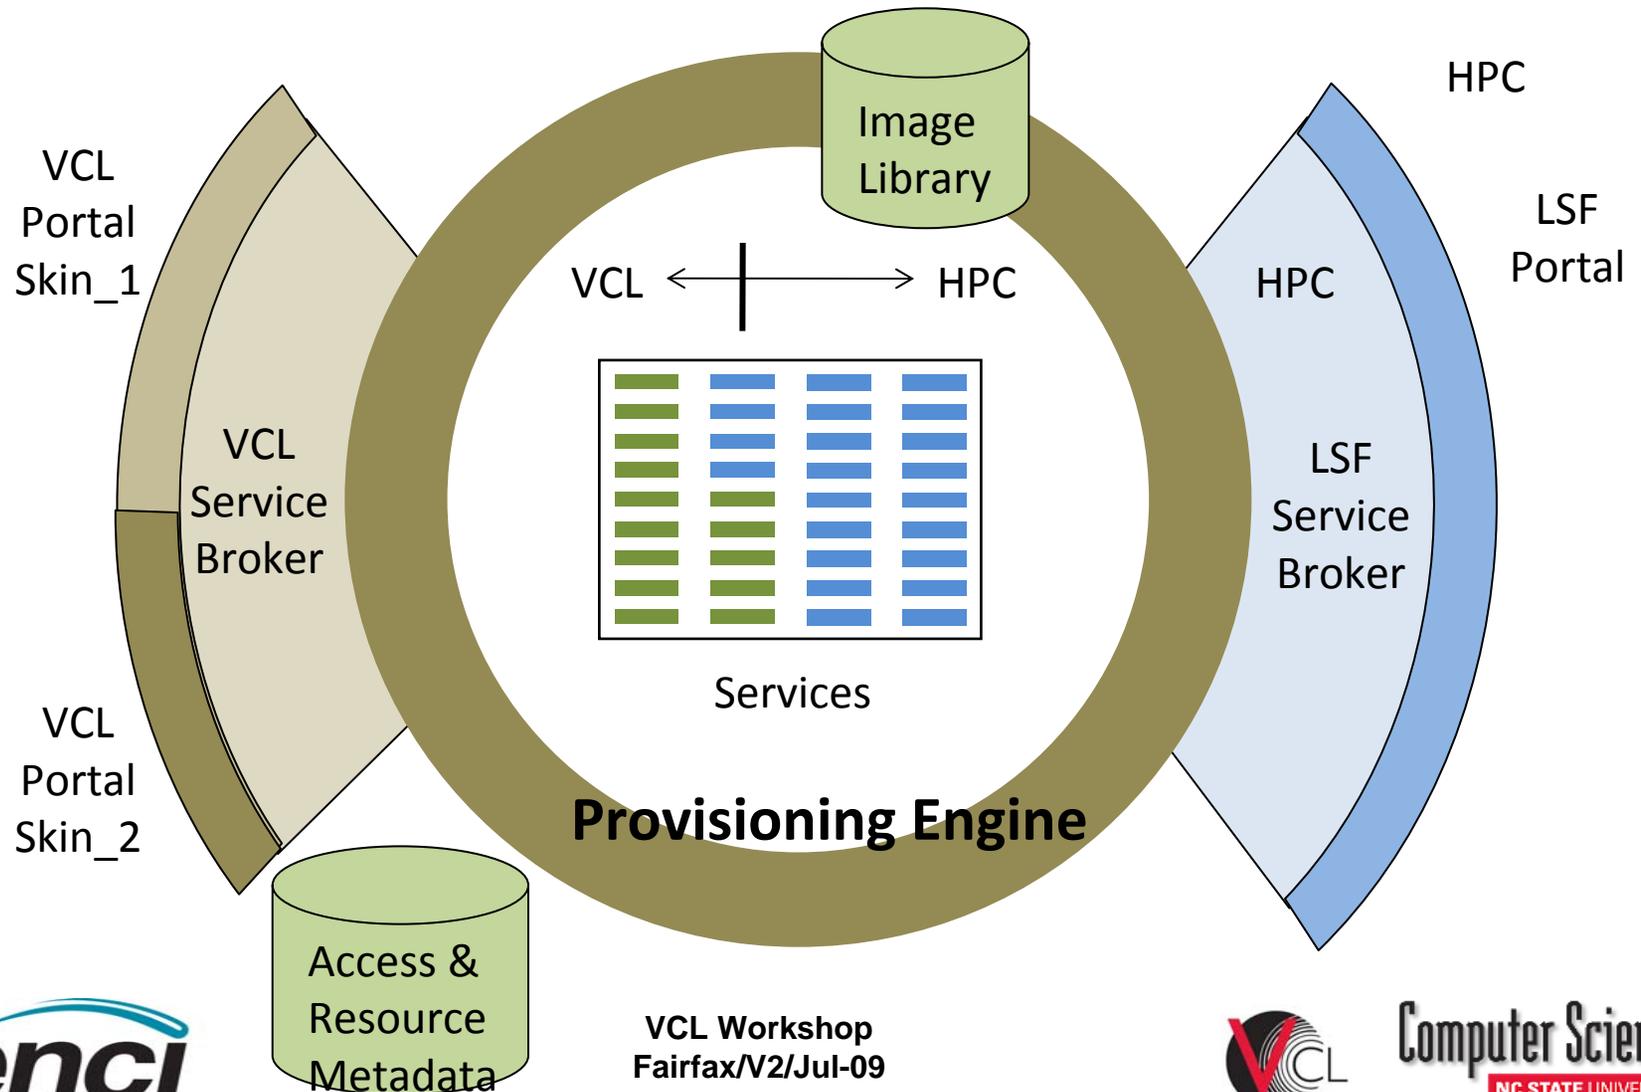
VCL-HPC CPU Hour Usage March 2008 – February 2009



HPC and Non-HPC Separate Total Cost Profiles

- **NCSU VCL cloud usage
(Jul 11, 2008 – Jul 11, 2009)**
 - > 7,000,000 HPC CPU hours
 - 358,154 non-HPC blade hours (~ 500,000 CPU hours)
- **HPC total cost profile**
 - Virtualization factor of 1
 - Yearly NCSU HPC cost ~ \$1,400,000
 - Total cost per HPC service hour is ~ \$.20
- **Non-HPC total cost profile**
 - Total annualized cost is about \$1,000,000
 - ~ \$2 per non-HPC CPU service hour

Goal is to Design a VCL Provisioning Engine



Combination of VCL Utility and HPC Cloud Computing

- Capabilities and functionalities can be combined and offered to both utility computing and HPC researchers
- **Combination mode**
 - Annualized total cost is about \$2,000,000 to deliver ~ 7,500,000 combined CPU hours
 - Savings recognized through common resources and personnel costs
 - Total cost of running combined non-HPC and HPC services is ~ \$0.27 per CPU service hour
- **Virtualization can decrease total VCL cost to ~ \$.05-.10 per hour**

Cost Savings and Economies of Scale

- **Options for utilization optimization**
 - HPC/non-HPC use combination
 - Sharing among different time-zones
- **Utilize time sharing, virtualization and integration of both non-HPC and HPC resources**
 - Up-front reduction in the total cost required to individually maintain both of these services separately
 - Increased average utilization of the total available VCL resources
 - Translates to lower overall cost/CPU-hr for both utility and HPC computing



Observations and Remarks

Non Financial Business Case Educational Needs - Requirements

- Computing environments provide a component of creative incubator for learning
- Historical silo strategy response is unsustainable
- Education has diverse
 - Missions
 - Requirements
 - Pedagogies
 - Cultures

Non Financial Business Case in the Educational Space

- **Key non-financial business case issues**
 - **Empower both the individual teaching and learning students, faculty, staff and researcher**
 - Build what they need
 - When they need it
 - Configured in a manner needed to get work done
 - **Trust**
 - Remove the “external gatekeeper” (no basis for trust)
 - Lack of trust is the single largest factor inhibiting adoption of cloud computing on campuses and institutions of higher education

Perspectives for the Future

- **Cloud Computing interactions within the educational environment**
 - **Impacts on students and researchers**
 - **Impacts on university**
- **Overcoming barriers to adoption**

Challenges to Adoption

- List of challenges from easiest to more challenging

- | | |
|--------------------------|-------------|
| • Physical distribution | Technology |
| • Technology Evolution | Portability |
| • Size – Uncertainty | Robustness |
| • Customer Awareness | Education |
| • Ownership Distribution | Sociology |



Questions

VCL Experience and Some Suggested Best Practices: NC State

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

Sam Averitt, Vice-Provost for Information Technology, Director of the Center of
Excellence for Cloud Computing

Aaron Peeler, VCL Program Manager

Henry Schaffer, Professor Emeritus Genetics and BioMathematics, and VCL
Special IT Projects

Eric Sills, Assistant Vice-Provost for Research Computing

Sarah Stein, Professor of Communications, VCL Faculty Outreach and
Collaboration

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



Outline

- Briefly about VCL and Clouds
- Planning and set-up
- Scaling up
- Configurations
- Consortium support infrastructure – Level 1, 2, 3 and 4 support



<http://vcl.ncsu.edu>

Powering Clouds

Virtual Computing Laboratory is Open Source

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

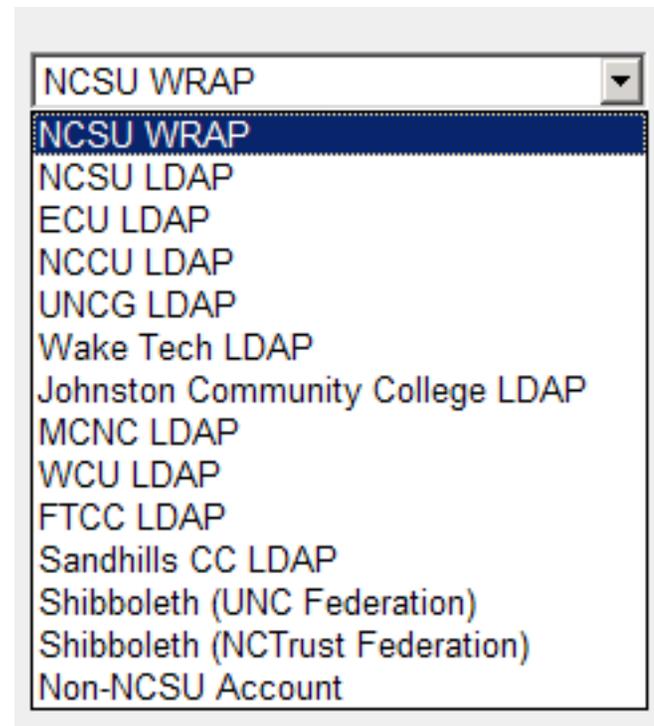
Current NC State University VCL installation:

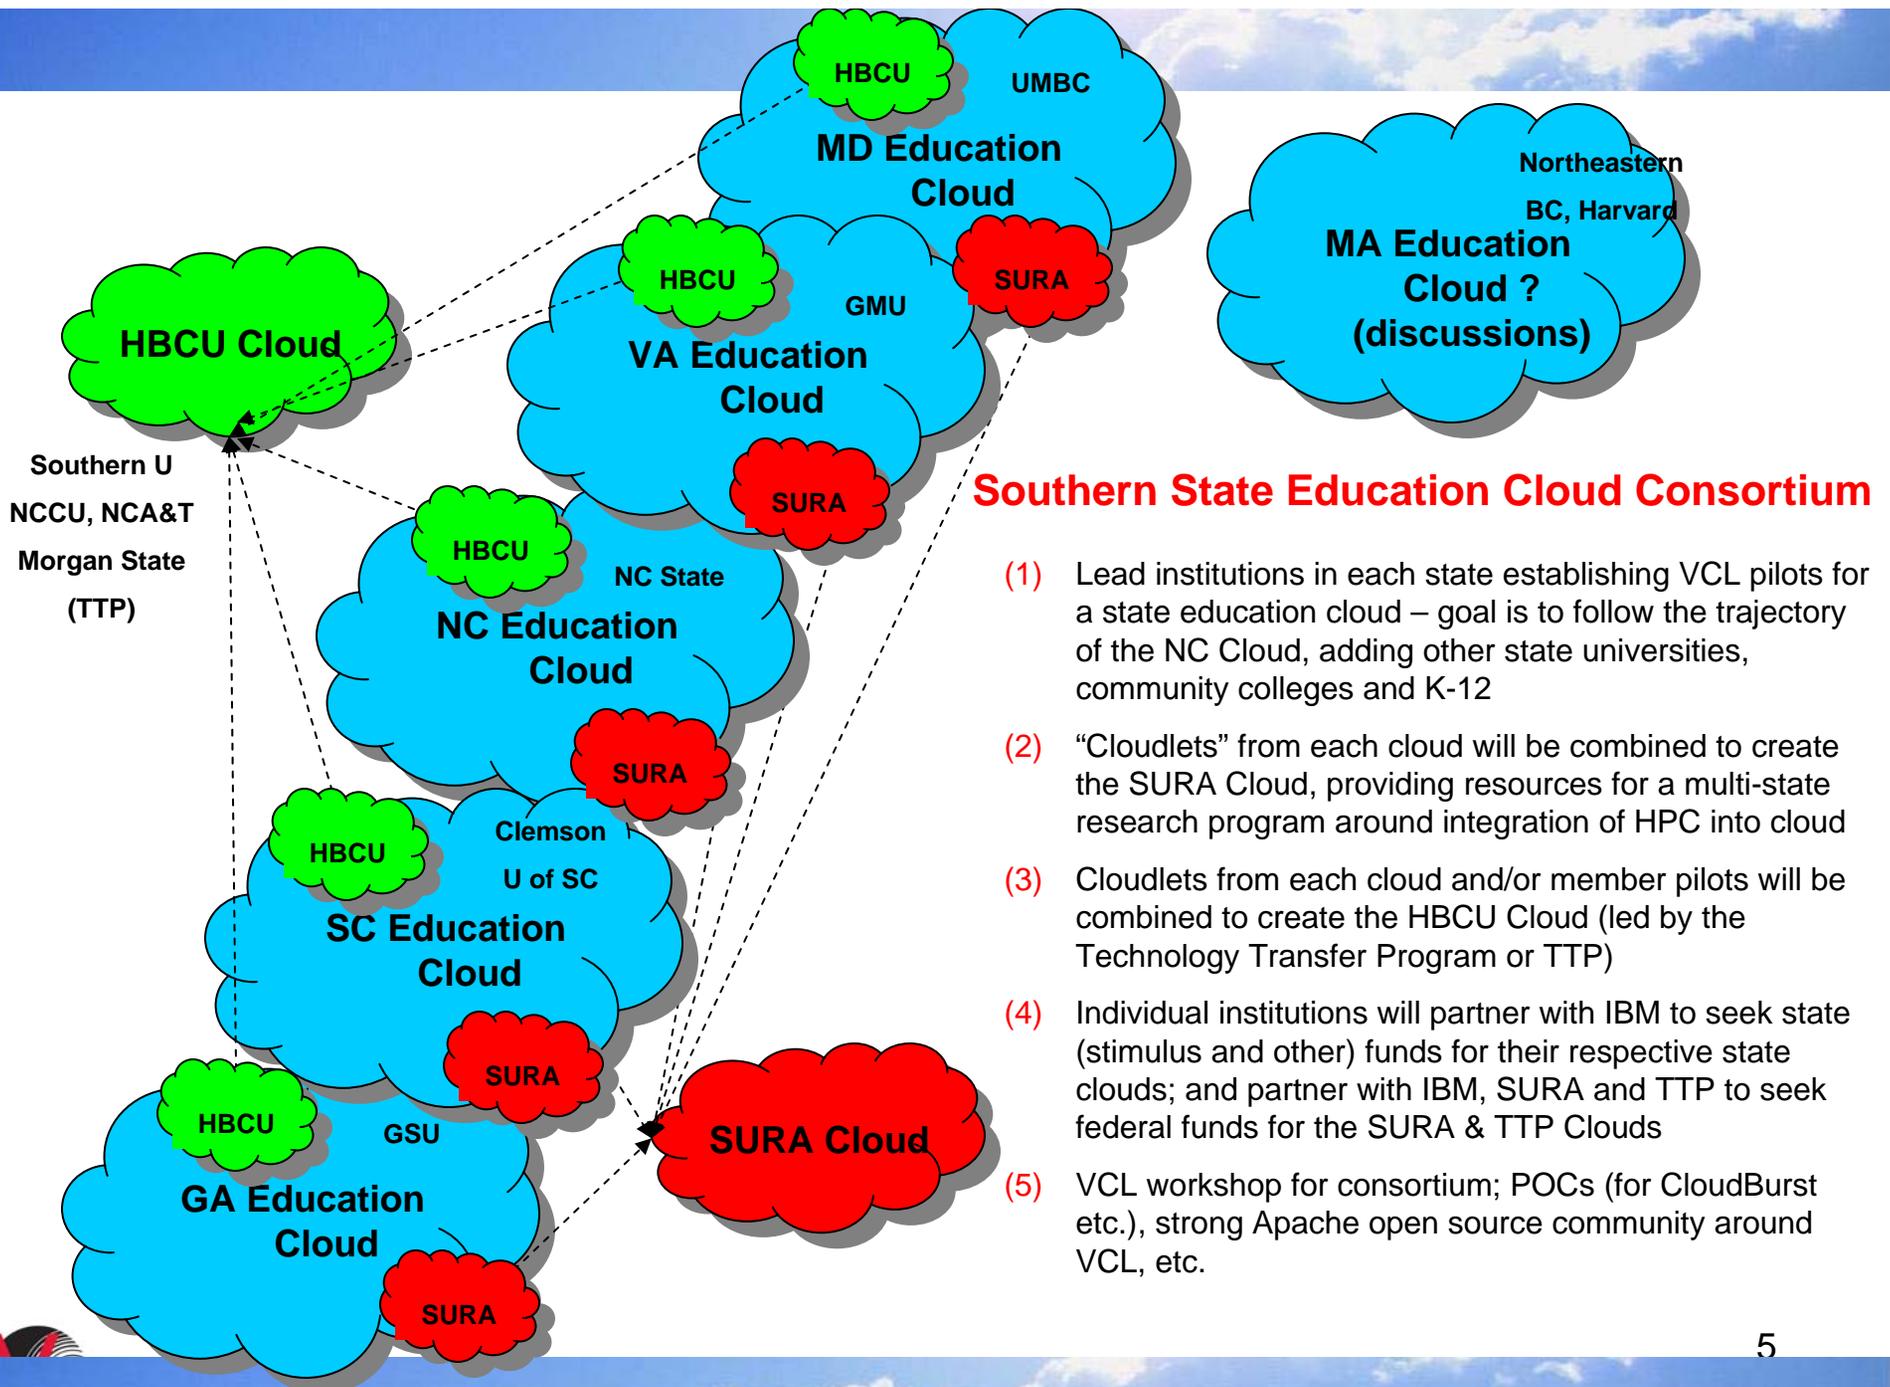
1. cca 1,800 blades, only about 1200 or so are in production
2. cca 500 to 600 in non-HPC mode, about the same in HPC mode, and one to two hundred in various test-beds
3. open to 30,000+ NCSU students and faculty, different pilots and partner accounts, through Shibboleth all UNC System campuses have access to VCL (cca 250,000 students).



Community

- NCSU (in 4 data centers, three on campus, one off campus at MCNC)
- Numerous partners and pilots
- A number of stand-alone facilities
 - Duke, ECU, GMU, RENCI, UNC-CH, NCCU, India, Old Dominion, Western Carolina, Kannapolis, etc.

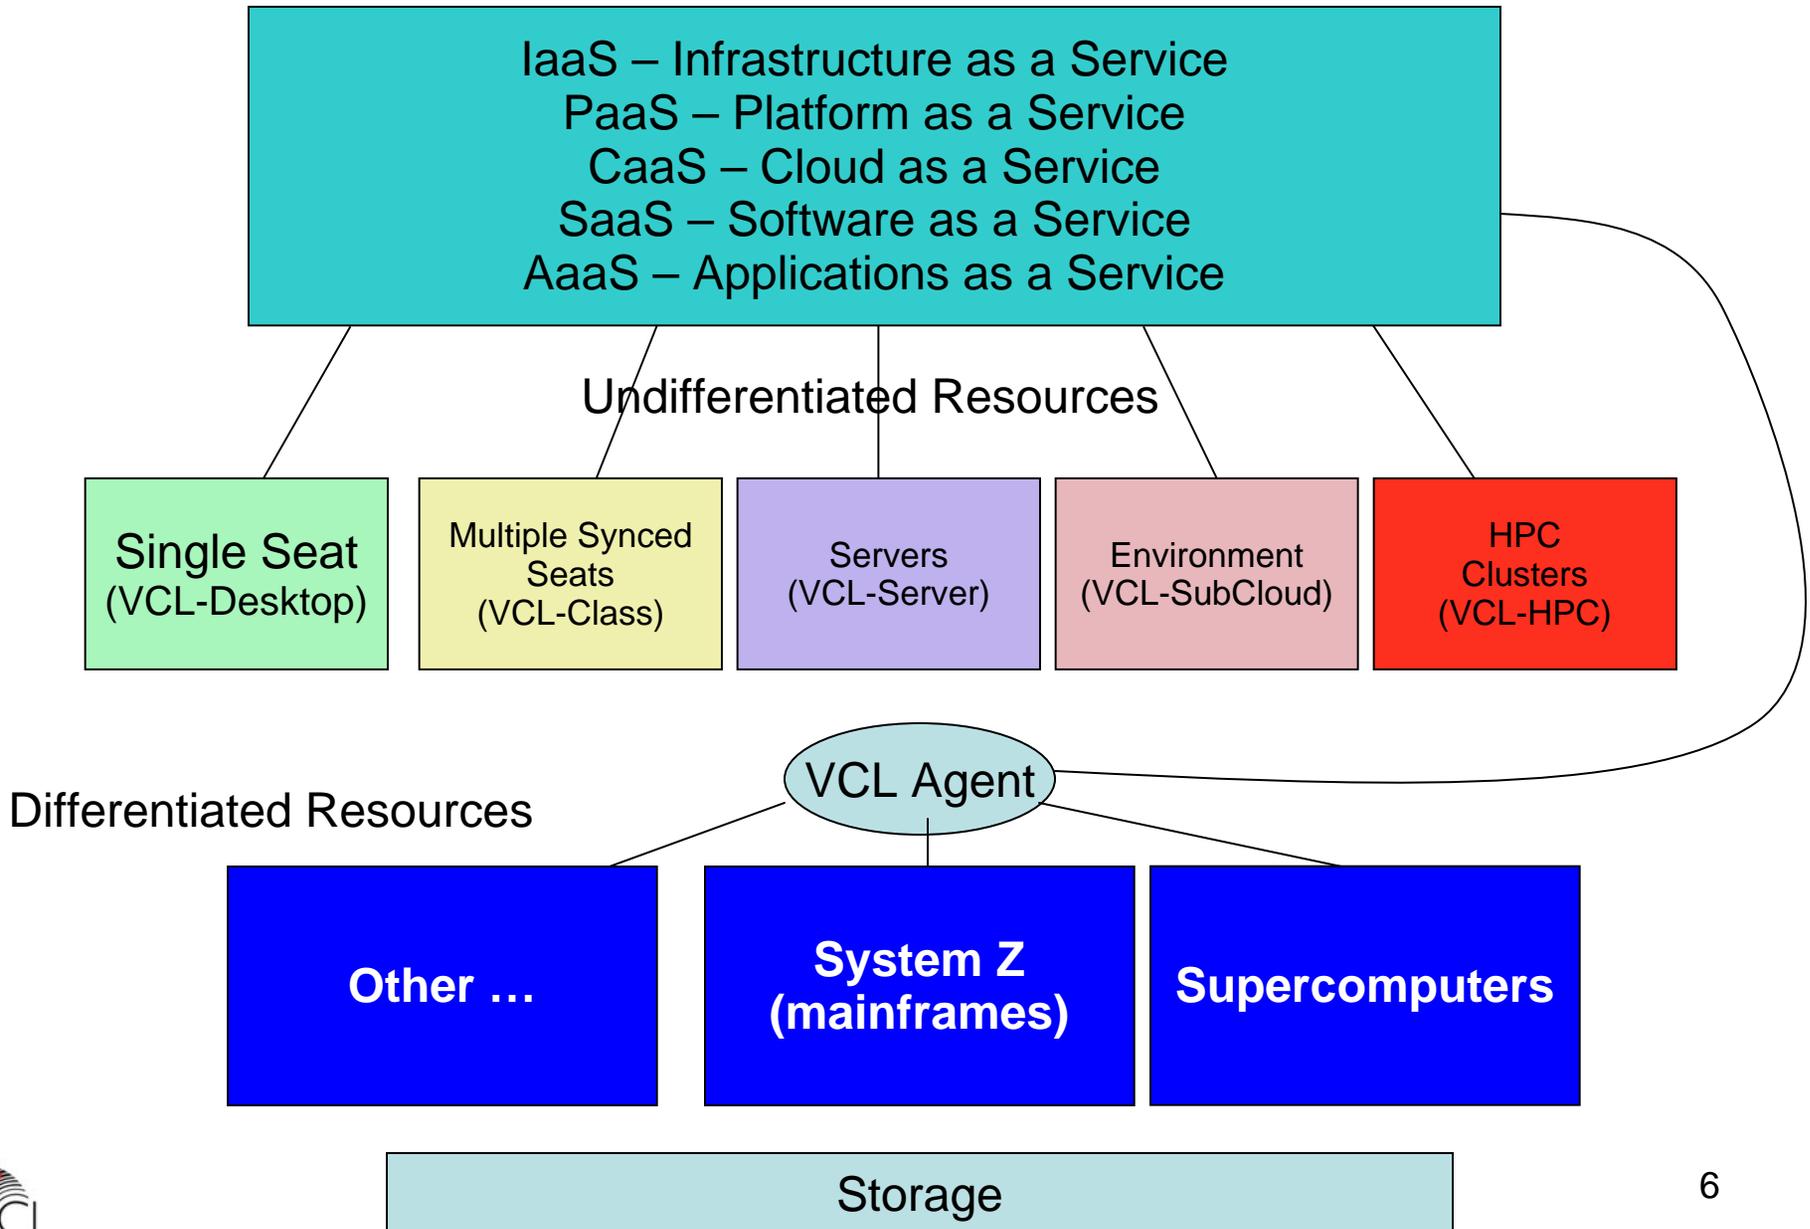


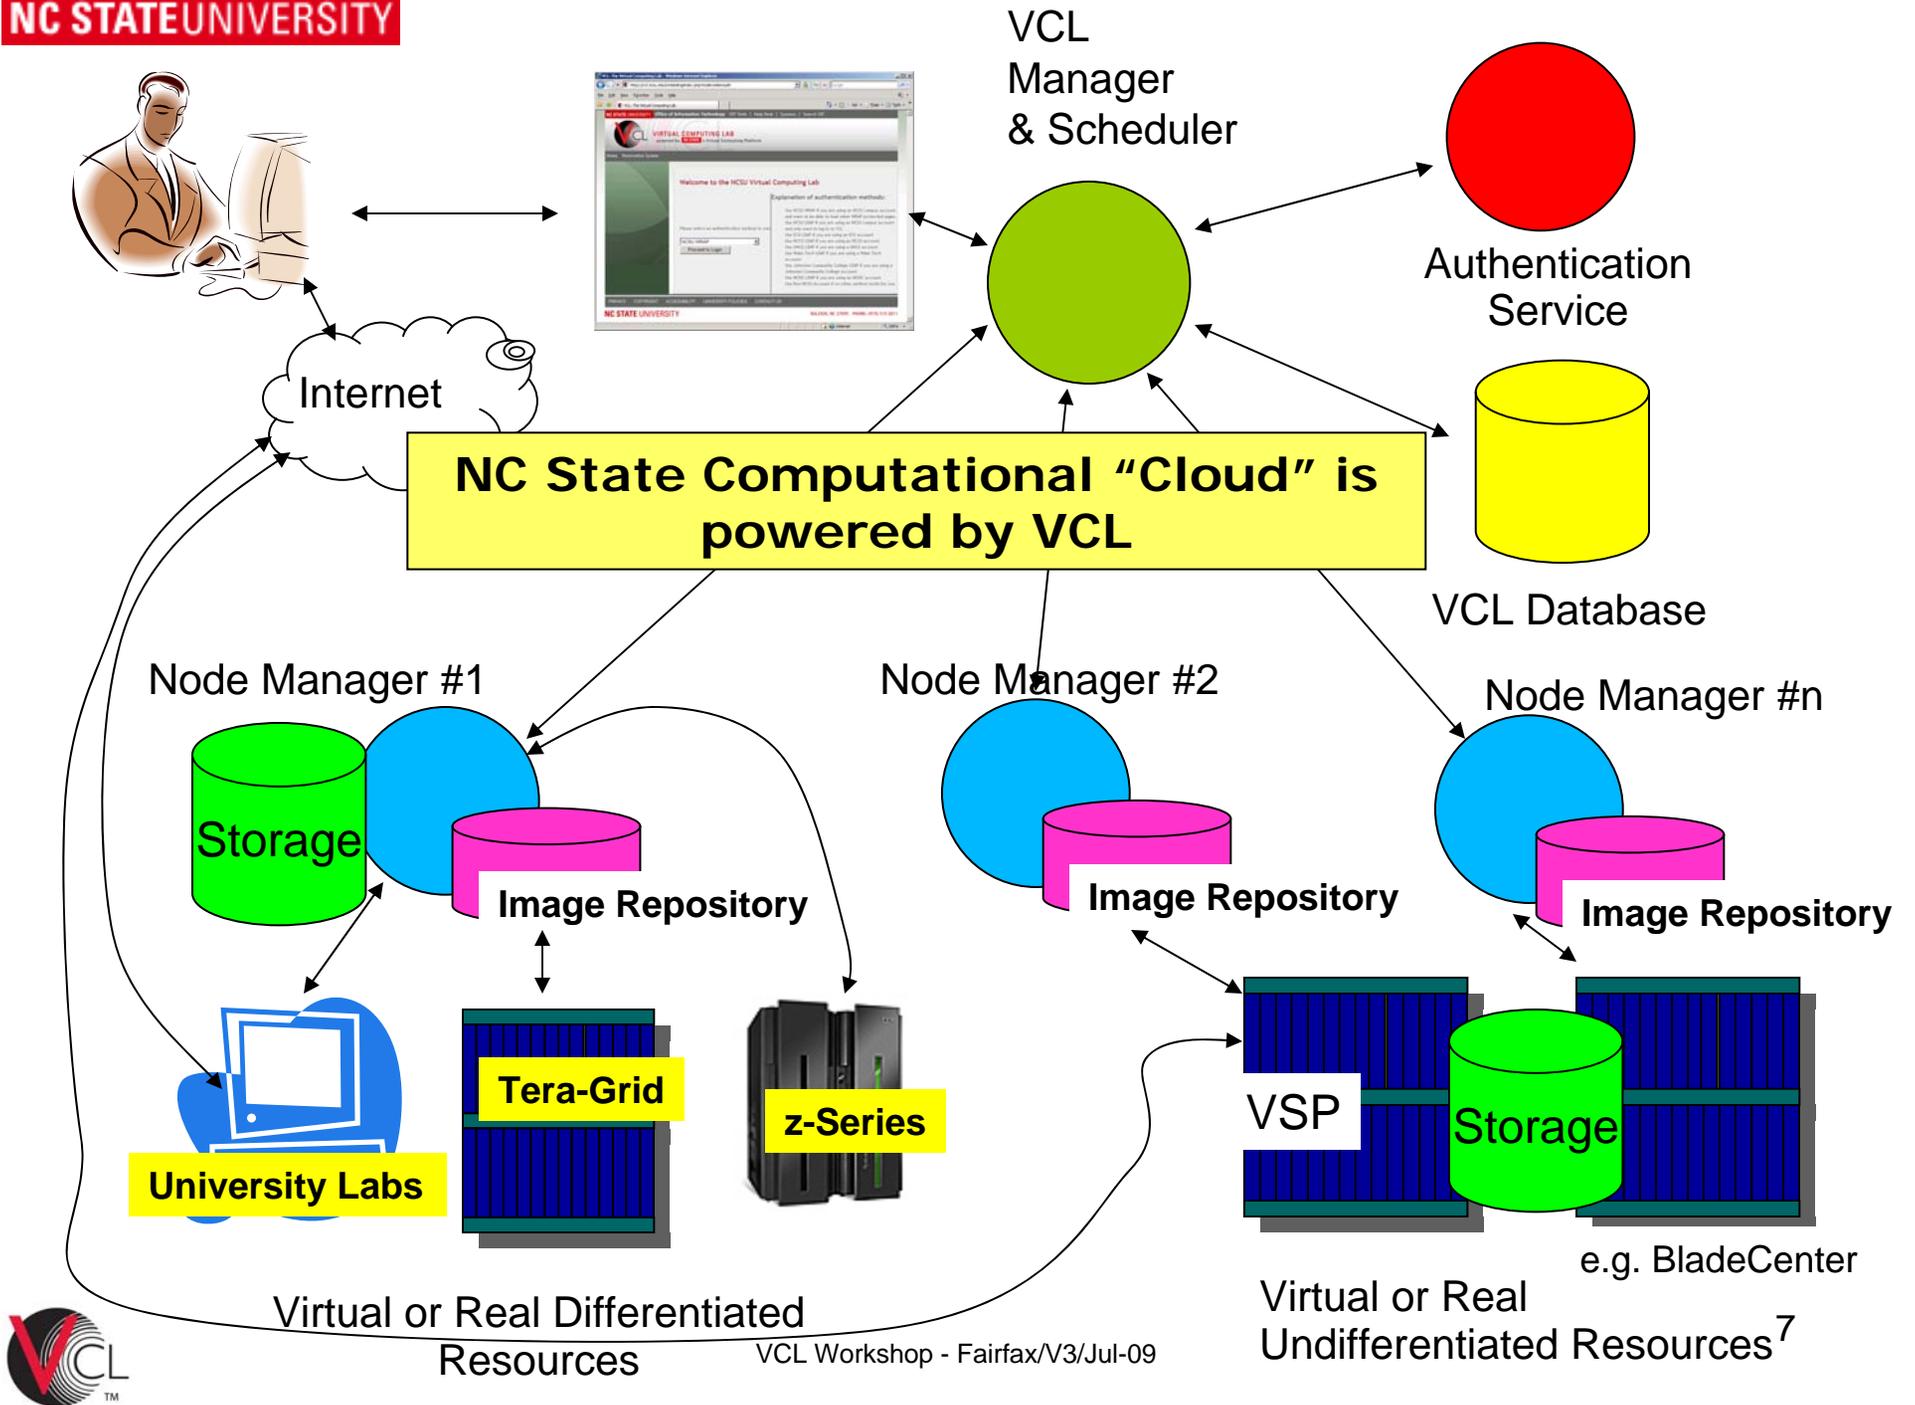


Southern State Education Cloud Consortium

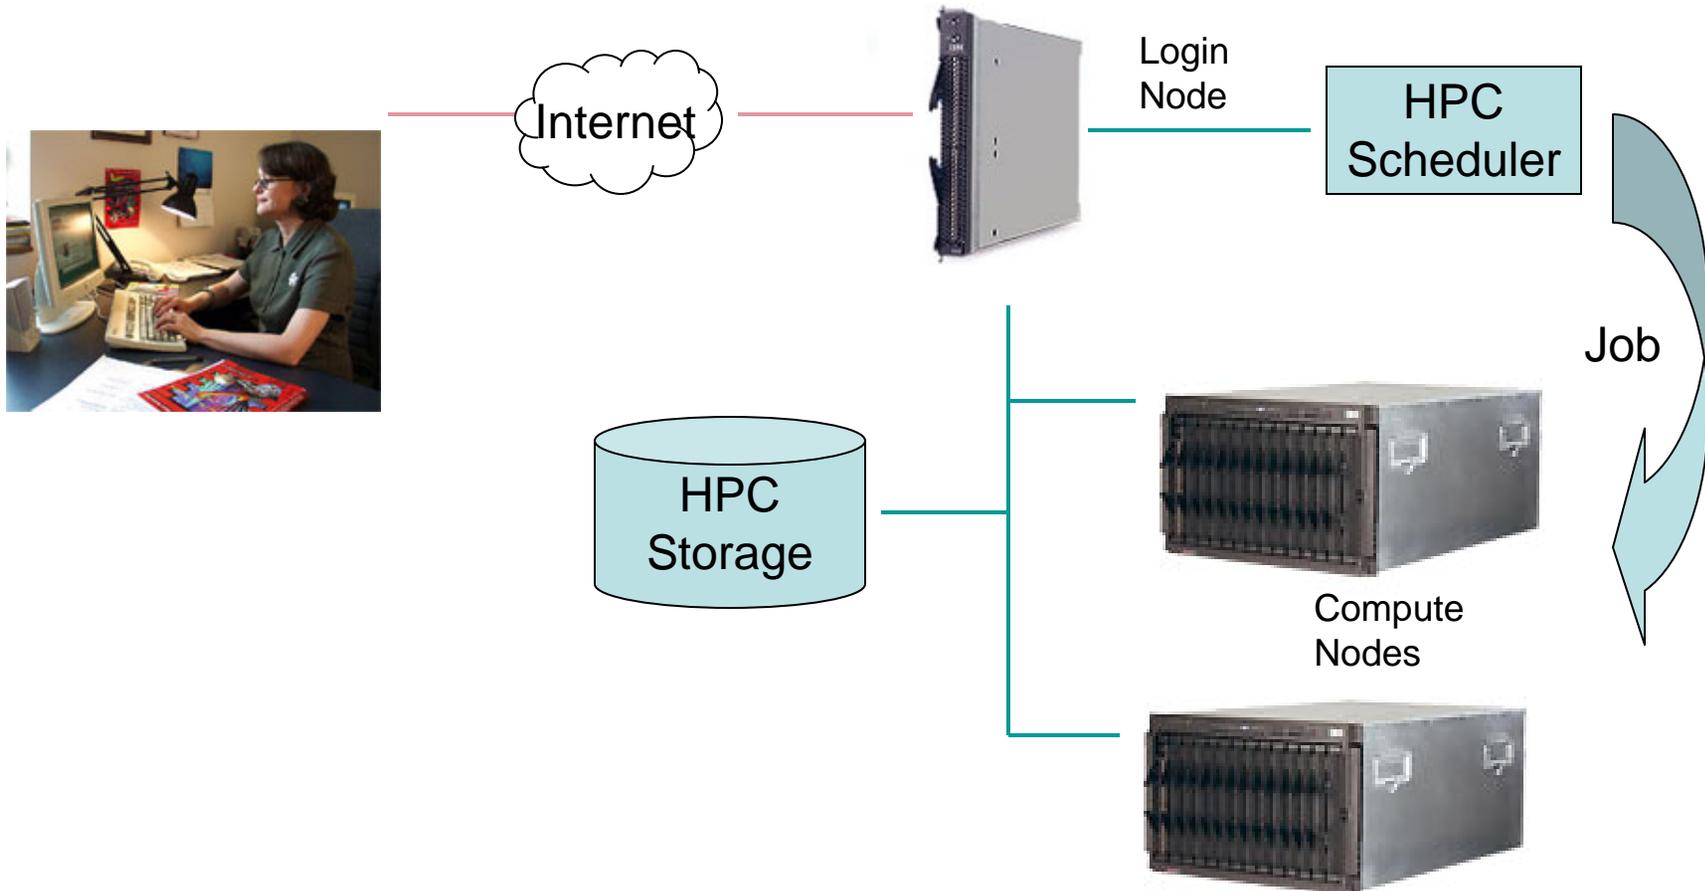
- (1) Lead institutions in each state establishing VCL pilots for a state education cloud – goal is to follow the trajectory of the NC Cloud, adding other state universities, community colleges and K-12
- (2) “Cloudlets” from each cloud will be combined to create the SURA Cloud, providing resources for a multi-state research program around integration of HPC into cloud
- (3) Cloudlets from each cloud and/or member pilots will be combined to create the HBCU Cloud (led by the Technology Transfer Program or TTP)
- (4) Individual institutions will partner with IBM to seek state (stimulus and other) funds for their respective state clouds; and partner with IBM, SURA and TTP to seek federal funds for the SURA & TTP Clouds
- (5) VCL workshop for consortium; POCs (for CloudBurst etc.), strong Apache open source community around VCL, etc.

VCL Cloud Services



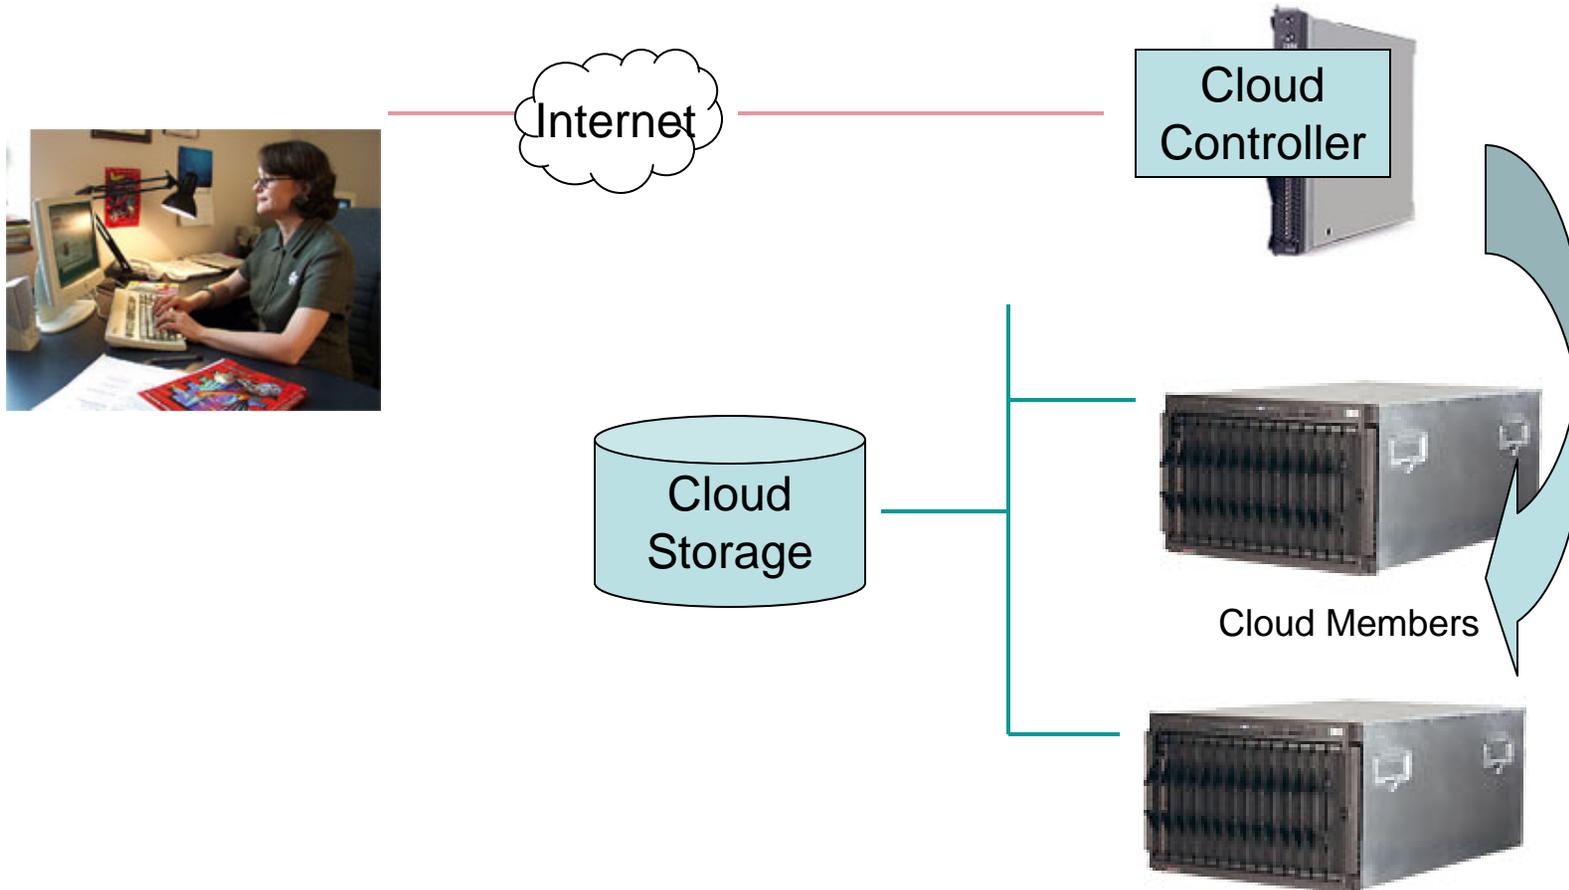


Typical HPC Use of VCL



Typical "Cloud" Use of VCL

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



Security

- Access to reservation page via HTTPS and appropriate authentication (LDAP, Shibboleth, other)
- IP lock on remote location
- Required confirmation of reservation
- Time-out
- Access to Windows images through one-time passwords
- Access to Linux images through campus-level authentication
- VLAN separation of non-HPC and HPC resources



Storage

- The image “knows”
- Enterprise-level storage (AFS-based at NCSU) – appears as K-drive or NFS attached storage
- Access to any remote storage (via SSH-based file transfer software)
- Local storage on blades
- Remote storage (disks on access client, memory key)
- HPC NFS storage (scratch, home and off-line)
- Storage “images”



Cost Factors

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



Planning and Set-up: Operational Profile

- At NCSU over 600 images
- About 150 are in use in a semester
- non-HPC and HPC modes
- A max of 300+ concurrent non-HPC reservations

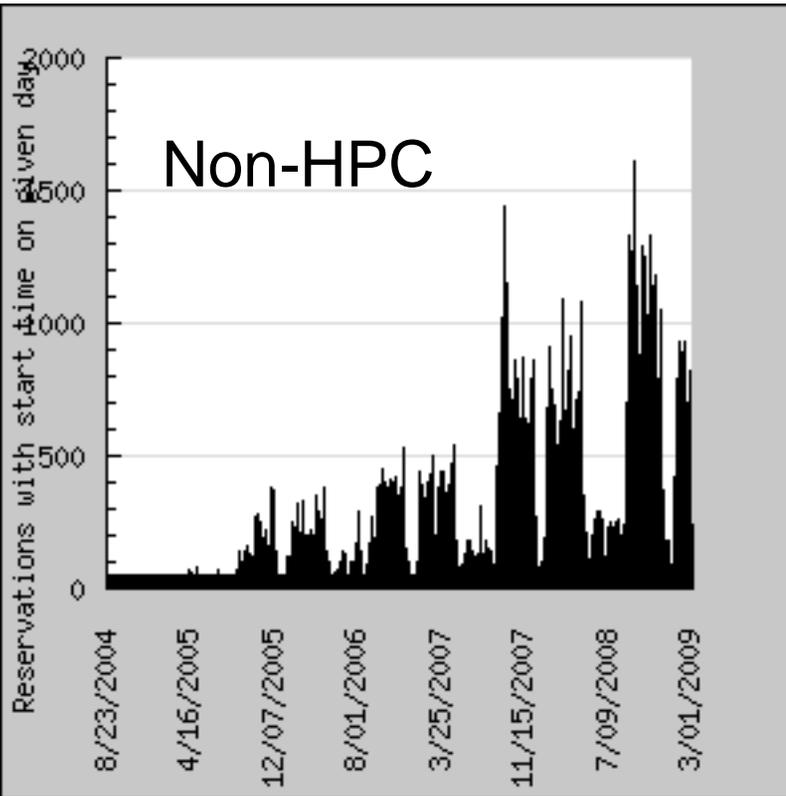


VCL Usage 2004-2008

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

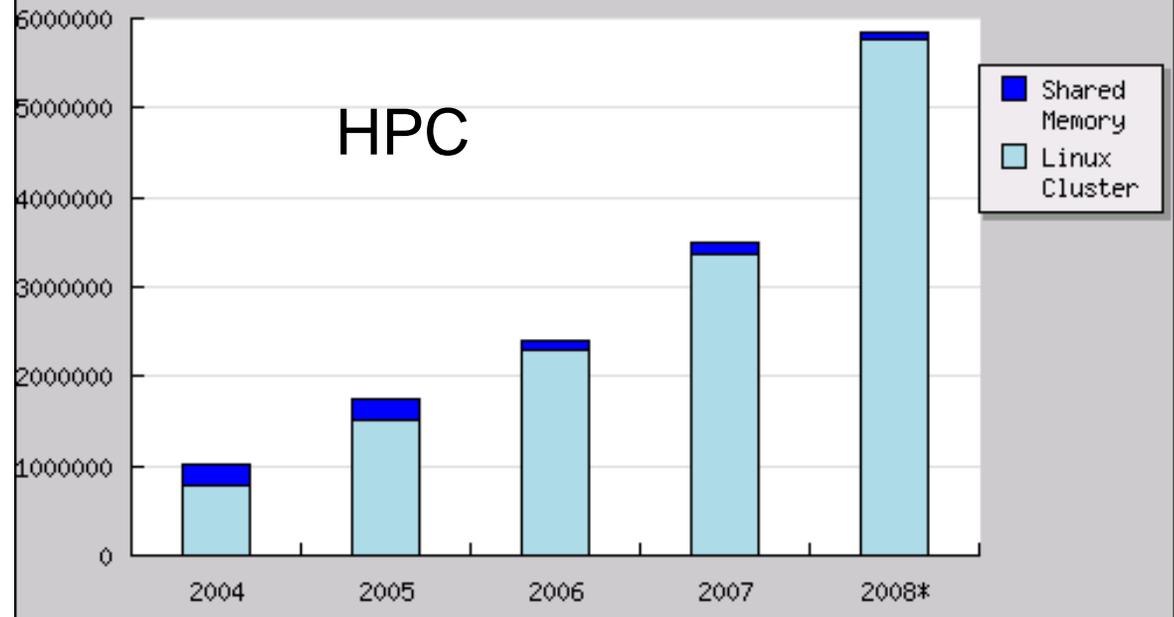


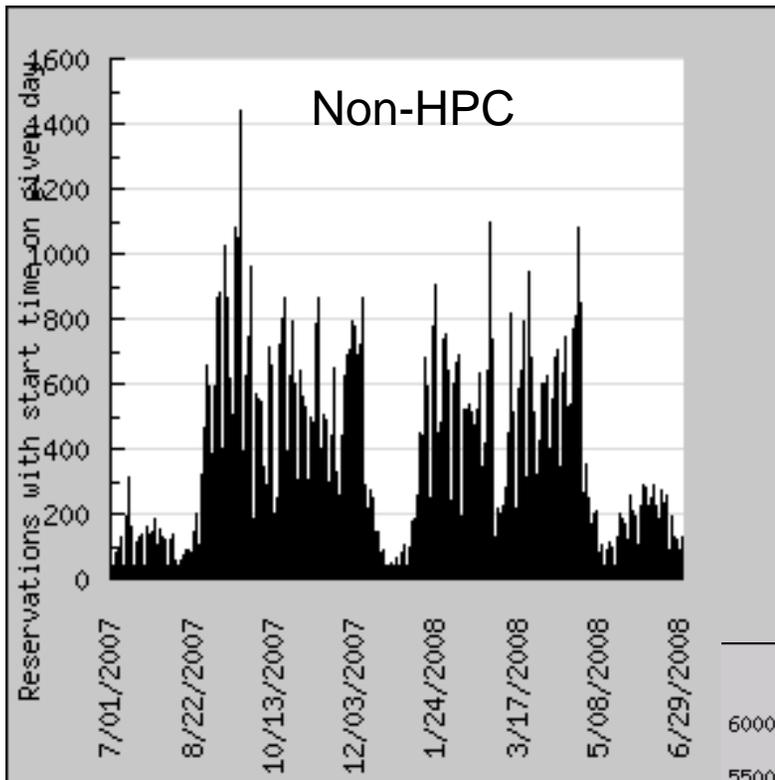
Non-HPC Reservations:

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



HPC CPU-Hrs Used



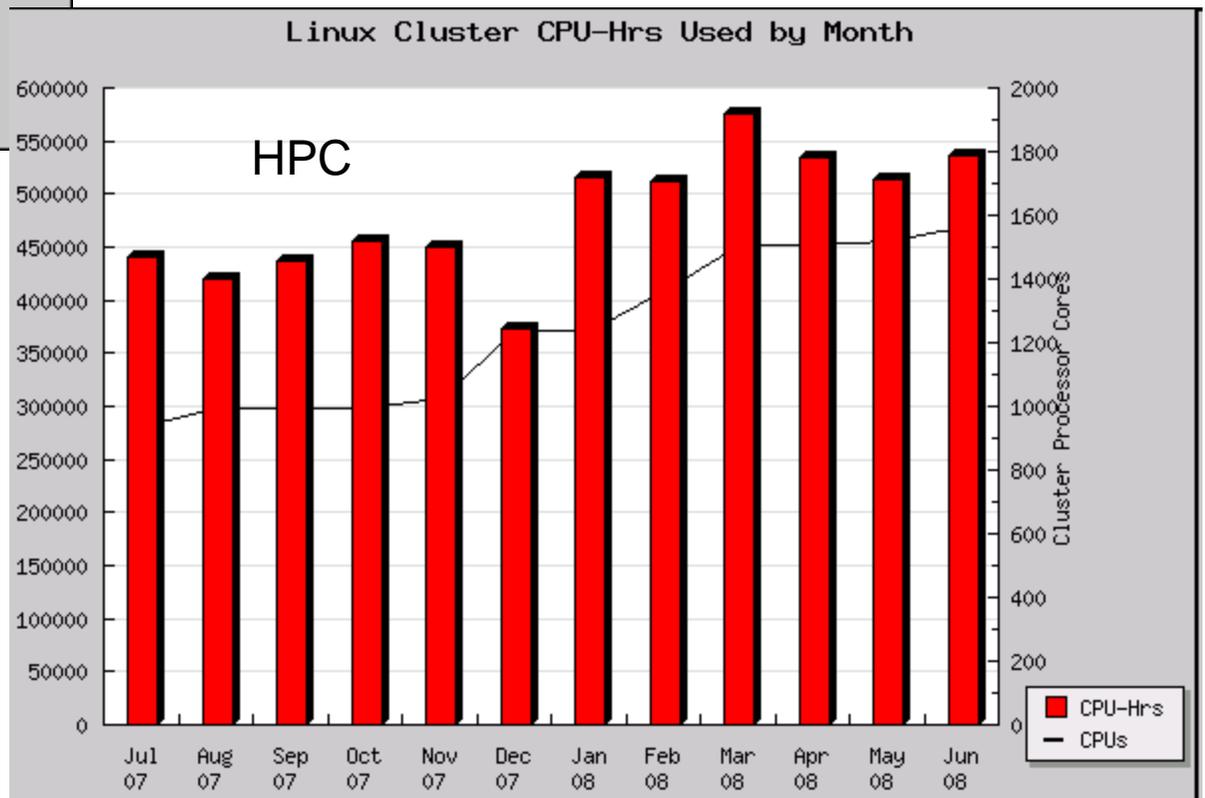


Non-HPC:

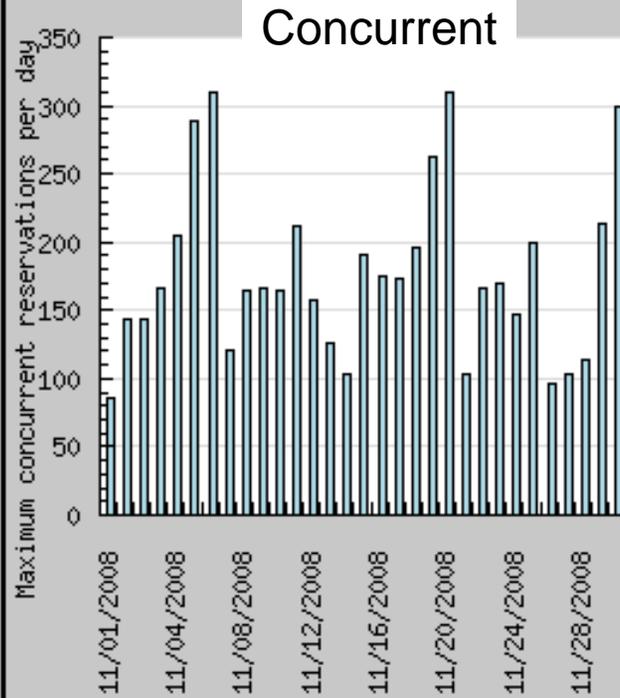
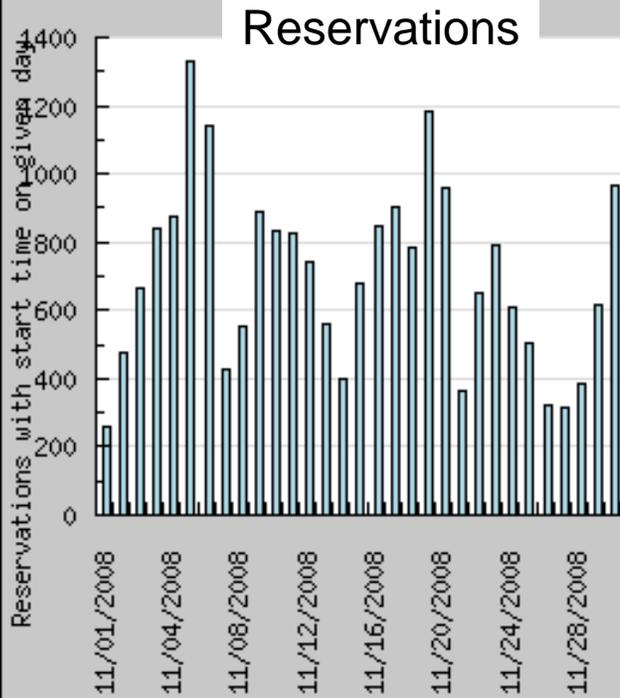
Total Reservations:	130,800
Total Hours Used:	198,583
"Now" Reservations:	125,278
"Later" Reservations:	11,436
Unavailable + Failed:	5,914
Failed:	1,611
Reliability:	0.955 – 0.988
Load times < 2 minutes:	109,223
Load times >= 2 minutes:	21,577

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

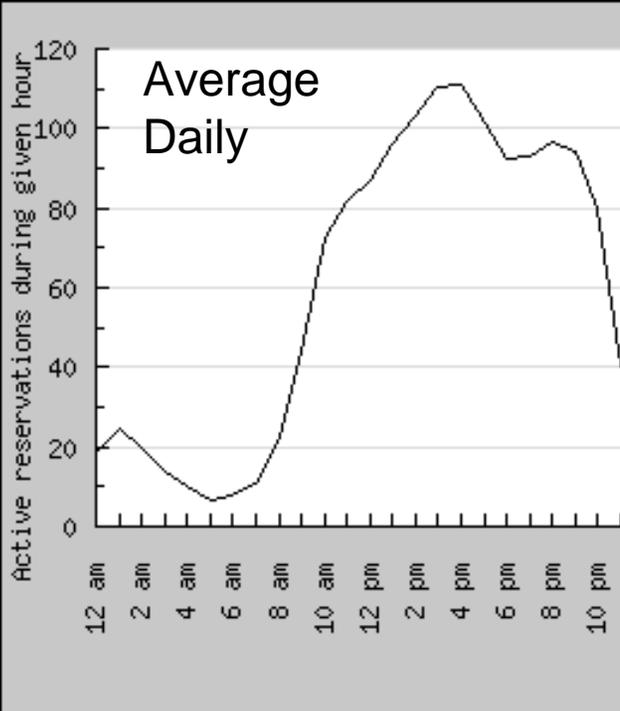


November 2008 Non-HPC



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

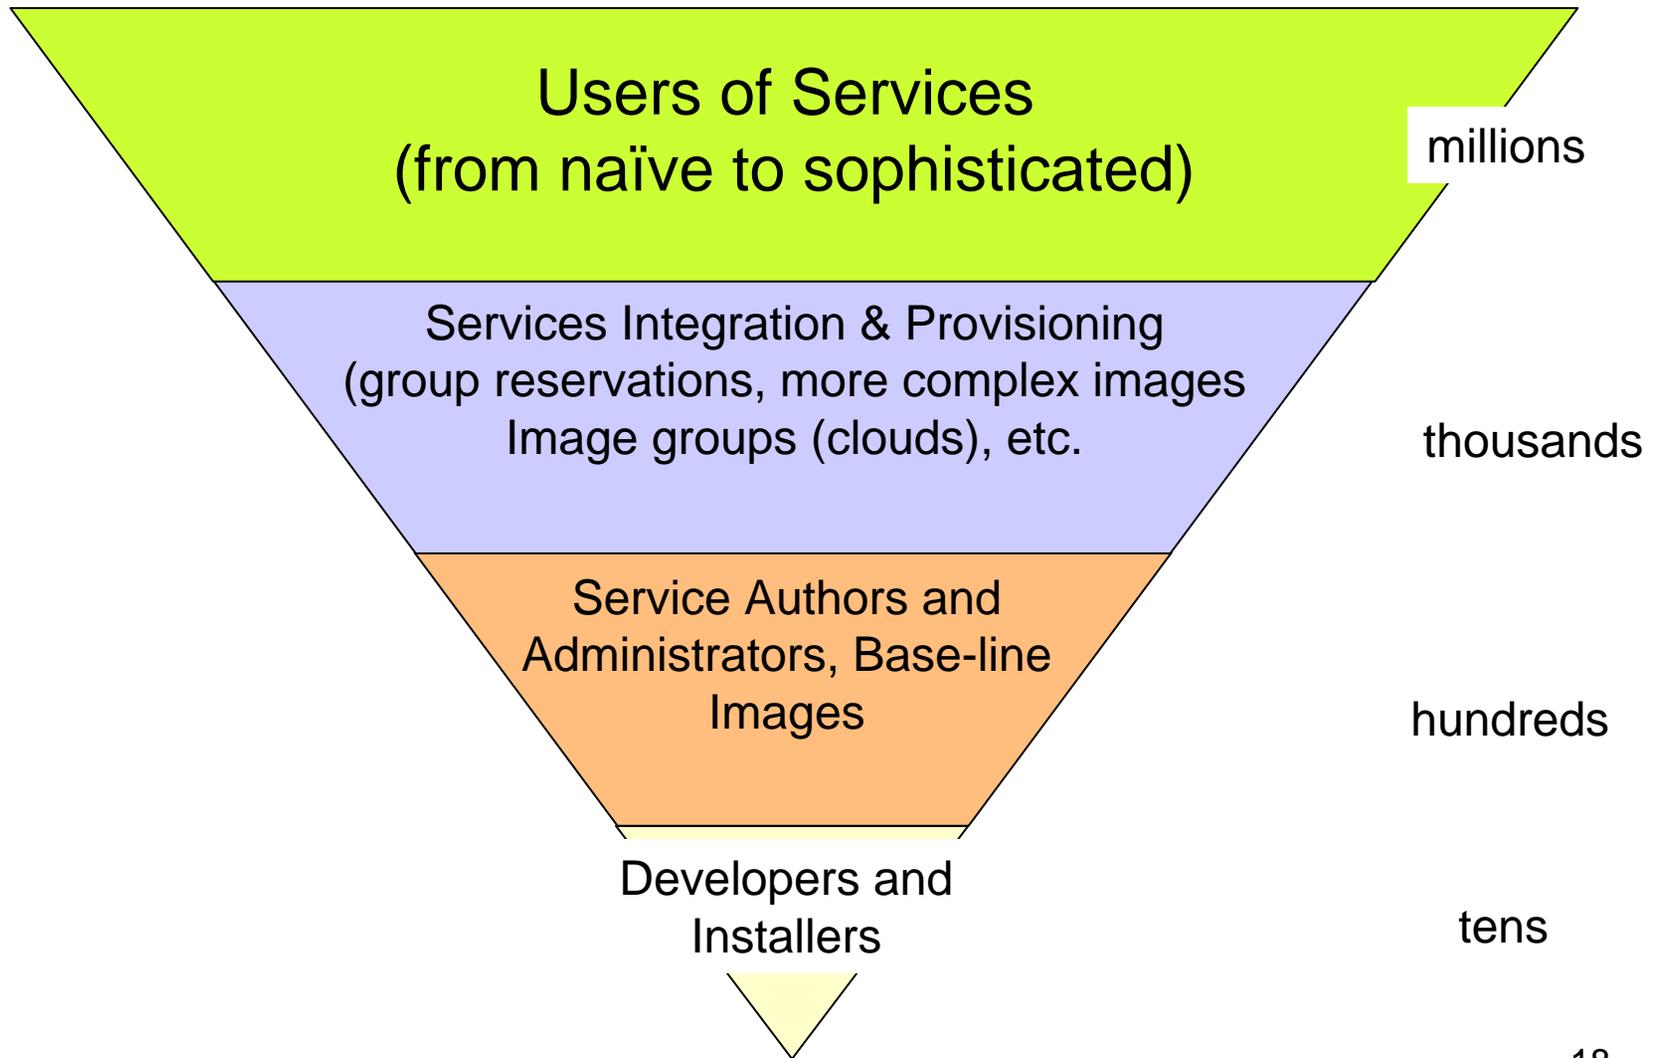


Total Reservations:	20,686
Total Hours Used:	31,853
"Now" Reservations:	19,770
"Later" Reservations:	1,933
Unavailable + Failed:	1,017
Failed:	429
Reliability:	0.950 - 0.979
Load times < 2 minutes:	17,013
Load times >= 2 minutes:	3,673
Total Unique Users:	4,095

Options



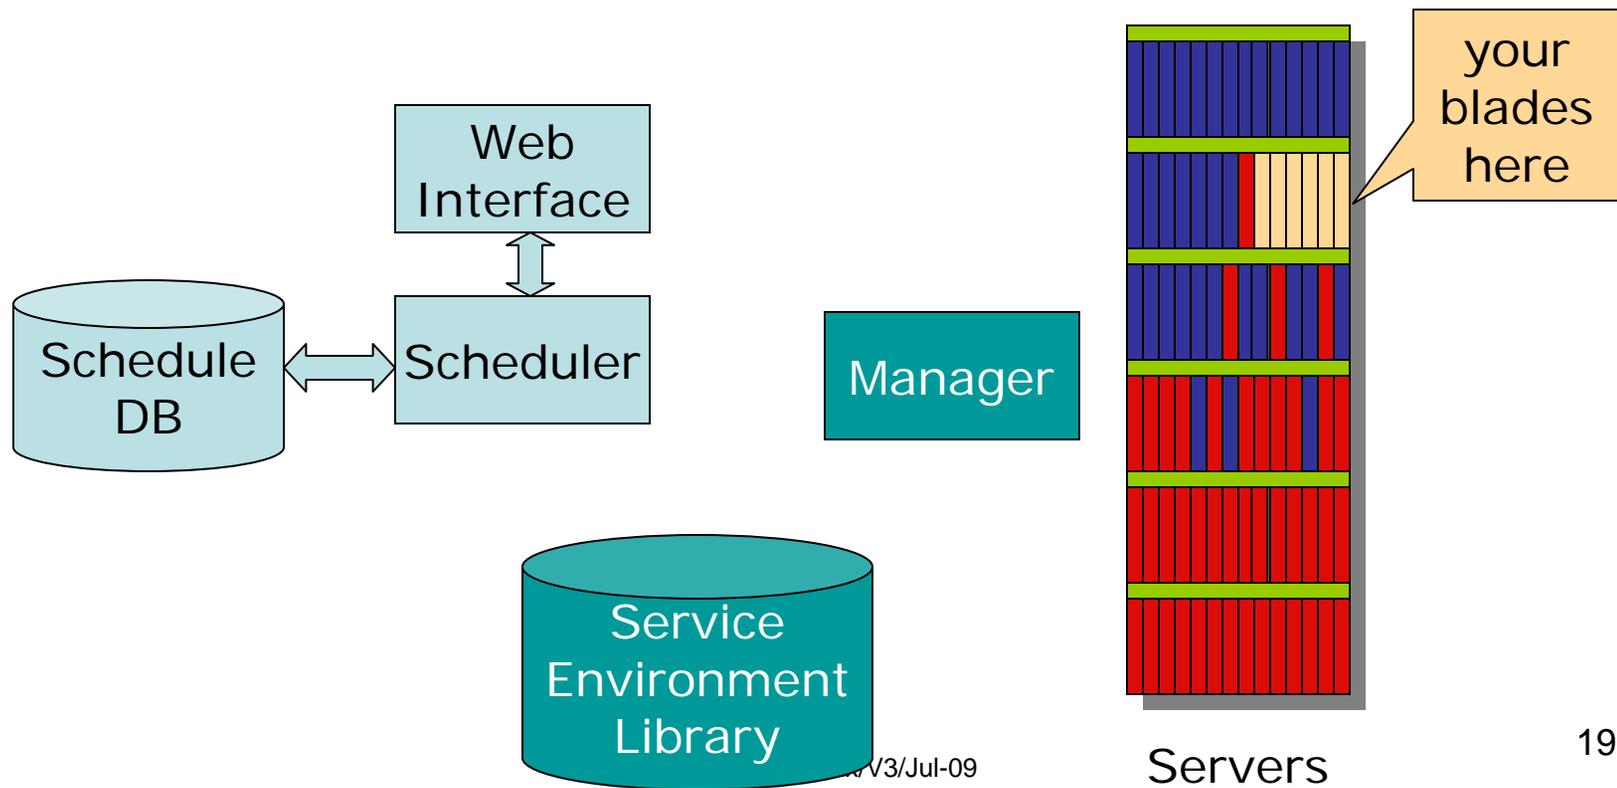
Service Composition



Participation Options

Level 1,2

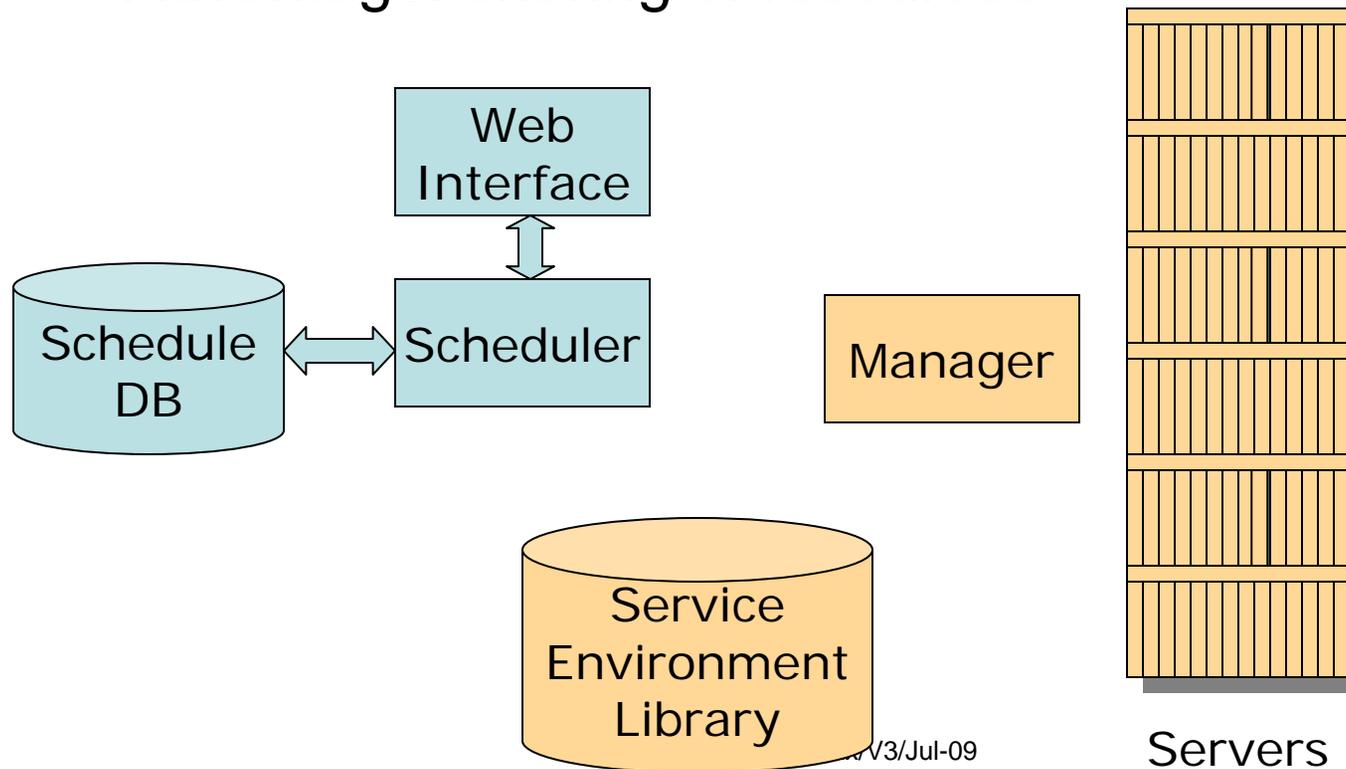
- Partner option
 - add server blades to existing VCL



Participation Options

Level 3,4

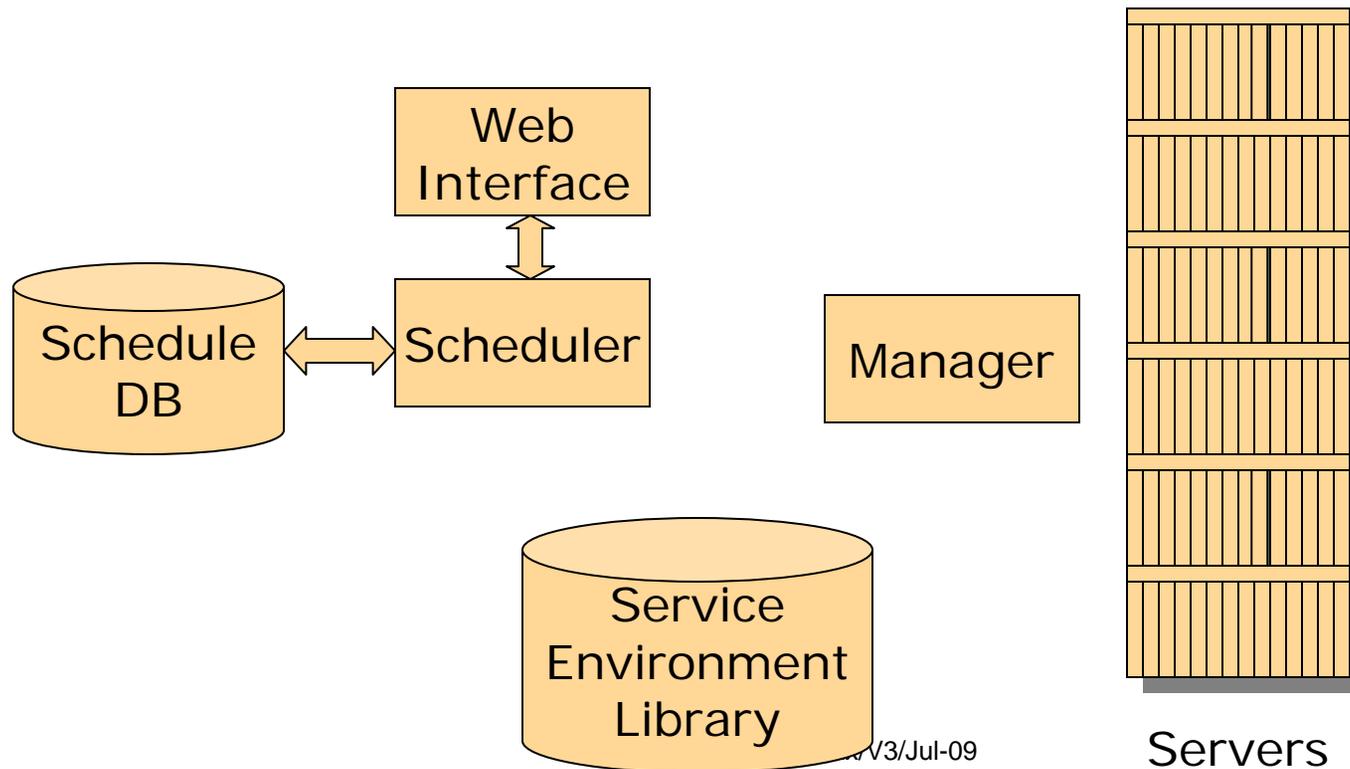
- Operate managed servers
 - use customized web interface and scheduler from existing site
 - Advantage: sharing of resources



Participation Options

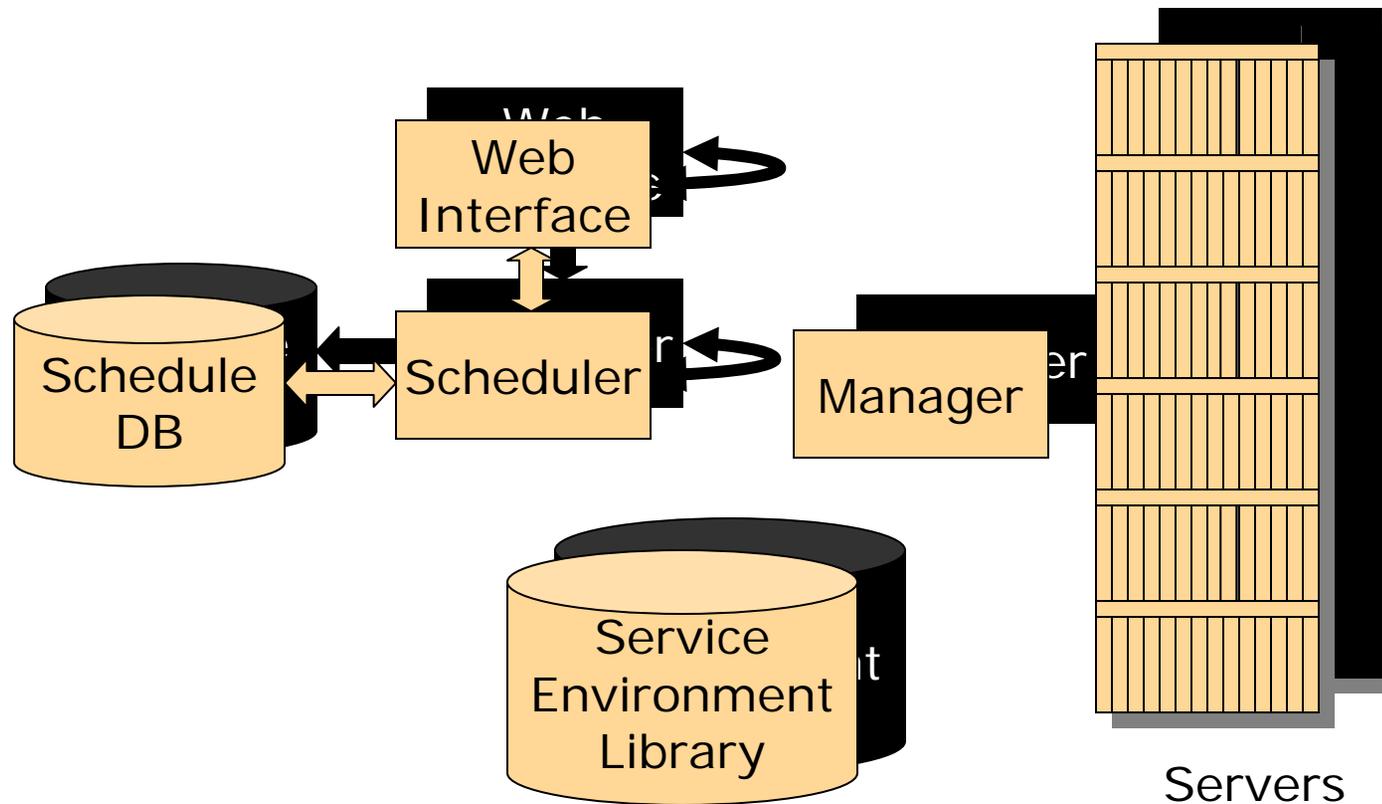
- Operate own standalone VCL
 - use VCL software on your system

Level 3,4



Participation Options

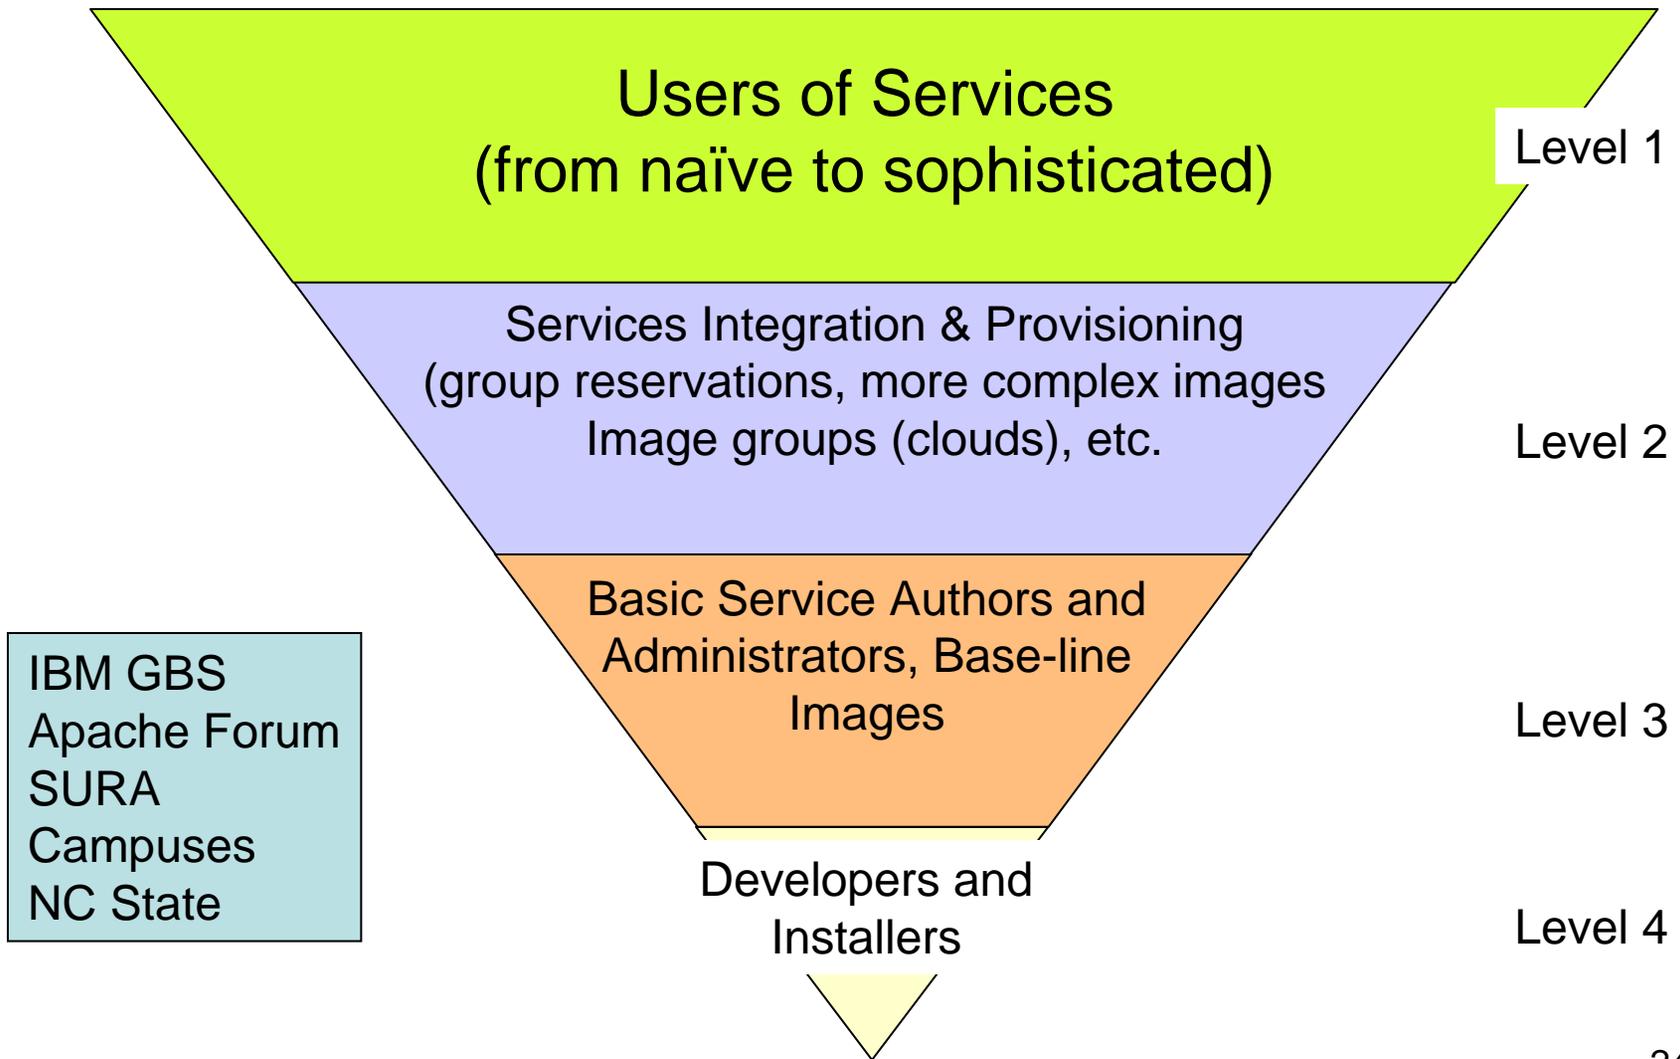
- Interoperate your VCL with other sites
 - Development effort - initial steps underway



Consortium Support Infrastructure



Support



VCL Options

- Full VCL install
 - <http://cwiki.apache.org/VCL/>
- VCL-in-the-box
 - In alpha-testing phase
- VCL VM appliance
 - 1.6 (<http://www.ibiblio.org/vclvm/>)
 - 2.1 (being retired in favor of 2.2)
https://wiki.oscr.ncsu.edu/wiki/index.php/The_VCL_2.1_Virtual_Appliance#SVN_linkages_of_this_VM)
 - 2.2 (being built)
- Cloudburst VCL
 - In preparation

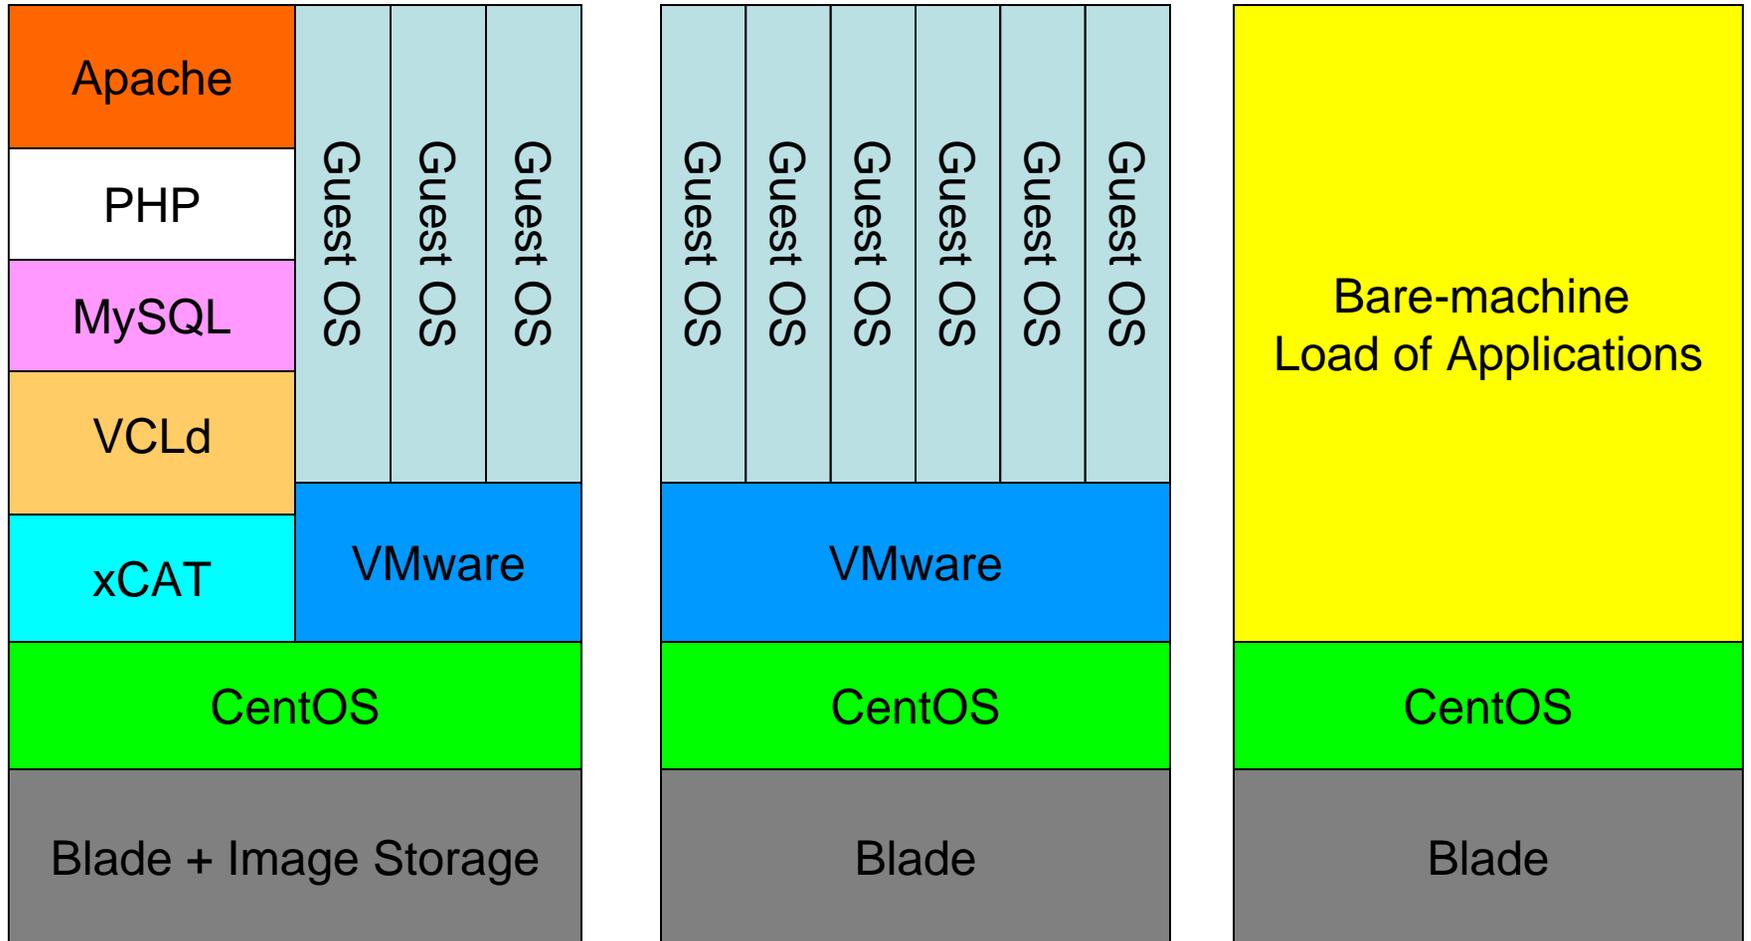


Full VCL Installation

- Blade 1 – Apache, PHP (access)
- Blade 2 – MySQL (database)
- Blade 3 – vclld, xCAT, VMware toolkit, Provisioning Manager X such as IBM Ensemble (VPS) (management node)
- Blades 4 to n - VM and bare-machine loads
- A management node for every 100 or so blades
- Management nodes can be distributed



VCL-in-the-Box



1

VCL Workshop - Fairfax/V3/Jul-09 **2-100 (any combination)** ²⁷

Cloudburst VCL

- CloudBurst will deploy VCL-Appliance
- It will allow construction of cloud-within-a-cloud solutions



Ensembles

- Collections (pools) of “like” managed resources
- Easy insertion of new hardware
- Autonomic behavior
- Fault-Tolerance
- Security
- Etc.



Shades of Things to Come



Plans

- Appliances and automation of VCL installations
- Help Desk (IBM, Apache)
- Transparent enterprise storage integration
- NAT support
- Active Directory support
- Virtualization variety (VMware, XEN, KVM, ...)
- Pro-active and predictive scheduling
- Automated image construction
- Government and military-level security options
- UNC build-out, East-Coast Consortium Build-out
- Community Colleges and K-12
- Increased performance
- Seamless resource sharing
- Dynamic resource re-allocation
- Power management
- Modularization
- Other ...





Installation Clinique



Some Categories of Help Requests

- Hardware configurations (including network, <http://cwiki.apache.org/VCL/>)
- Web server
- Database
- PHP
- Management node (vclid)
- xCAT
- VMware and other hypervisors
- Image creation (under VMware and other hypervisors, bare-machine images, Windows images, Linux images)
- VCL-in-the-Box
- Etc.



Please Send Questions to:

- vouk@ncsu.edu
- aaron_peeler@ncsu.edu



VCL Configurations



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

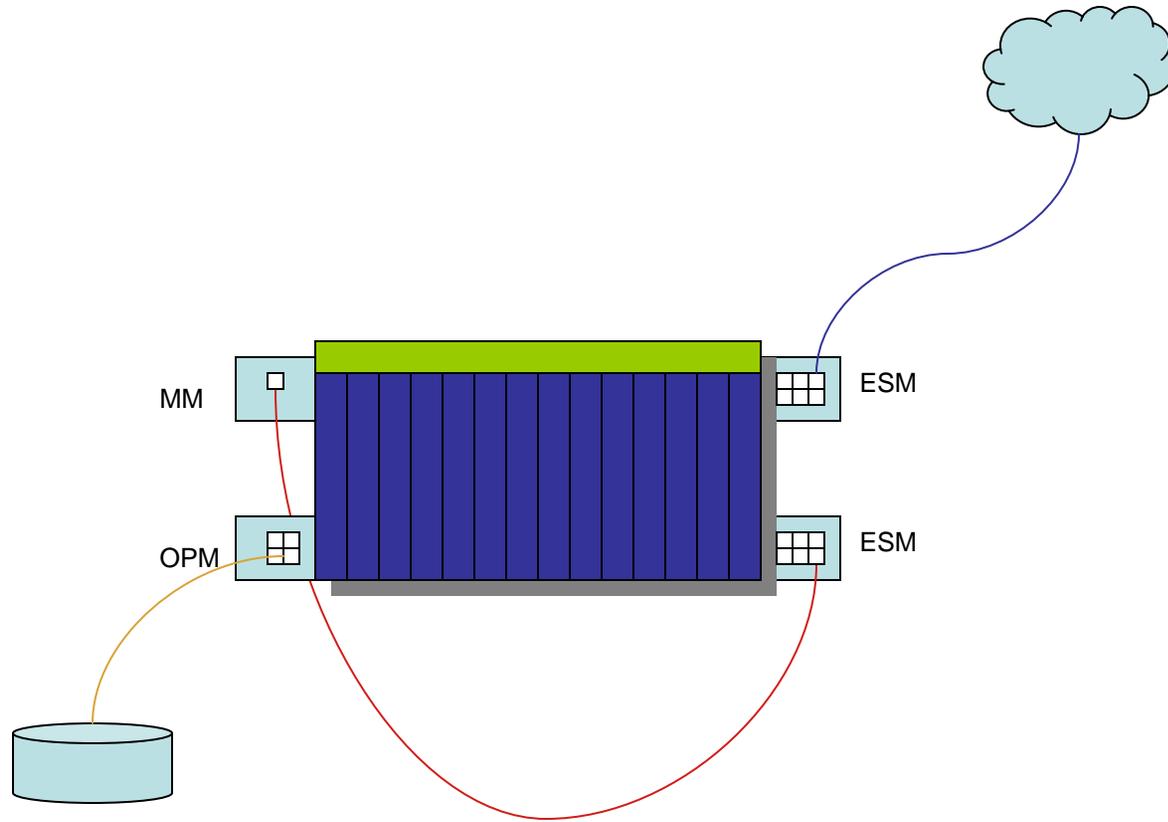


Small VCL Configuration

- 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 one 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

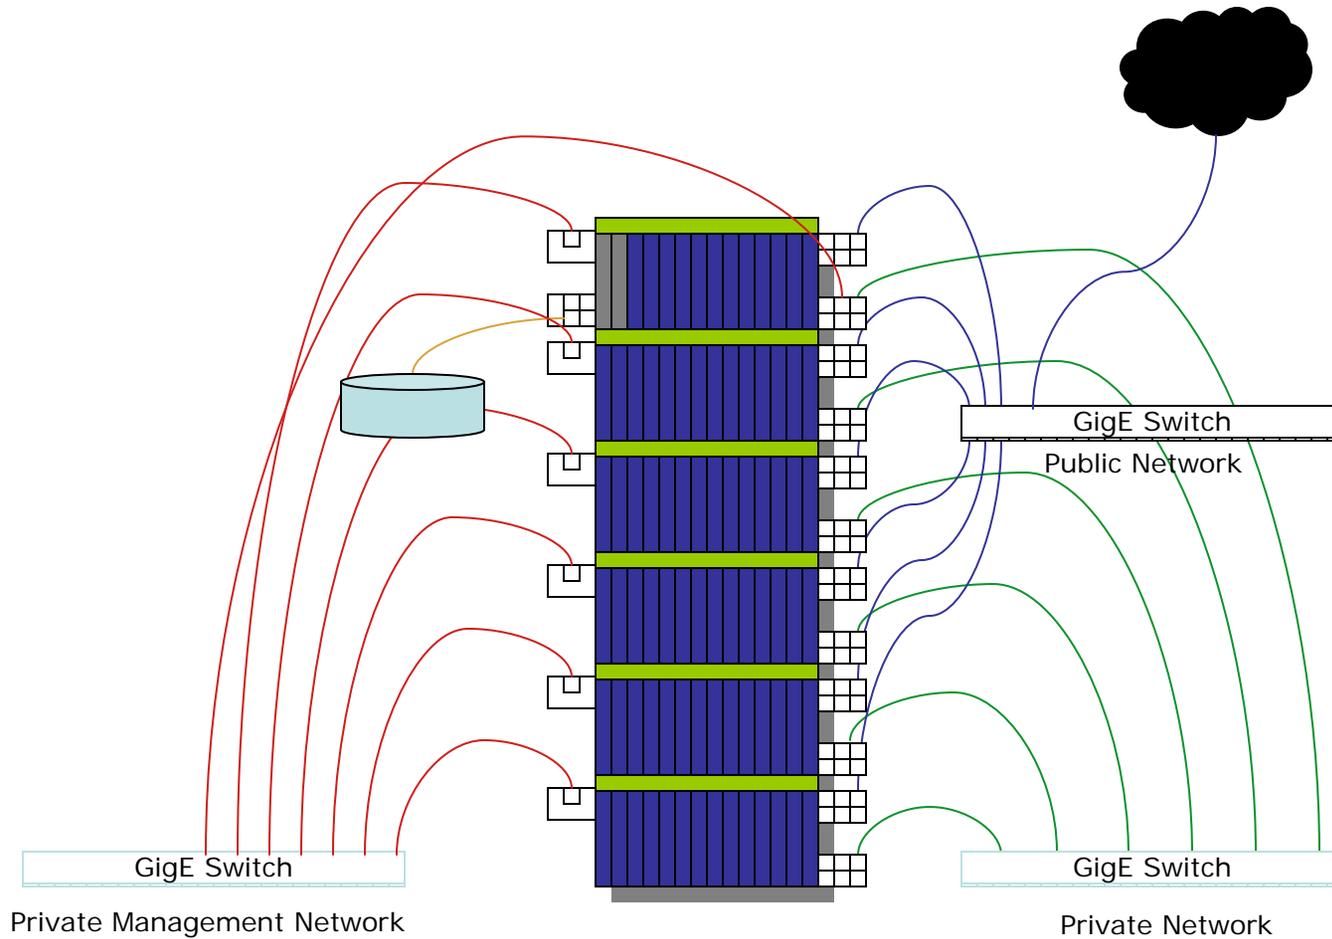


Scaling BladeCenter VCL Configuration

- 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

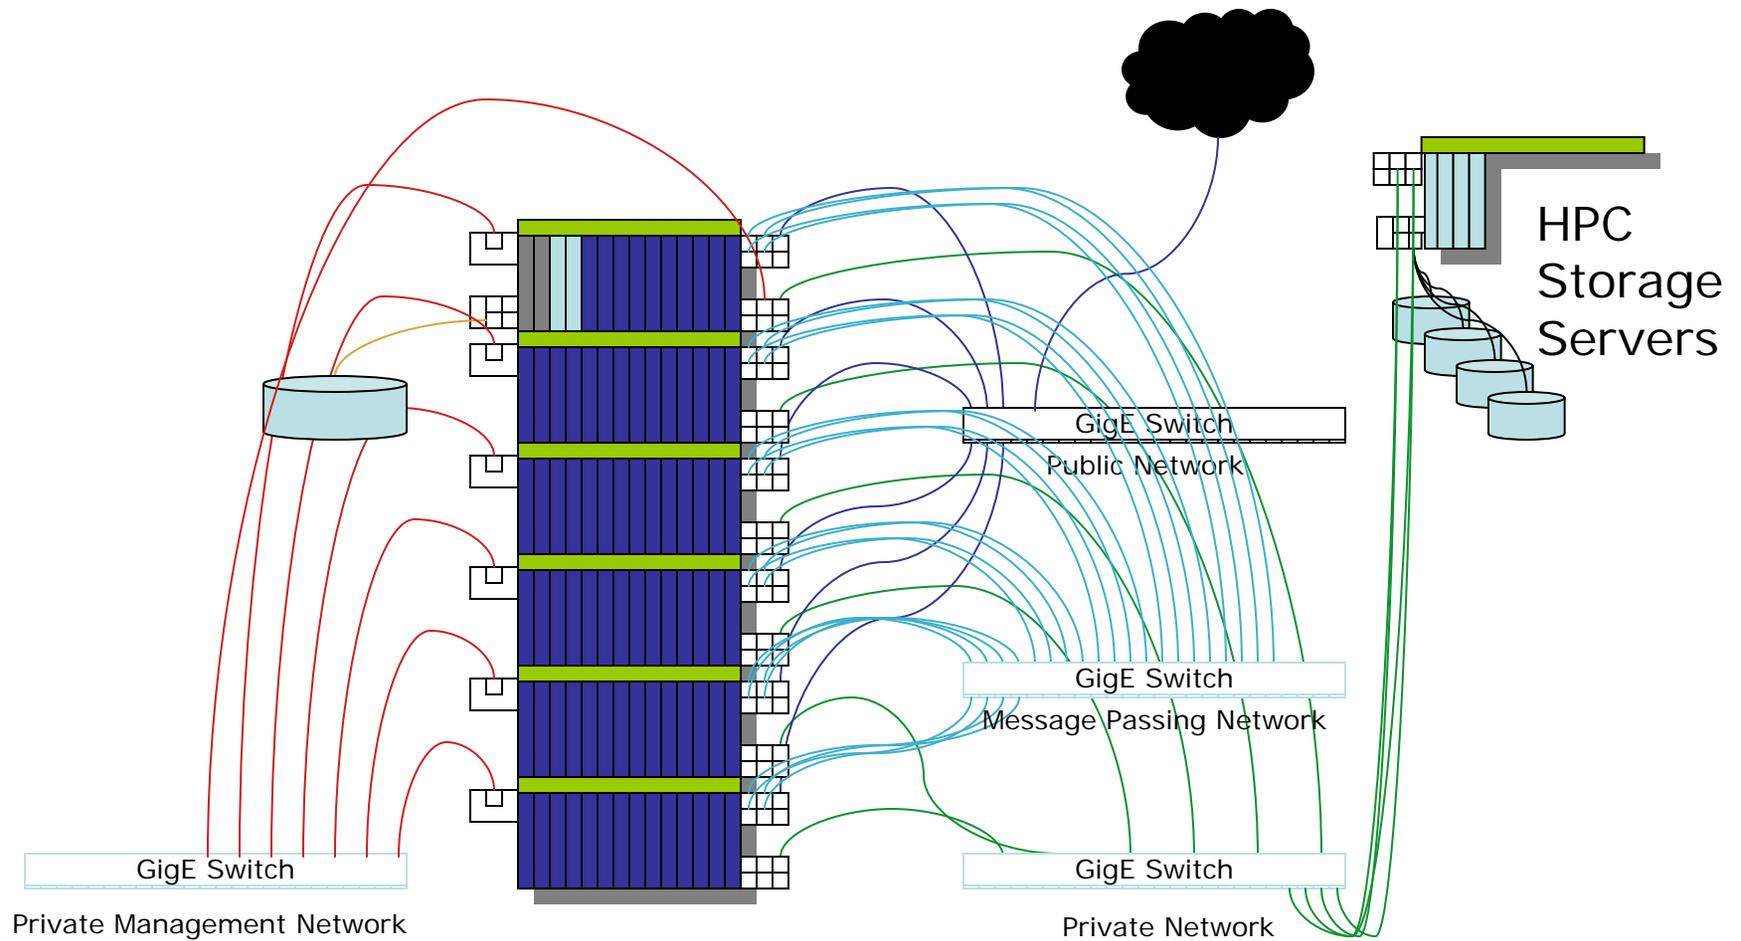


HPC Cluster in VCL

- 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

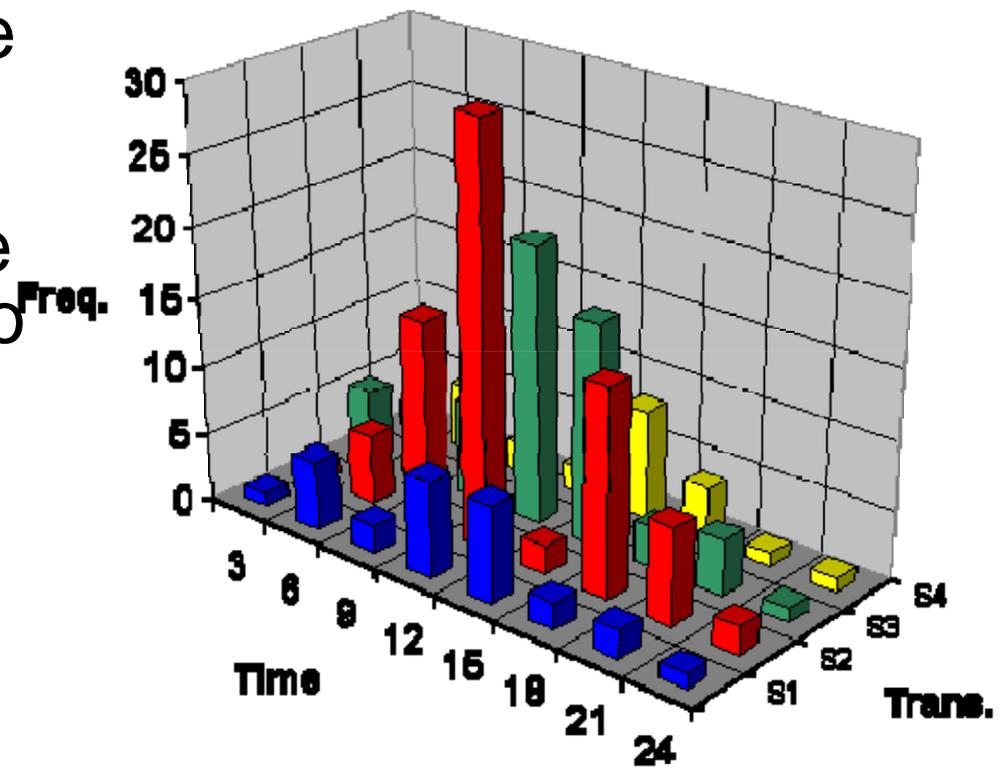


Capacity Planning



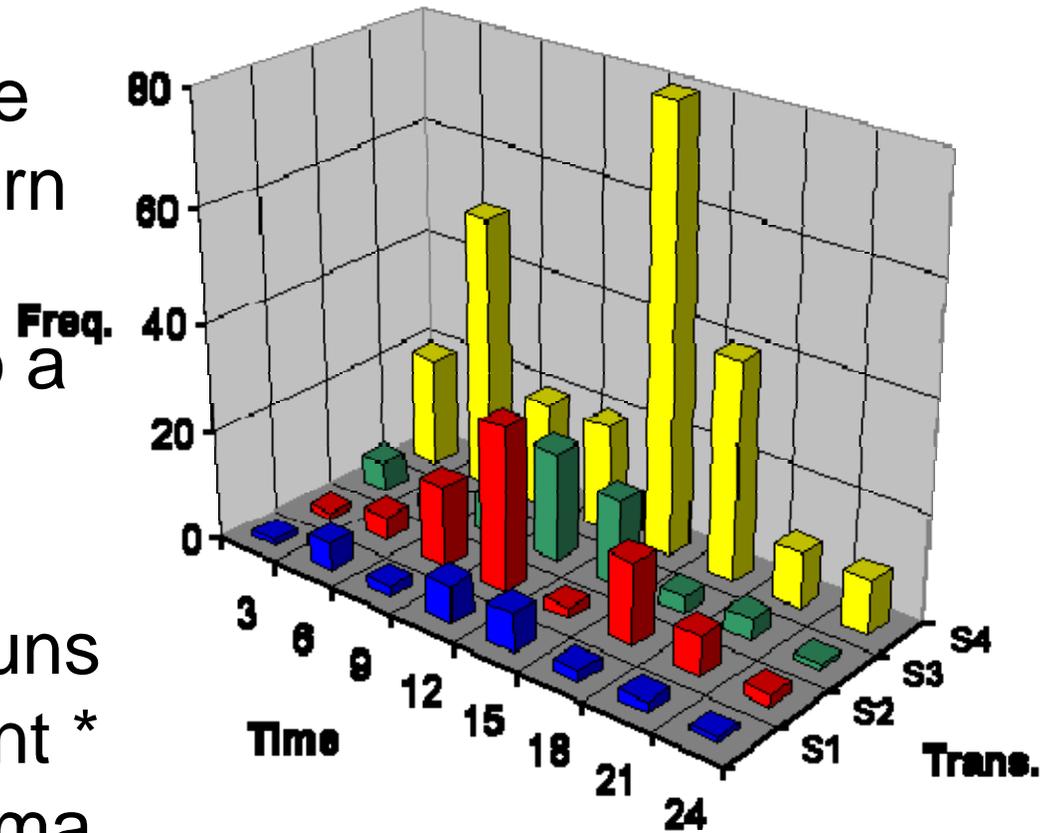
Operational Profile

- Operational profile (OP) is the set of relative frequencies of occurrence of the operations/functions used, usually expressed as fractions of the total number of such (in which case, we have probabilities).



Risk Profile

- Risk-based profiles combine the usage pattern with the cost or loss factors into a risk-profile.
- Risk = Probability_of_unsatisfactory_event * Cost_or_Loss_magnitude



VCL Profile (some elements)

- Image types and usage frequency
 - Operating Systems
 - Applications
 - Who would be using an image
 - How long
 - How many at a time
 - How far away from campus
 - Bandwidth needs (access method)



VCL O/P (2)

- End-to-end
- Bandwidth: few kbytes/sec to several MB/sec
- Image memory and CPU footprint (how many, virtualization, etc.)
- Peak vs. average
- Virtualization level will depend on the hardware and application characteristics.
- Examples of functions (demo)



VCL O/P (3)

- Classes using VCL
 - Image
 - Frequency
 - Concurrency
 - Other
- Examples of functions (demo)



Level 1 Support

- Who
- Training
- Cca 1 to 2 hours
- Examples of functions (demo)



Level 2 Support

- Who
- Functions
- Image creation
- Admin management
- Etc.
- Cca 1 t 2 day training
- Examples of functions (demo)



Level 3 and 4 Training

- More complex functions of VCL
- 3 to 4 weeks of training
- Examples of functions (demo)

