NC STATE UNIVERSITY

"Powered by VCL" - Using Virtual Computing Laboratory (VCL) Technology to Power Cloud Computing

Mladen A.Vouk, Sam Averitt, Michael Bugaev, Andy Kurth, Aaron Peeler, Henry Shaffer, Eric Sills, Sarah Stein, Josh Thompson

North Carolina State University

Copyright Mladen A. Vouk et al. 2008. This work is the intellectual property of the authors. Permission is granted for this material to be shared for non-commercial, educational purposes, provided that this copyright statement appears on the reproduced materials and notice is given that the copying is by permission of the authors. To disseminate otherwise or to republish requires written permission from the authors.

Cloud Computing?

- **The next step in the evolution of distributed access to computational resources.**
- **A trusted, dependable, usable, pervasive and economical on-demand network-based broad-capability utility/service for access to (and delivery of) computer, storage, network, and software solutions, products and services primarily based on virtualized resource.**
- **Service-Oriented Architecture**
  - Tightly and loosely coupled systems and services.
- **Scalable - Exascale data sets are (almost) here, exascale computing capabilities are in range.**
- **Workflow-oriented**
  - "Flavoured" – e.g., e-Learning cloud, or e-Government cloud, or an Analytics cloud, …
- **Easy to use**

Brief History

- 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"

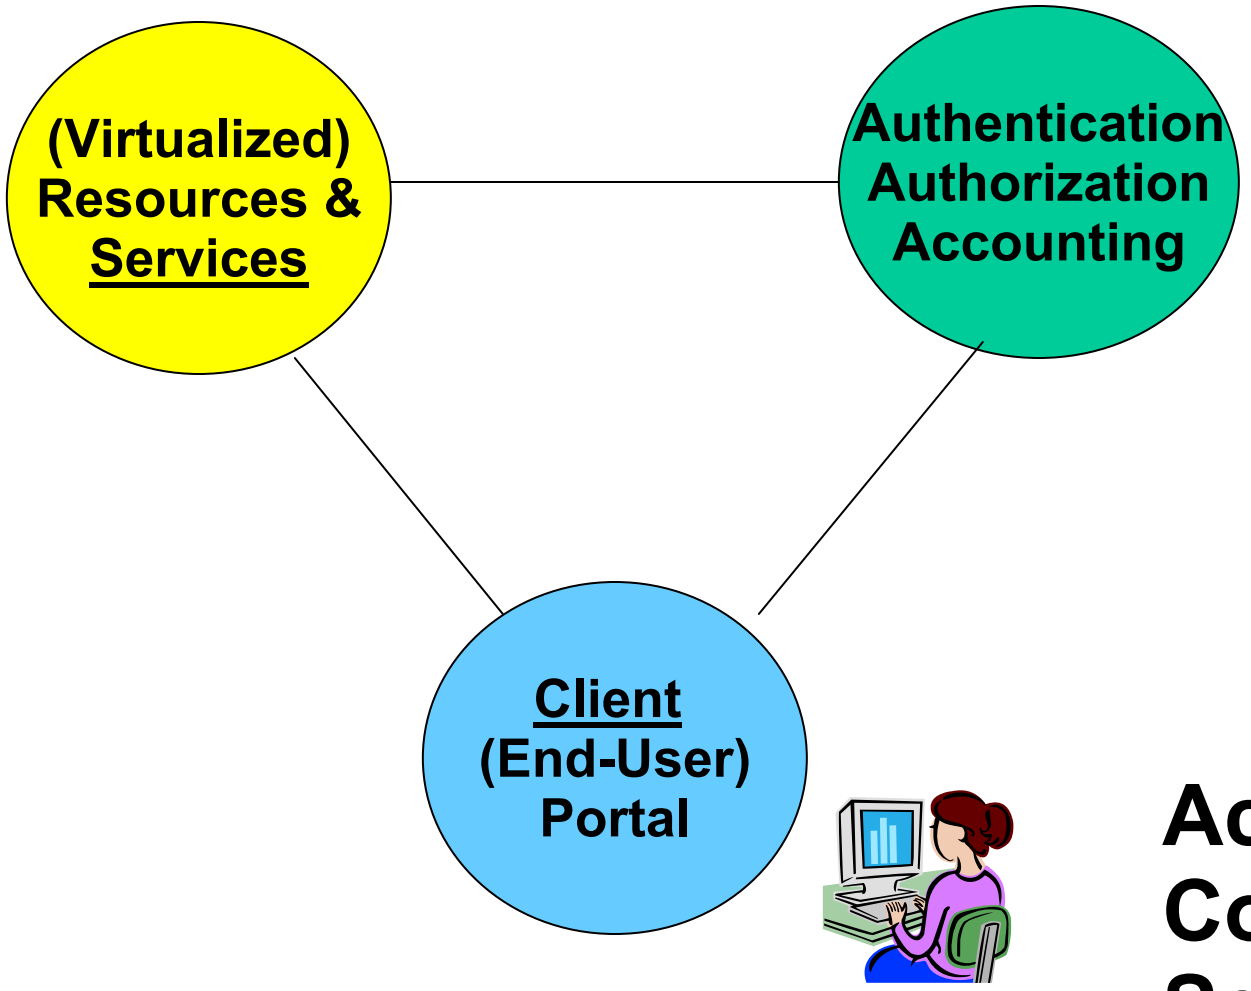

Access
Content
Services

VCL
Manager
& Scheduler

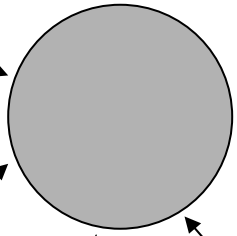

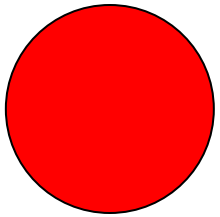

Authentication
Service

Internet

NC State Computational "Cloud" is powered by VCL

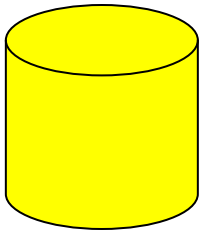

VCL Database

Node Manager #1

Node Manager #2

Node Manager #n

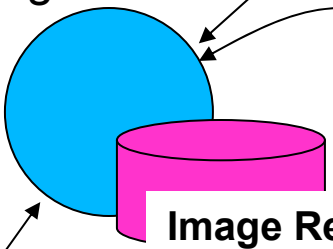

Image Repository

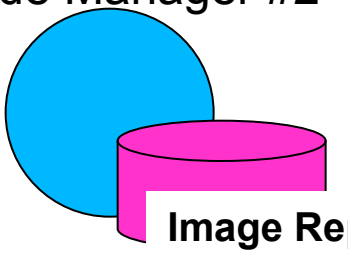

Image Repository

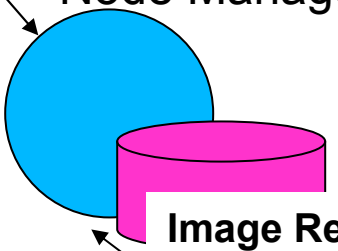

Image Repository

University Labs

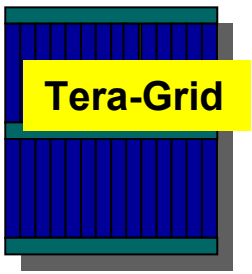

Tera-Grid

z-Series

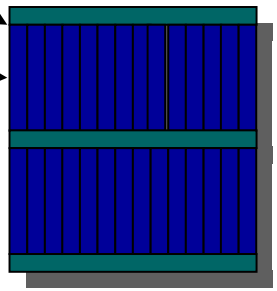

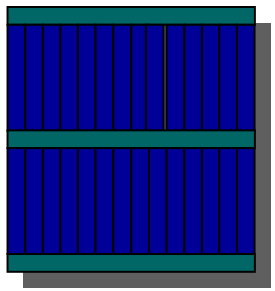

Virtual or Real Differentiated
Resources

Virtual or Real
Undifferentiated Resourses

Dynamic Re-Configuration

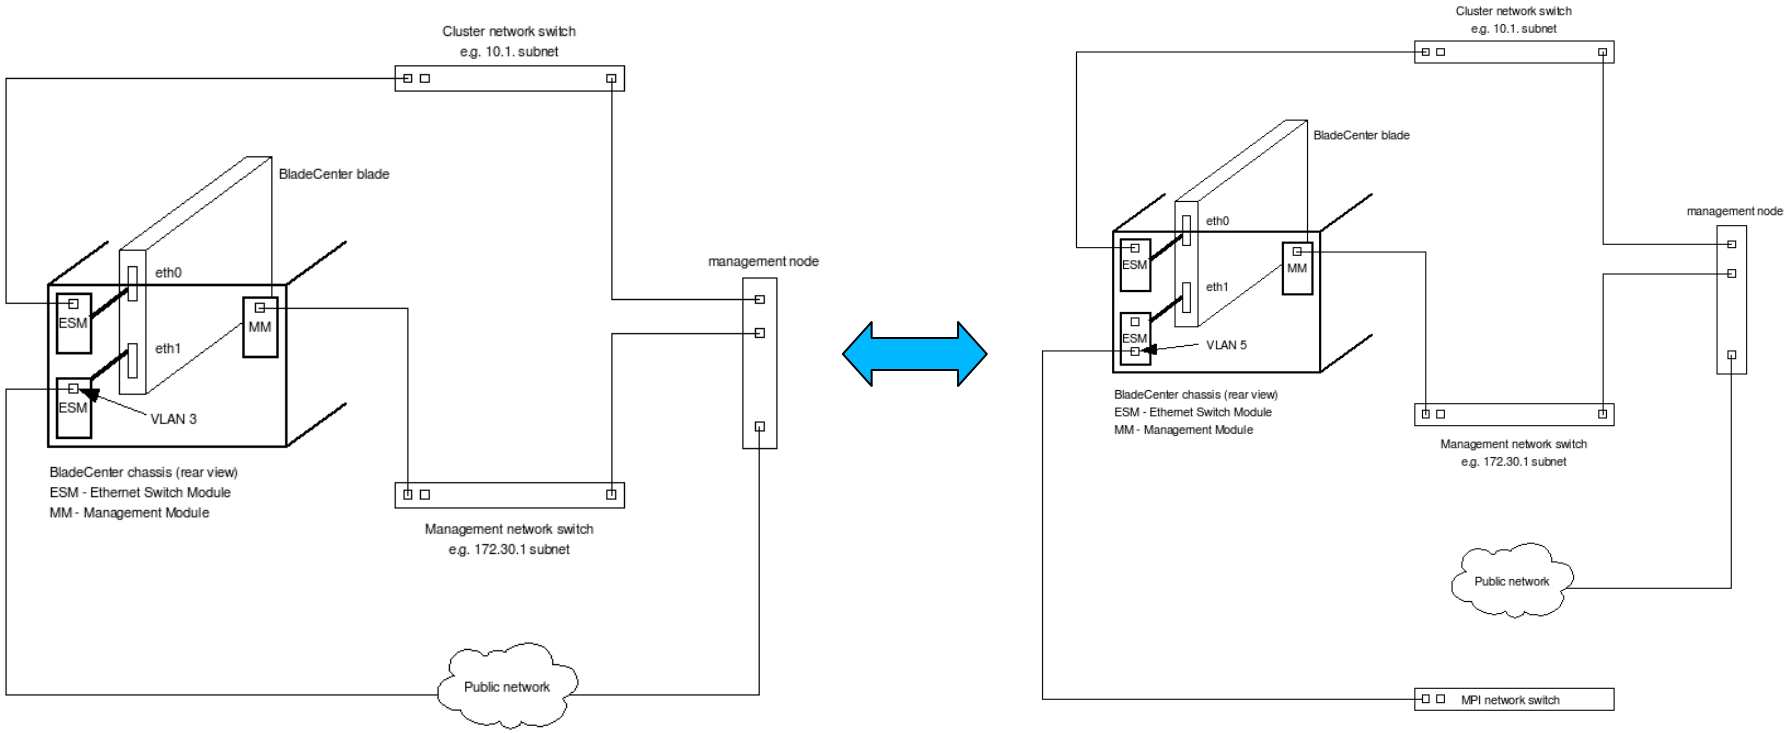

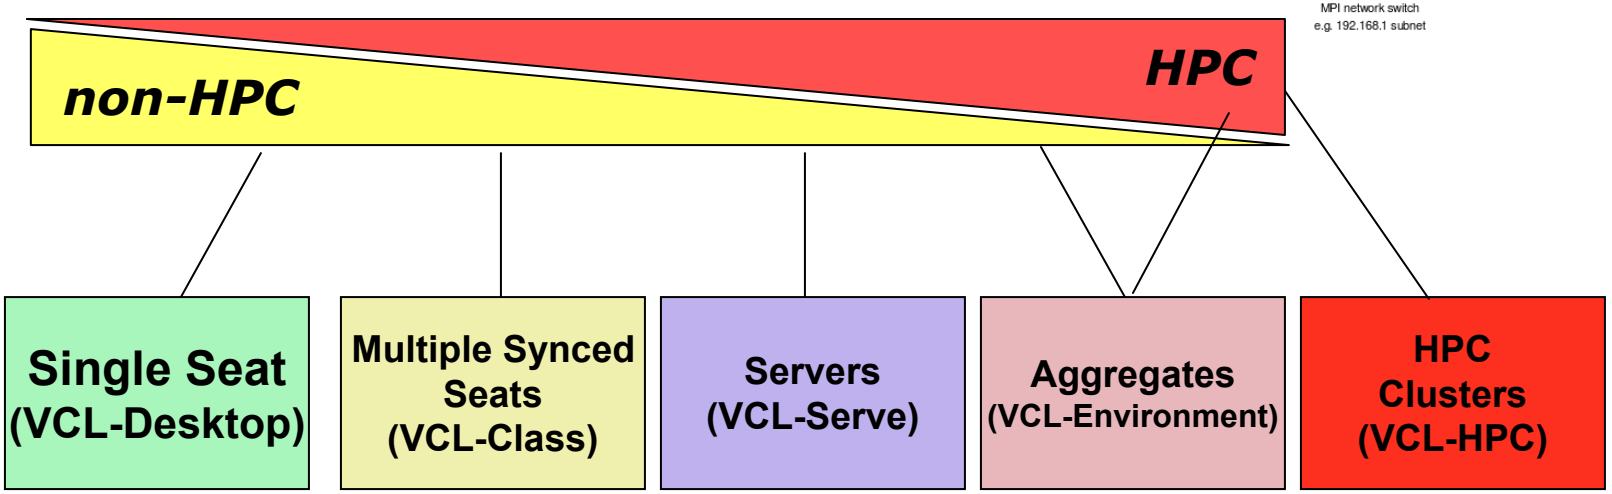

Service Composition

- **Image blocks**
- **More complex images**
- **Image groups**
- **Workflow construction**
- **Base-line images (e.g., XP, Linux)**

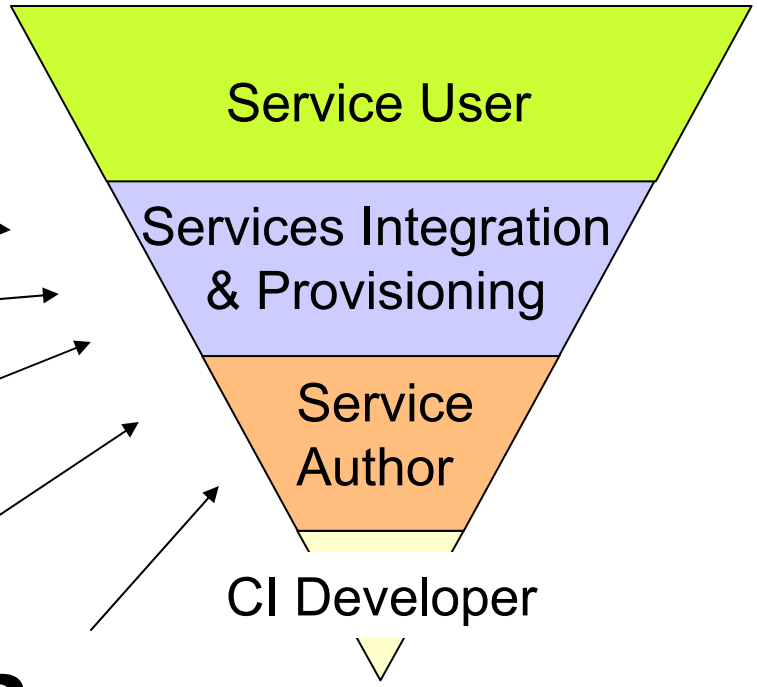

Components

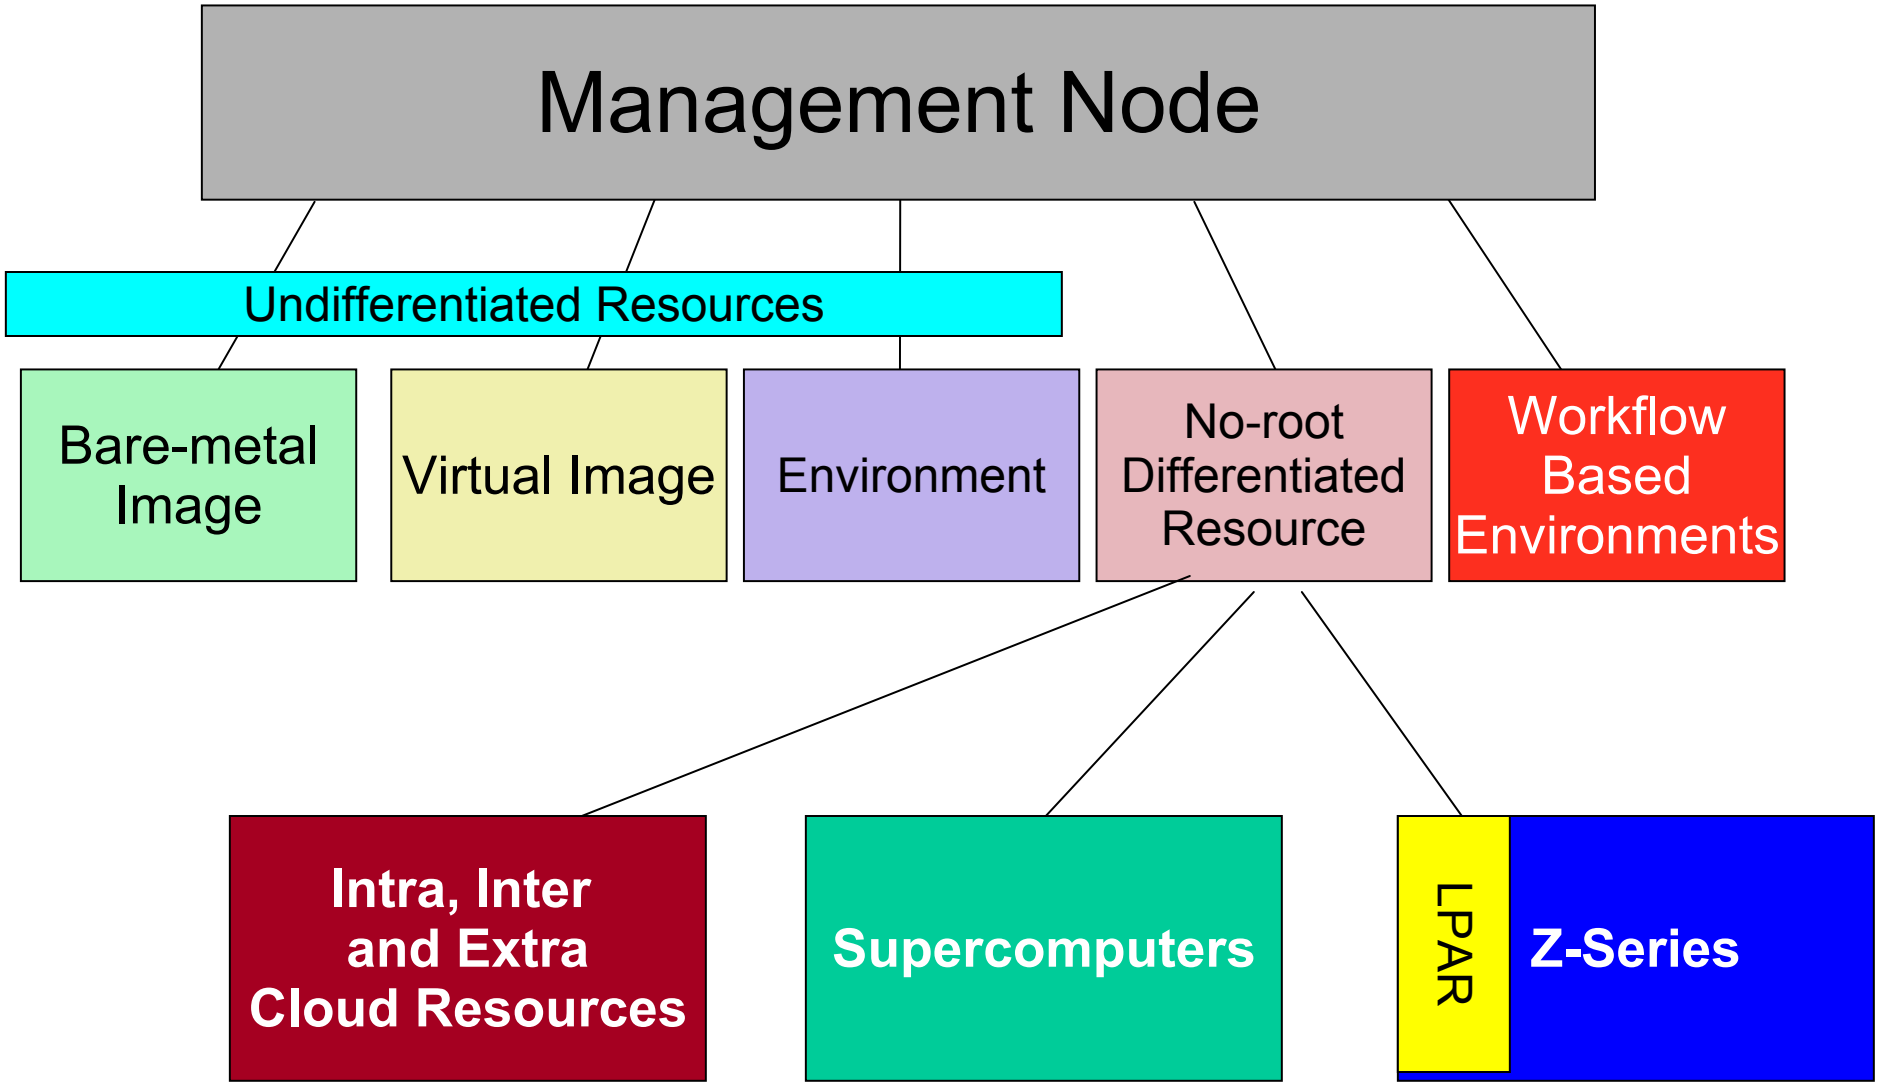
Management Node
Undifferentiated Resources
Bare-metal Image
Virtual Image
Environment
No-root Differentiated Resource
Workflow Based Environments
Intra, Inter and Extra Cloud Resources
Supercomputers
LPAR
Z-Series

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

Users

"Images"

"Application" Image Stack

Images
&
Environments

VCL
Manager

H/W Resources

Hardware
Blades, servers, desktops, storage …

Undifferentiated Local or distributed

xCAT	VCL code	IBM TM
WebServer	DataBase	Etc.

Reliability, Component-Based,
Scalability, Economy

Meta-Data & Monitoring

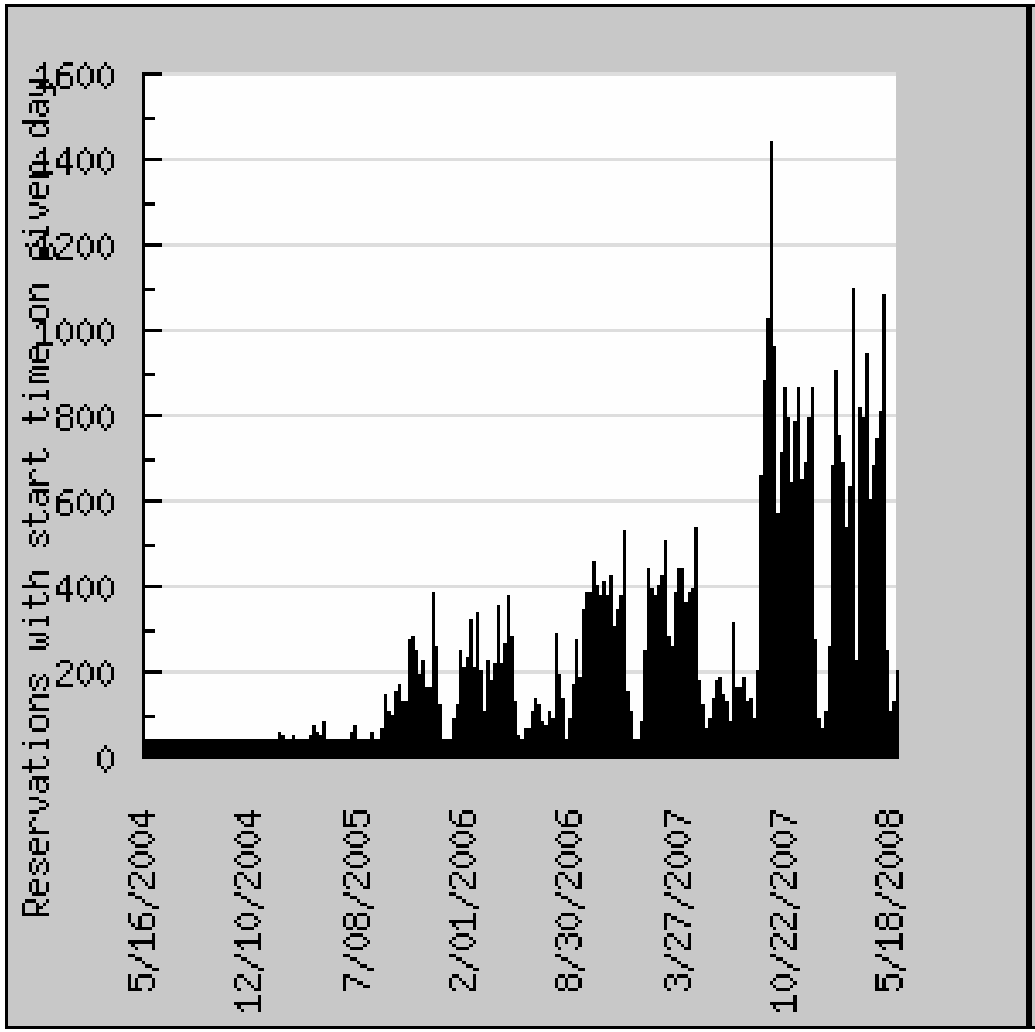

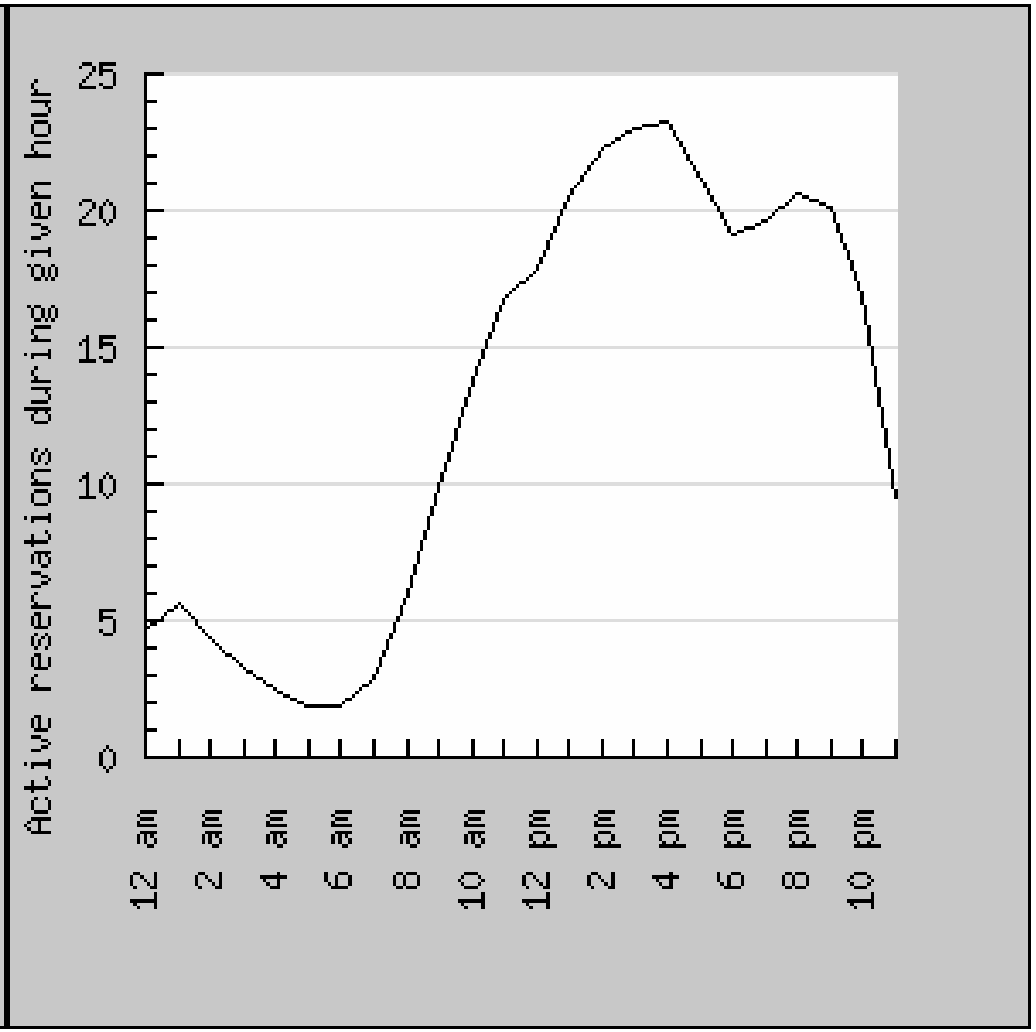

Some Stats

- About 1800+ blades (cca 200-500 used for VCL individual seats and services, the rest for VCL HPC cycles), plus several hundred idle student laboratory machines.
- 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.
- Currently Cca 600 images, cca 120 in use per semester.
- About 60-100,000 image reservations per semester.
- Serving population of 30,000+ students (in a semester there may be about 10,000 unique users).
- Most of the "individual seat" requests are on-demand "Now" reservations: cca 90% of requests
- System availability: about 99%

Economics

- Typical NC State bare-metal blade serves about 25 students seats – 25:1 ratio – considerably better than tradtional labs at 5:1 to 10:1.
- Gains come from time and diversity utilization.
- Hypervisors and server-apps increase utilization by another factor of 2 to 40 depending on the application and user profile.
- Avg. 1 FTE maintenance for about 1000+ nodes.
- Avg. 3+ FTE in development
- Typical user reservation is 1-2 hours

VCL™
VT4c/Oct08/Copyright © Vouk 2008

Top Requirements

- **Reliability, Agility, Usability, "Green" …**
- **Efficient "image" and service construction**
  - Portability of images and solutions (OVM format)?
- **Provenance and meta-data**
- **Workflow automation (Kepler?)**
- **Security and Policy**
- **Return on Investment (ROI) and Total Cost of Ownership (TCO)**
- **Etc.**

Development and Research

- Security and networking (e.g., end-to-end isolation, image security/"water marking", security of hypervisors)
- Service composition and management
- Image format, re-usability, deployment
- Next generation education paradigms and applications
- Other …

V14c/Oct08/Copyright © Vouk 2008

Awards

- *"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*

What is VCL?

- Hardware abstraction
  - Can deliver environments:
    - on bare metal
    - on top of hypervisors
    - moves virtualized environments between machines as needed
  - Users don't need to be concerned with how their environments are being delivered

What Makes Up VCL

- Backend system
  - **Web server**
    - **Linux, Apache, PHP**
    - **frontend - VCL UI, scheduler, administration**
  - **Database server**
    - **Linux, MySQL 5**
  - **Management Node(s)**
    - Linux, perl, XCAT
    - vcld – backend; touches hardware, makes things happen
  - **Blades or standalone servers**
  - **Separate servers or combined on single server**

What Makes Up VCL

- Compute nodes
  - IBM BladeCenter Blade Servers
    - Housed in a datacenter
  - Standalone workstations
    - Housed anywhere; we include our lab machines when the labs are closed
  - Working on Sun Blade servers
  - VCI partners are working Dell and HP blades
  - Can easily be moved between HPC cluster and VCL system
    - We move nodes to HPC during student breaks

License

- **Open Source**
  - Eclipse (approved)
  - (L)GPL (under consideration)
  - Apache (approved)