

Powering “Clouds” - the Value Proposition -

Mladen A. Vouk

Professor and Department Head of Computer Science, and
Associate Vice-Provost for Information Technology

North Carolina State University
Raleigh, NC 27695






About the Speaker

Department Head and Professor of Computer Science, and Associate Vice Provost for Information Technology at N.C. State University, Raleigh, N.C., U.S.A.
vouk@ncsu.edu

- **Mladen A. Vouk** received Ph.D. from the King's College , University of London , U.K.. He has extensive experience in both commercial software production and academic computing. He is the author/co-author of over 300 publications. His research and development interests include software engineering, scientific computing and workflows, information technology (IT) assisted education, and high-performance computing and networks. Dr. Vouk has extensive professional visibility through organization of professional meetings, membership on professional journal editorial boards, and professional consulting. Dr. Vouk is a member of the IFIP Working Group 2.5 on Numerical Software, and a recipient of the IFIP Silver Core award. He is an IEEE Fellow, and a member of several IEEE societies, ASQ , ACM , and Sigma Xi.

● cloud computing ● grid computing ● high performance com...




“Clouds”

A **seamless** component-based architecture that can deliver an integrated, orchestrated and **rich** suite of both loosely and tightly coupled on-demand information technology functions and services, and **significantly reduce overhead** and total cost of ownership and **services**.



Server consolidation, hardware abstraction via virtualization, resource management, reliability and availability, security, cost reduction, ...



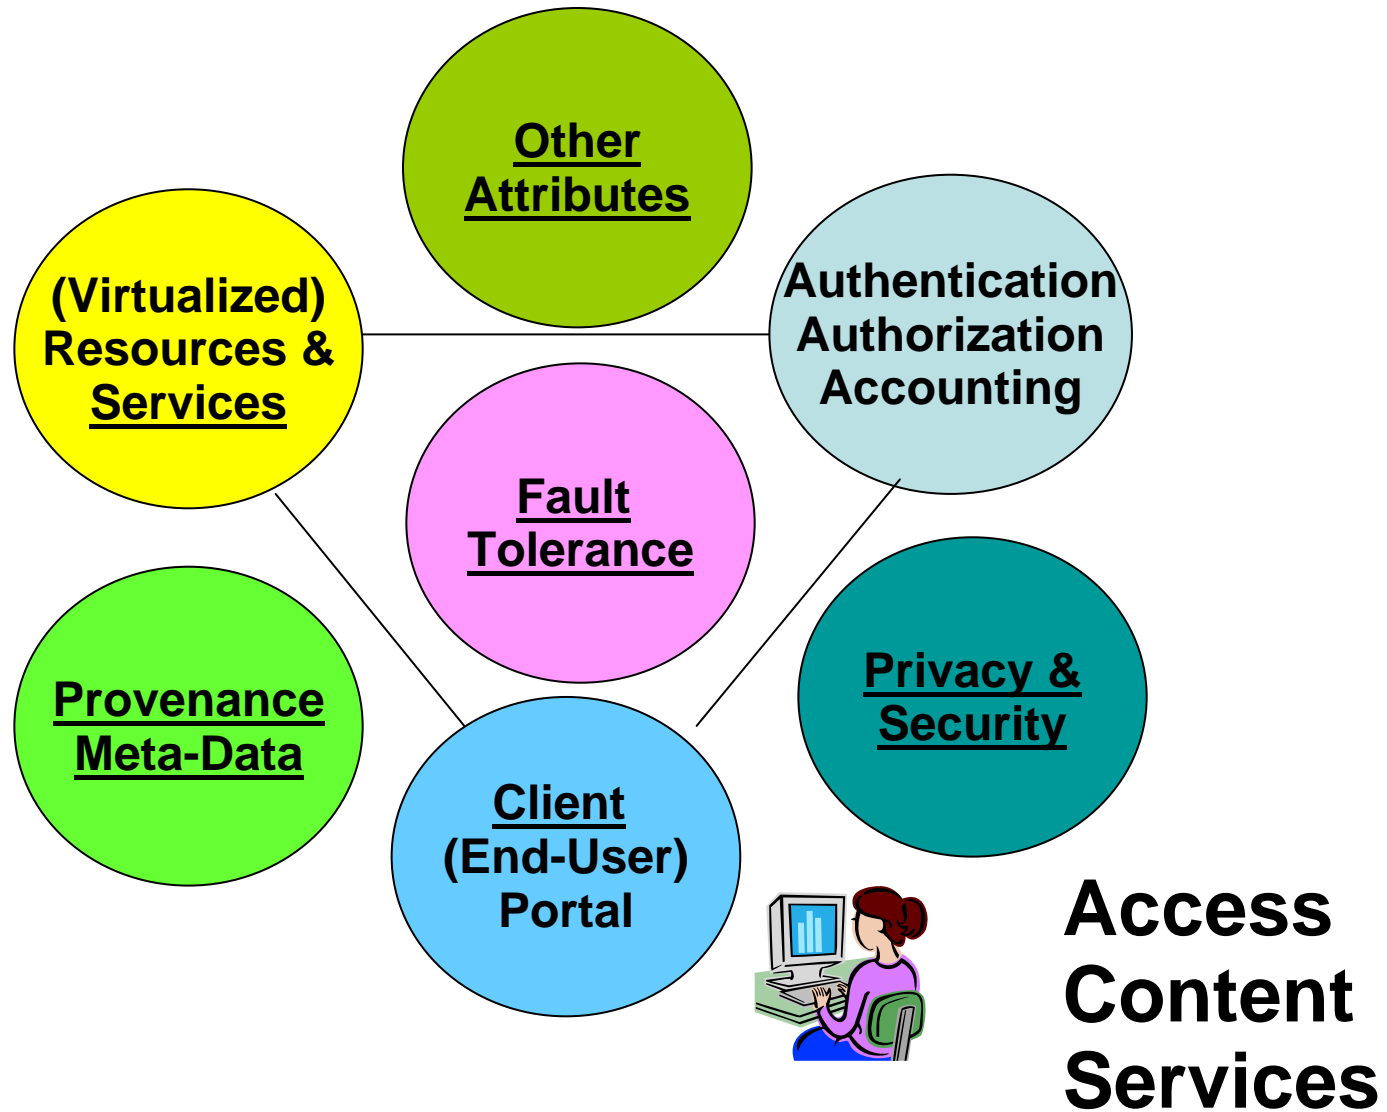
Disruptive Information Technologies



Brief History

- “Cloud” computing – builds on decades of research in virtualization, distributed computing, utility computing, grids, and more recently networking, web and software services.
 - Virtualization (since 1960s)
 - Distributed Computing (1988-1990)
 - Web (1989-1993)
 - Service Oriented Architectures (1995-2005)
 - Grids (1996-1999)
 -  Virtual Computing Laboratory – Aug 2004
 - Amazon Elastic Compute Cloud – Aug 2006
 - Hadoop/MapReduce (cca 2007)
 - IBM/Google Cloud (Oct 2007)
 - IBM Blue Cloud (Nov 2007)
 - Many other “Clouds” 

“Cloud Architecture”



<http://vcl.ncsu.edu>

An Implementation – NC State University Virtual Computing Laboratory (VCL)

VCL is Open Source – developed by NCSU OIT, COE and CSC

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

Partnerships with IBM, Intel, NetApp, Cisco, SAS, UNCGA, State of NC, NCCCS, MCNC, Friday Institute, SOSI labs, and others.

Bootstrapping reference:

<http://vcl.ncsu.edu/news/awards-and-recognition/apache-vcl-ncsu-featured-ieee-computer-magazine>



VCL Research and Development Team

- **Core NCSU team: Sam Averitt, Michael Bugaev, Patrick Dreher (RENCI), Andy Kurth, Marc Hoit, Aaron Peeler, Henry Shaffer, Eric Sills, Sarah Stein, Josh Thompson, Mladen Vouk, Brian Bouterse, John Bass, Shawn VanHulst, ...**
- **Many others at NCSU (faculty, students, staff)**
- **Many others at other sites and other organizations**





VCloud Community

NC Community College System
 NC K-12

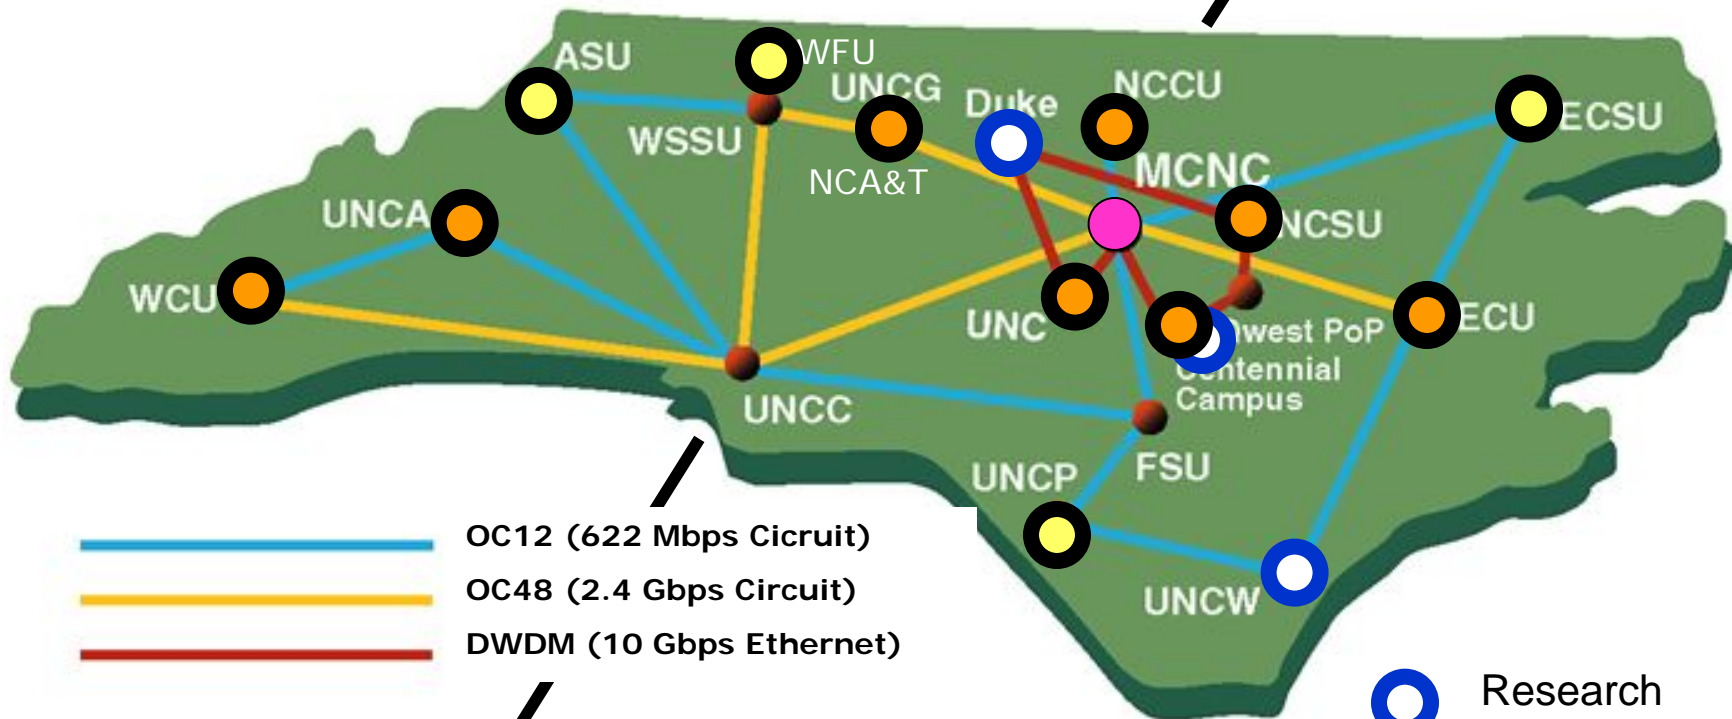
India
 Amrita U.
 U. Hyderabad
 HBTI-UPTU

Toronto
 Queens
 Waterloo
 Carleton

UMBC

BC

VTech
 ODU
 CCV
 MSU
 GM

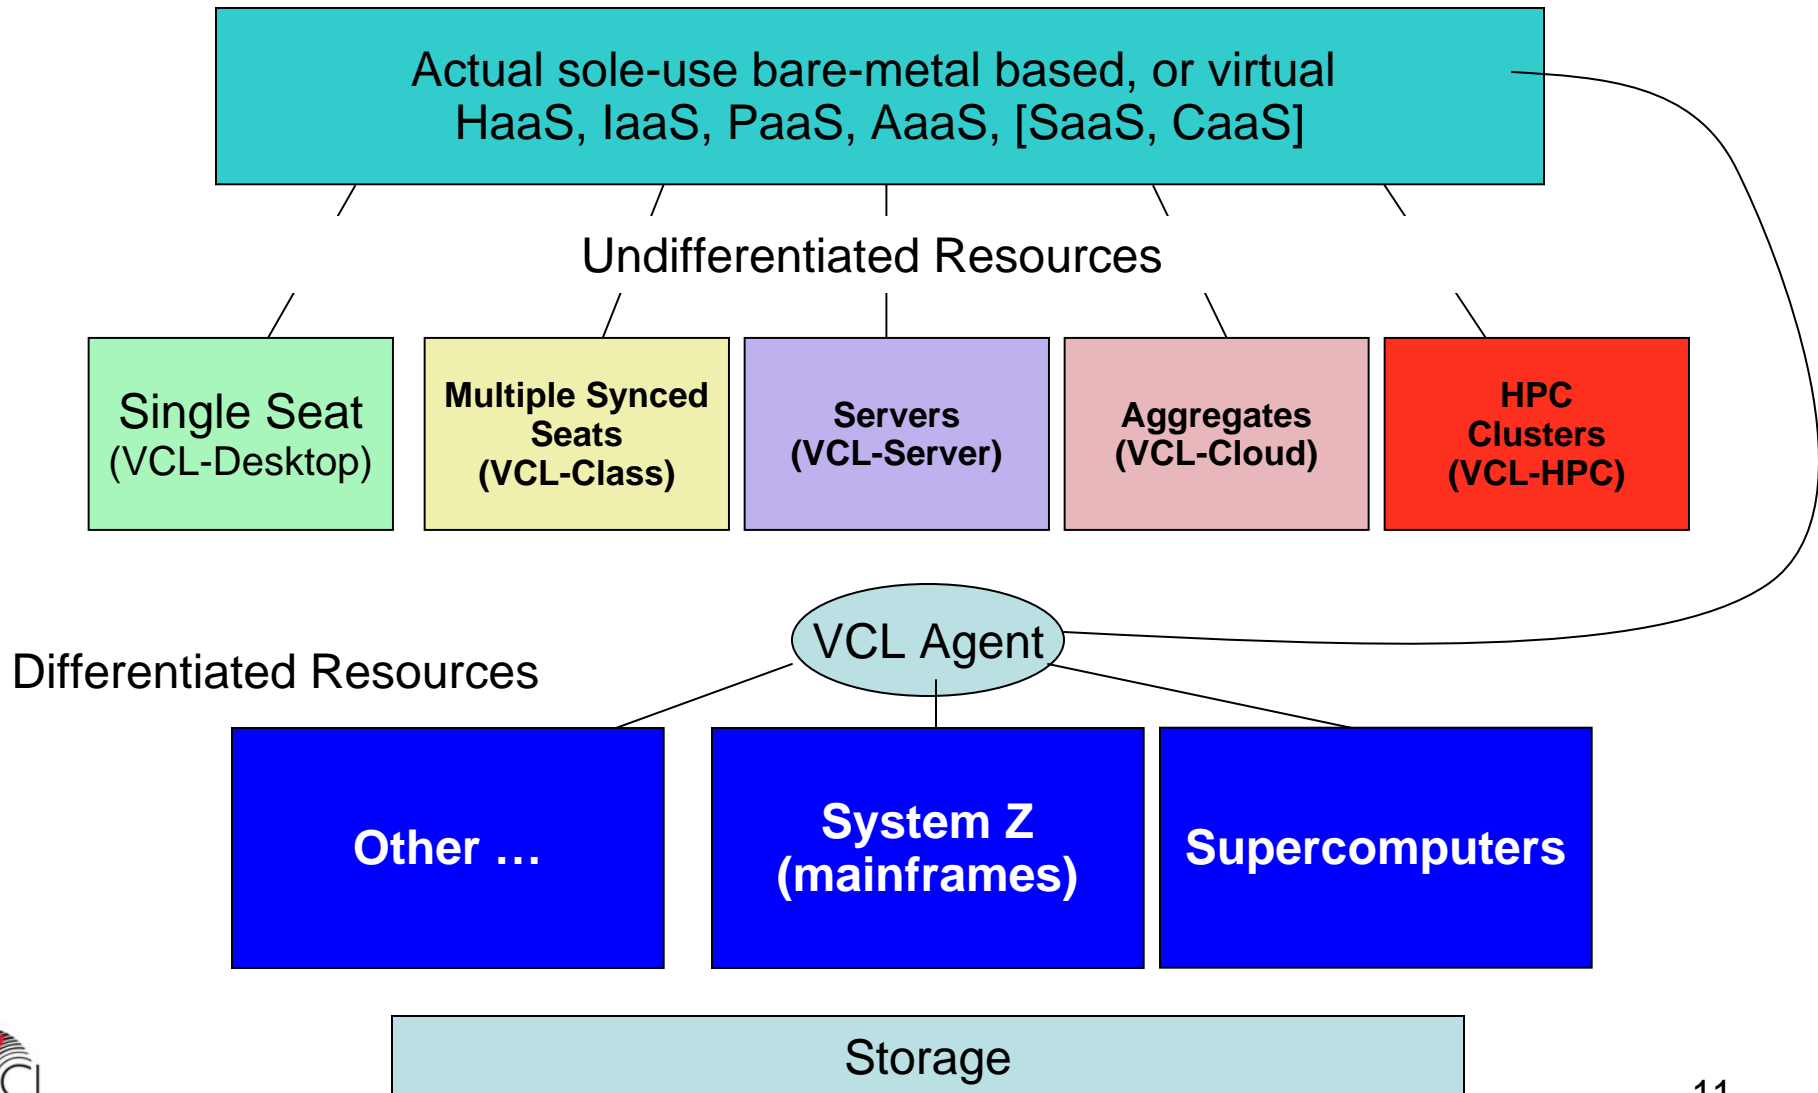


U. South Carolina,
 Clemson

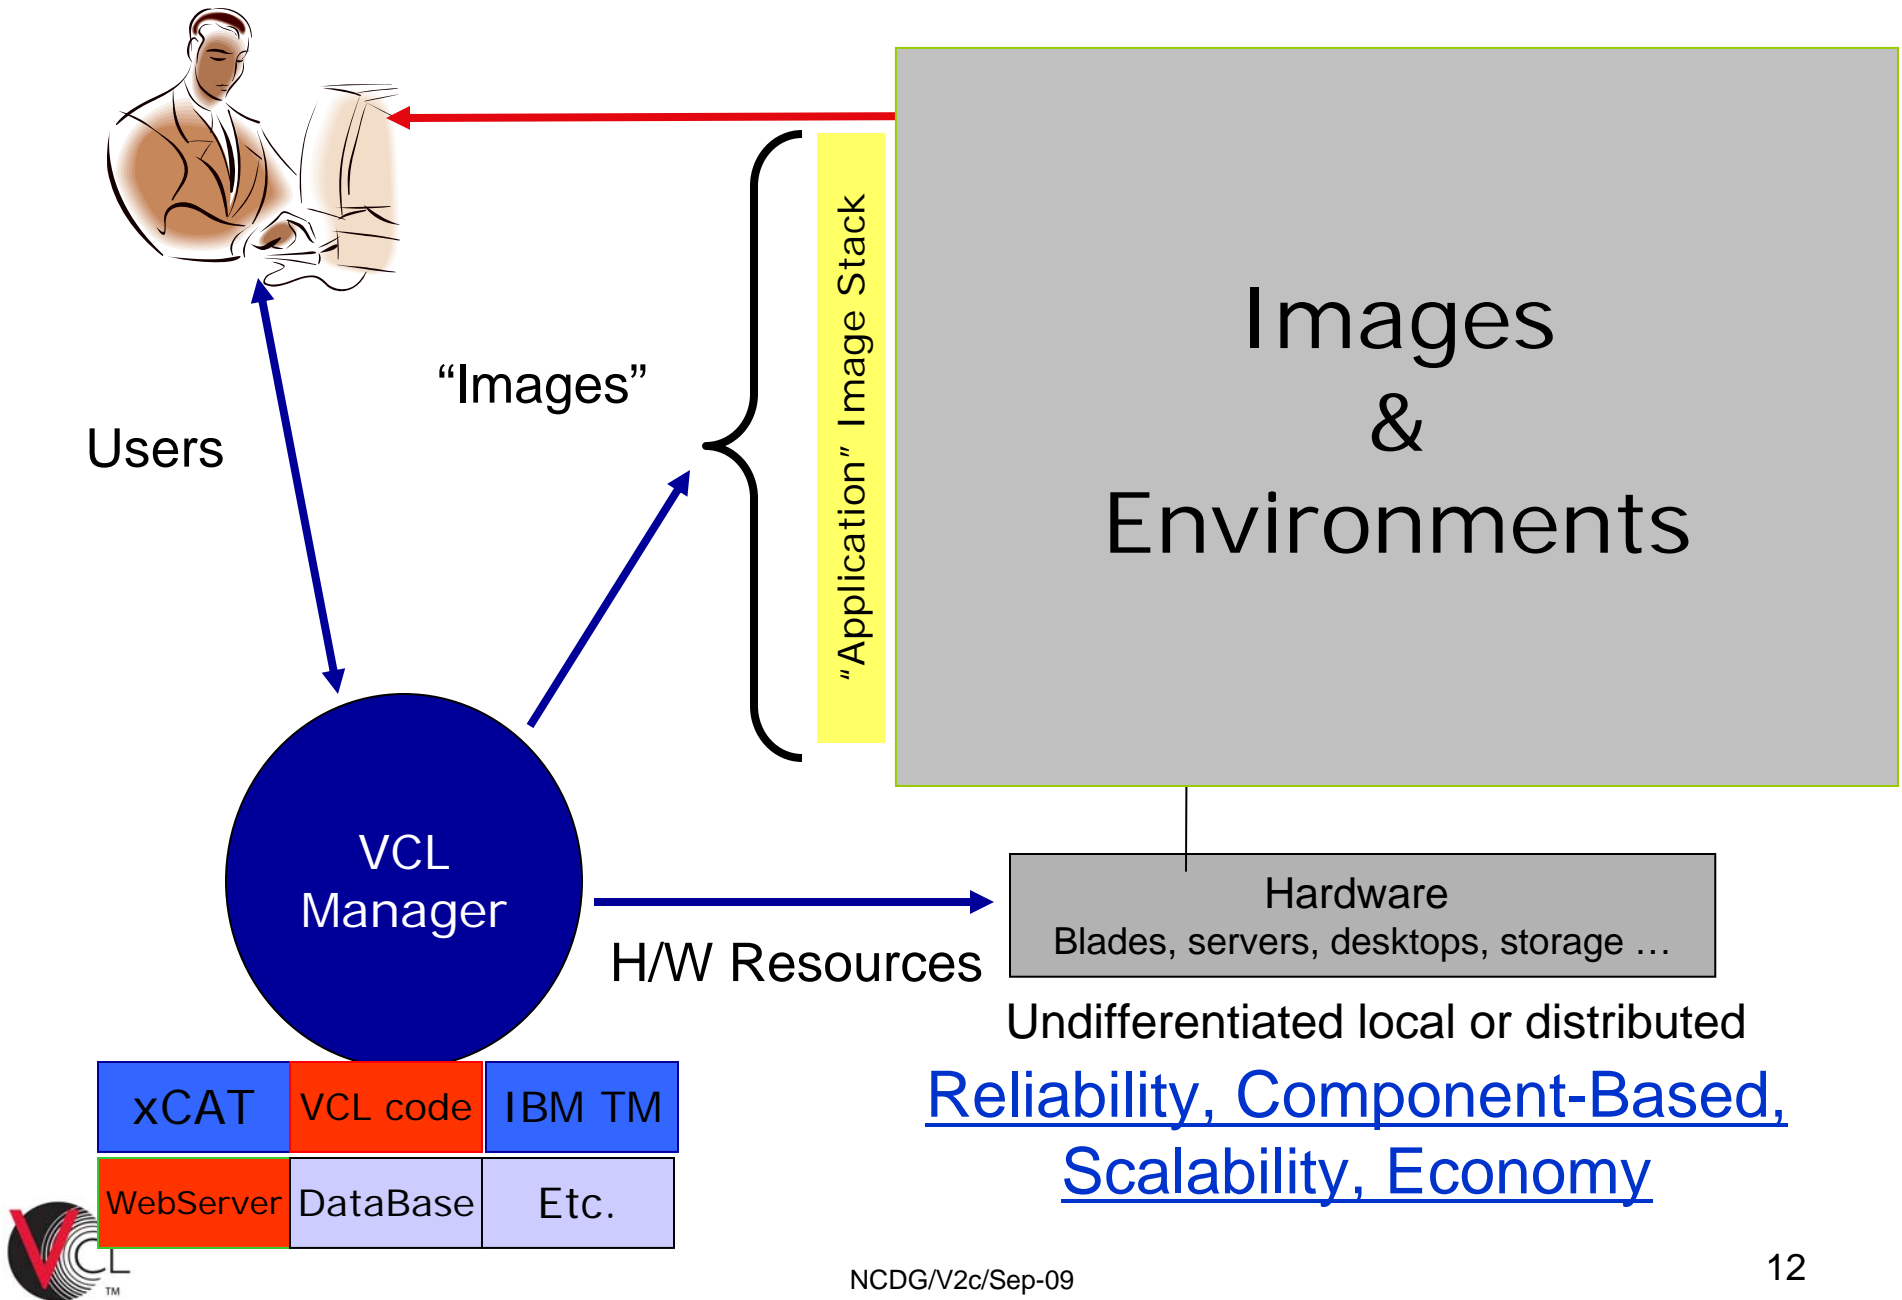
Research
 Production/
 Pilots/Users
 Interest/Plans



NC State Cloud Services



Differentiator: User to Image to Resource Mapping, Management & Provenance



Business Model

Current VCL (at NC State University):

1. cca 2,000 blades
2. open to 30,000+ students and faculty
3. cca 500 to 700 in non-HPC mode, the rest in HPC mode
4. Environment base-lines are typically Windows and Linux with a variety of applications. Depending on how demanding an application is, service may be virtualized (VMWare) or bare-metal.
5. Currently Cca 600 images, cca 120 in use per semester.
6. About 80-100,000 image reservations per semester.
7. Most of the “individual seat” requests are on-demand “Now” reservations: cca 90% of requests
8. System availability: exceeds 99%




Cost Factors


- Utilization (real-time + batch mix), operational profile
- Lab spaces (25:1) – in 2008/09 cca 160,000 non-HPC reservations (real-time), cca 7 million HPC CPU hrs (batch)
- 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 (1-2 FTEs for about 2,000 blades)
- One stop shopping (augmentation)
- Distributed burden of image creation (600+ images)
- “Green”
- Other ...




VCL Use



Capacity Planning




Green & Cost-Effective



Average daily active reservations

High-Performance Computing



Economics

- In 2008, about 7,200,000 CPU hours (about 6.9 million on HPC and about 300,000 on non-HPC) on about 1,500 blades (cca 3,000 processors) – upto about 1000 in HPC mode.
- About 70-80% utilization on the average, but lower on non-HPC side (over provisioned to handle peak loads), high on the HPC side.
- About \$2 million annually (refresh, management and maintenance, improvements, personnel, ...).
- About 27 cents or less per CPU hour (cca 3 cents HPC, 24 or less cents non-HPC).
- This can come down to 10 to 15 cents per CPU hour – and lower - with scale-up, large-scale virtualization, and new hardware (moving to quad-core processors).



Case-Study: Wake Tech Community College


- 60,000+ students
- Pilot project with cca 800 students
 - Several introductory class laboratories.
 - Using VCL with about 60 blades, no bare-metal loads (virtualization using VMware)
- Lab cost savings: cca 50%
- NCCCS ramping up VCL to 14+ Community Colleges



Contact

- M. Vouk
- vouk@ncsu.edu
- 919-515-7886



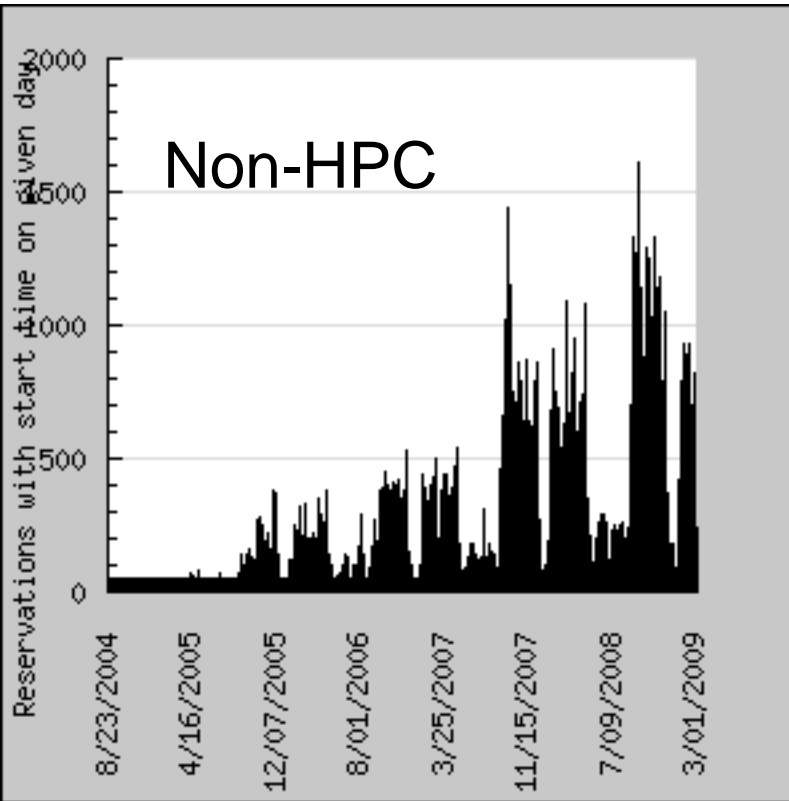


VCL Usage 2004-2008

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

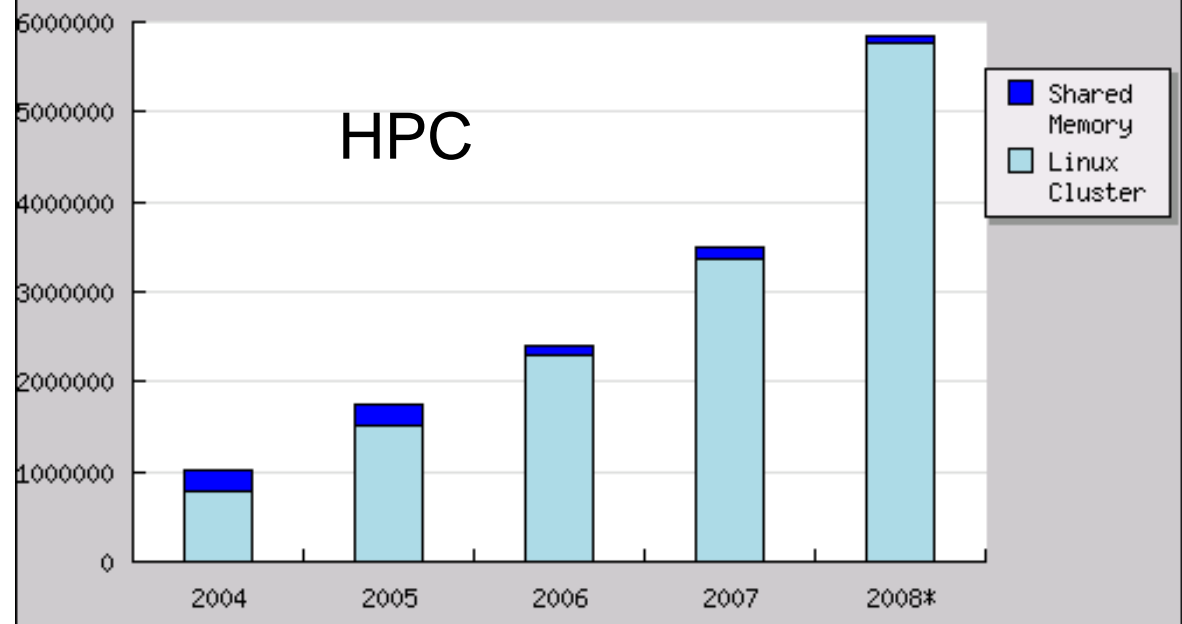



Non-HPC Reservations:

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:


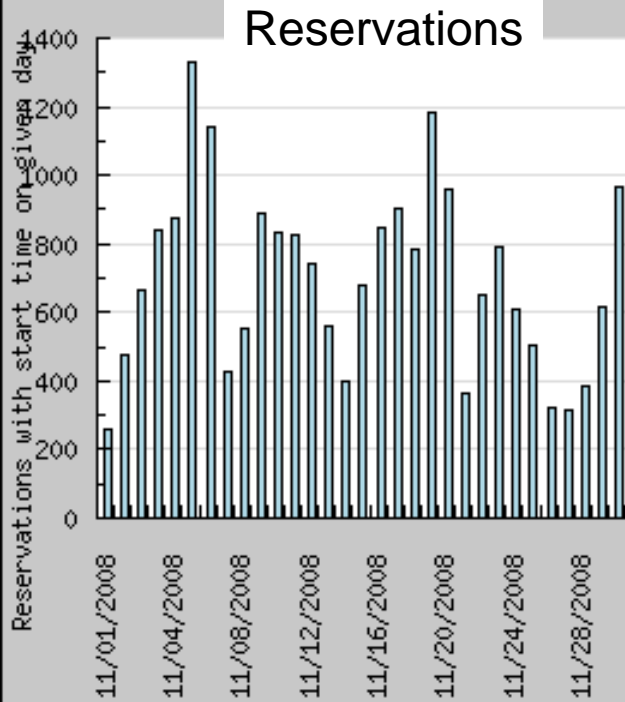
Total Reservations:	130,800
Total Hours Used:	198,583
"Now" Reservations:	125,278
"Later" Reservations:	11,436
Unavailable + Failed:	5,914
Failed:	1,611
Reliability:	0.955 – 0.988
Load times < 2 minutes:	109,223
Load times >= 2 minutes:	21,577

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

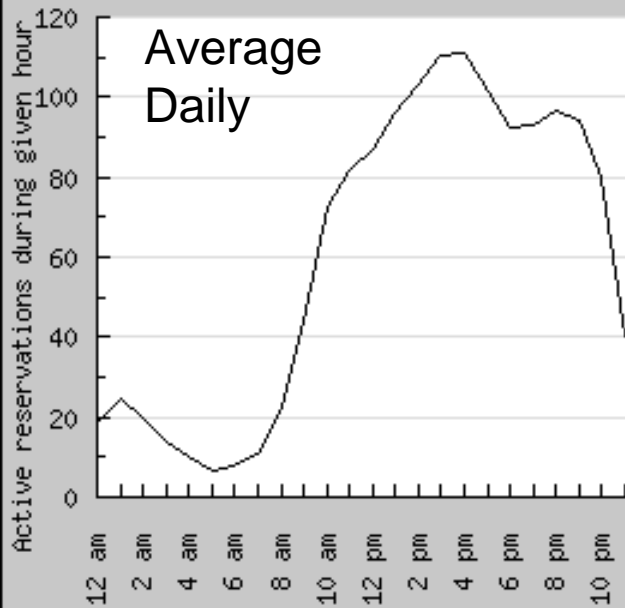


November 2008 Non-HPC



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

Awards

- ***"NC State Cloud Computing Services" received 2009 "Laureate Medal" from the Computerworld Honors Program, Computerworld Information Technology Awards Foundation.***
- ***"Virtual Computing Laboratory (VCL)" received 2007 "Laureate Medal" from the Computerworld Honors Program, Computerworld Information Technology Awards Foundation.***
- ***Finalist in the 2007 Best Practices in Infrastructure Management – Computerworld – Infrastructure Management World***



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

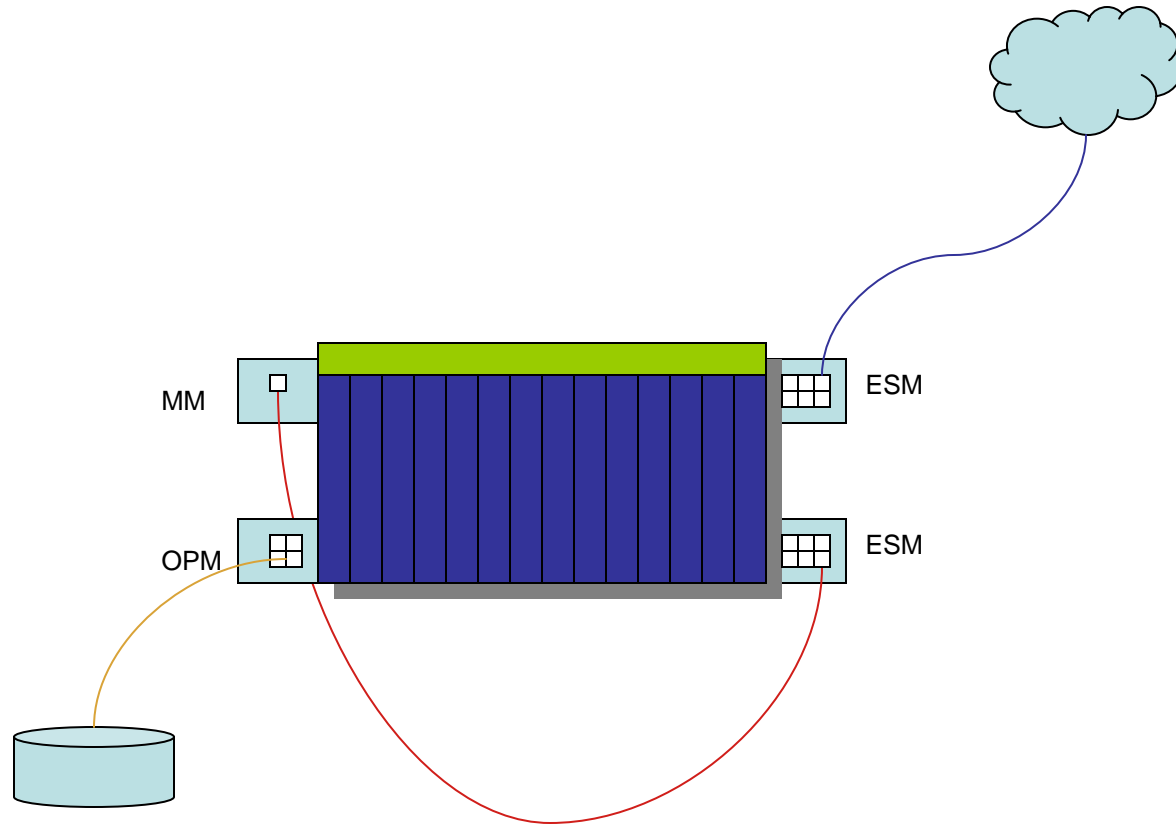


Small VCL Configuration

- 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

- IBM BladeCenters or iDataPlex - ~84 physical servers/rack, Dell and HP equipment can also be used
- LAMP & Management node servers
- Network switch(es)
 - Possibly 1 less network - no separate management network port (combined with one of two GbE ports and/or 10Gbps ports)
 - Server switches in iDataPlex rack, if iDataPlex is used.
 - High-security version requires VPN and VLANs to individual VMs.
- Storage



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
- UNC build-out
- Community Colleges and K-12
- Increased performance
- Seamless resource sharing
- Modularization
- Other ...



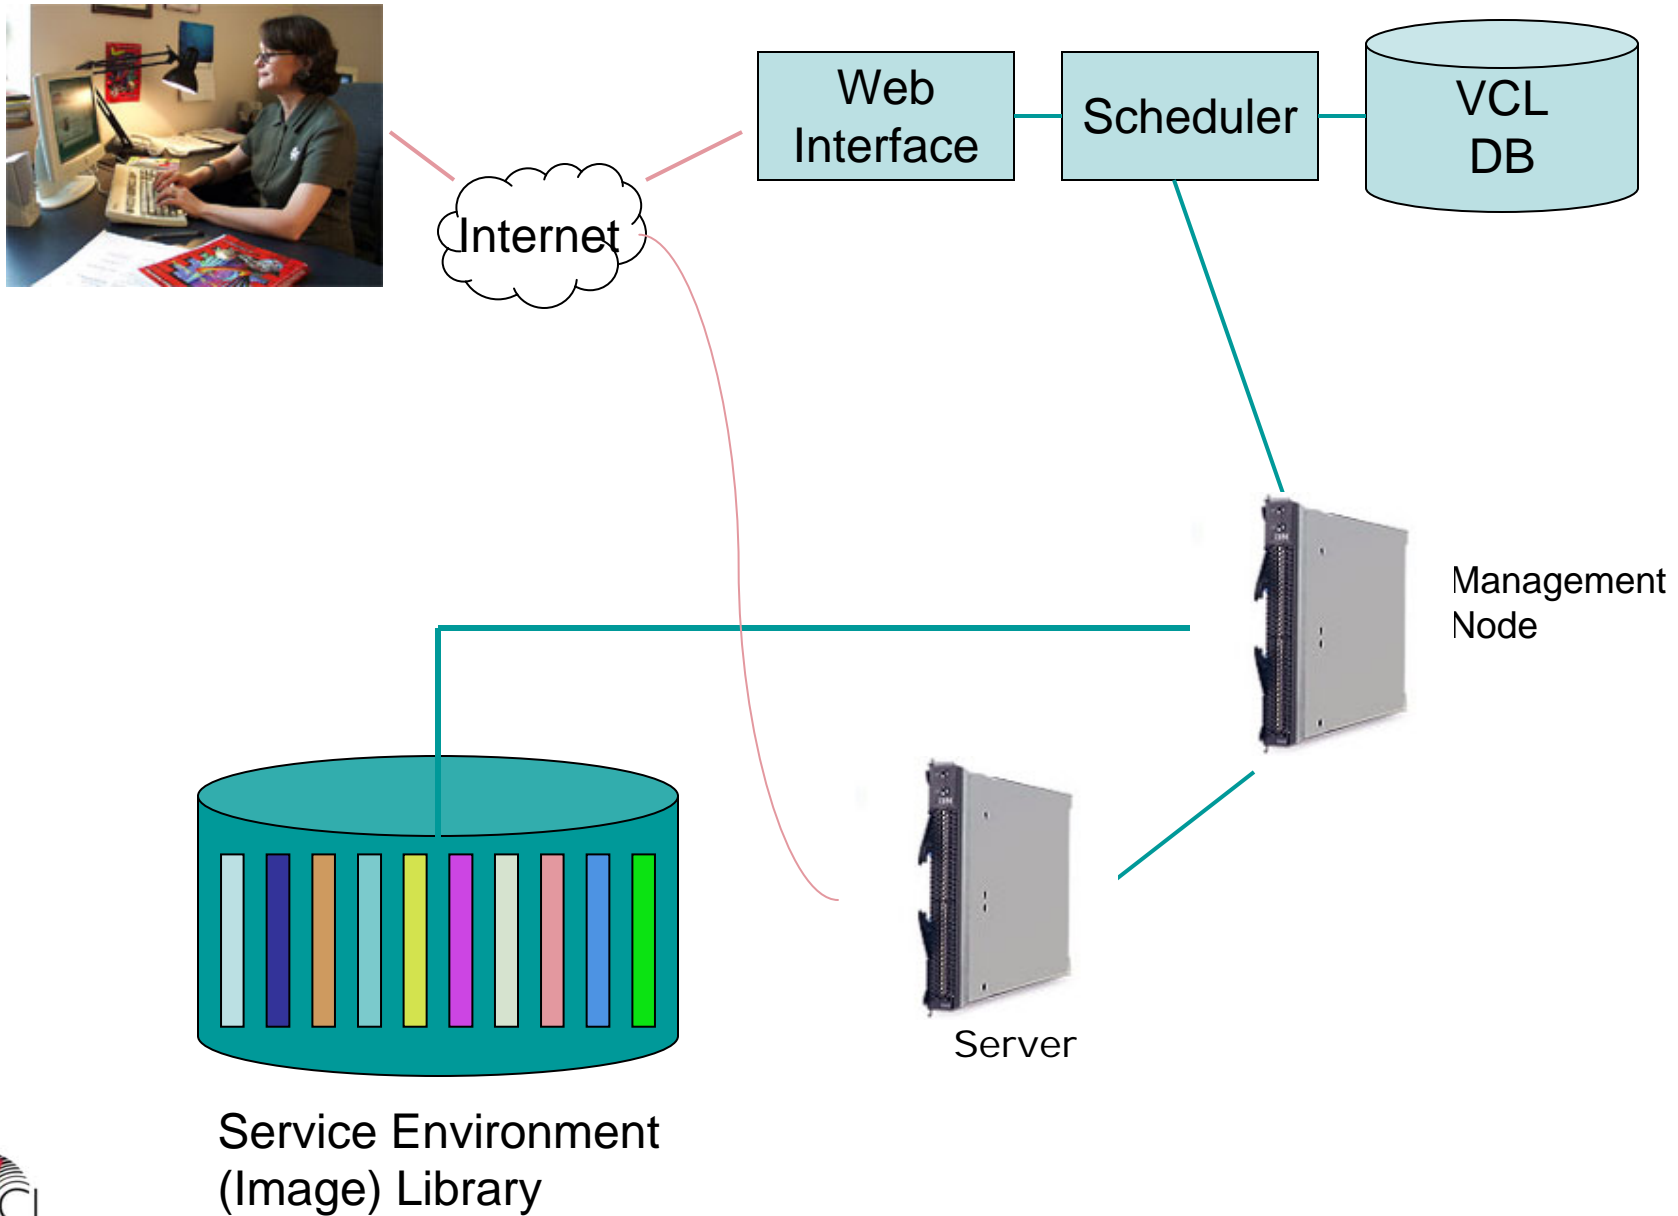
Desktop, Cloud, HPC

Write-ups


<http://vcl.ncsu.edu/papers-publications>



Typical Student Computing, Desktop Augmentation, Use of VCL



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.

