UNIVERSITY OF KENTUCKY®

Curriculum in the Cloud: Using OpenStack to Transform Computer Science Education

Dr. Brent Seales, Dr. Cody Bumgardner, C. Lowell Pike
University of Kentucky

UK
UNIVERSITY OF KENTUCKY

https://www.pcworld.com/article/2972895/computers/9-awesome-photos-of-school-computer-labs-from-the-1980s.html

VT-100 series dumb terminals (1986)

With the rise of personal computers, it became very rare in the mid-1980s to witness a scene like this one from Boston College High School in Boston, MA. Here we see a lab full of DEC VT-100 series dumb terminals, which were not computers unto themselves, but mere gateways to a bigger computer somewhere else. They would have been connected to a mainframe computer such as a DEC PDP-11 or VAX via RS-232 serial links. This was fairly advanced for a high school at the time.

After the 1980s, school computer labs in the U.S. were dominated mostly by Apple and IBM machines. Today, computer labs are becoming an endangered species, with more schools relying on notebook PCs and tablets that don't have to be tied to a particular room. But wherever they're used, computers remain an important part of education in the U.S.

UNIVERSITY OF KENTUCKY

The Options

- Conventional hardware labs
- The Cloud (3rd party “public cloud”)
  - Microsoft Azure
  - Amazon Web Services (AWS)
  - Google Cloud Platform
- OpenStack Private Cloud

Moving CS Curriculum to the Cloud

- Student access to VMs
  - Need to use IDEs from varied platforms in the classroom
  - Campus infrastructure: wifi, back-end, admin
- Authentication
  - Security, privacy
  - Integration with University systems
- Resources
  - Will it cost more than computer labs?
- Needs of the CS curriculum

Table 5-1: Courses Term by Term with the average enrollments in the recent semesters

Course (Department, Number, Title)	R/E/SE	*Subject Area (Credit Hours)* M/S	CS	GEd	Oth	Semester	Enrollment
Freshman Year							
CS 100 The Computer Science Profession	R		F (1)			F(14),F(15)	179
CS 115 Introduction to Computer Programming	R		F (3)			F(15),S(16)	Lec: 120 Lab: 18
CIS/WRD 110 Composition and Communication I	R			G (3)		F(15),S(16)	CIS/WRD 24/21
MA 113 Calculus I	R	M (4)				F(15),S(16)	Lec: 150 Rec 31
UK Core [U]	SE			G (3)		F(15),S(16)	
CS 215 Introduction to Program Design Abstraction	R		F (4)			F(15),S(16)	Lec: 93 Lab: 20
Natural Science Elective [N]	E	S (3)				F(15),S(16)	
MA 114 Calculus II	R	M (4)				F(15),S(16)	Lec: 114 Rec: 25
UK Core Statistical/Inferential Reasoning [U]	SE			G (3)		F(15),S(16)	
Sophomore Year							
CS 216 Intro to Software Engineering Techniques	R		F (3)			F(15),S(16)	Lec: 127 Lab: 29
EE 280 Design of Logic Circuits	R		F (3)			F(15),S(16)	75
MA 213 Calculus III	R	M (4)				F(15),S(16)	Lec: 104 Rec: 30
PHY 231 General University Physics	R	S (4)				F(15),S(16)	Lec: 169 Rec: 36
PHY 241 General University Physics Laboratory	R	S (1)				F(15),S(16)	28
CIS/WRD 111 Composition and Communication II	R			G (3)		F(15),S(16)	CIS/WRD: 25/23
CS 275 Discrete Mathematics	R	M (4)				F(15),S(16)	Lec: 99 Lab: 50
CS/EE 380 Microcomputer Organization	R		F (3)			F(15),S(16)	CS/EE 380: 47/81
PHY 232 General University Physics	R	S (4)				F(15),S(16)	Lec:127 Rec: 29
PHY 242 General University Physics Laboratory	R	S (1)				F(15),S(16)	26
STA 281 Probability and Statistics	R	M (3)				F(15),S(16)	59
UK Core [U]	SE			G (3)		F(15),S(16)	
Junior Year							
CS 315 Algorithm Design and Analysis	R		F (3)			F(15),S(16)	46
CS/MA 321 Introduction to Numerical Methods	R		F (3)			F(15),S(16)	CS/MA 321: 30/16
UK Core [U]	SE			G (3)		F(15),S(16)	
Elective [E]	E				O (3)	F(15),S(16)	
Elective [E]	E				O (3)	F(15),S(16)	
CS 375 Logic and Theory of Computing	R		F (3)			F(15),S(16)	64
Computer Science Elective [C]	SE		A (3)			F(15),S(16)	
Technical Elective [T]	SE				O (3)	F(15),S(16)	
UK Core [U]	SE			G (3)		F(15),S(16)	
Natural Science Elective [N]	E	S (3)				F(15),S(16)	
Elective [E]	E				O (3)	F(15),S(16)	
Senior Year							
CS 470G Introduction to Operating Systems	R		F (3)			S(15),S(16)	48
Computer Science Elective [C]	SE		A (3)			F(15),S(16)	
Technical Elective [T]	SE				O (3)	F(15),S(16)	
UK Core [U]	SE			G (3)		F(15),S(16)	
Elective [E]	E				O (4)	F(15),S(16)	
CS 499 Senior Design Project	R		A (3)			F(15),S(16)	38
Computer Science Elective [C]	SE		A (3)			F(15),S(16)	
Technical Elective [T]	SE				O (3)	F(15),S(16)	
Technical Elective [T]	SE				O (3)	F(15),S(16)	
Elective [E]	E				O (3)	F(15),S(16)	
TOTAL-ABET BASIC-LEVEL REQUIREMENTS		35	41	24	28		
OVERALL TOTAL CREDIT HOURS FOR COMPLETION OF PROGRAM	128						

UNIVERSITY OF KENTUCKY

Figure 5-1: Prerequisite Flowchart

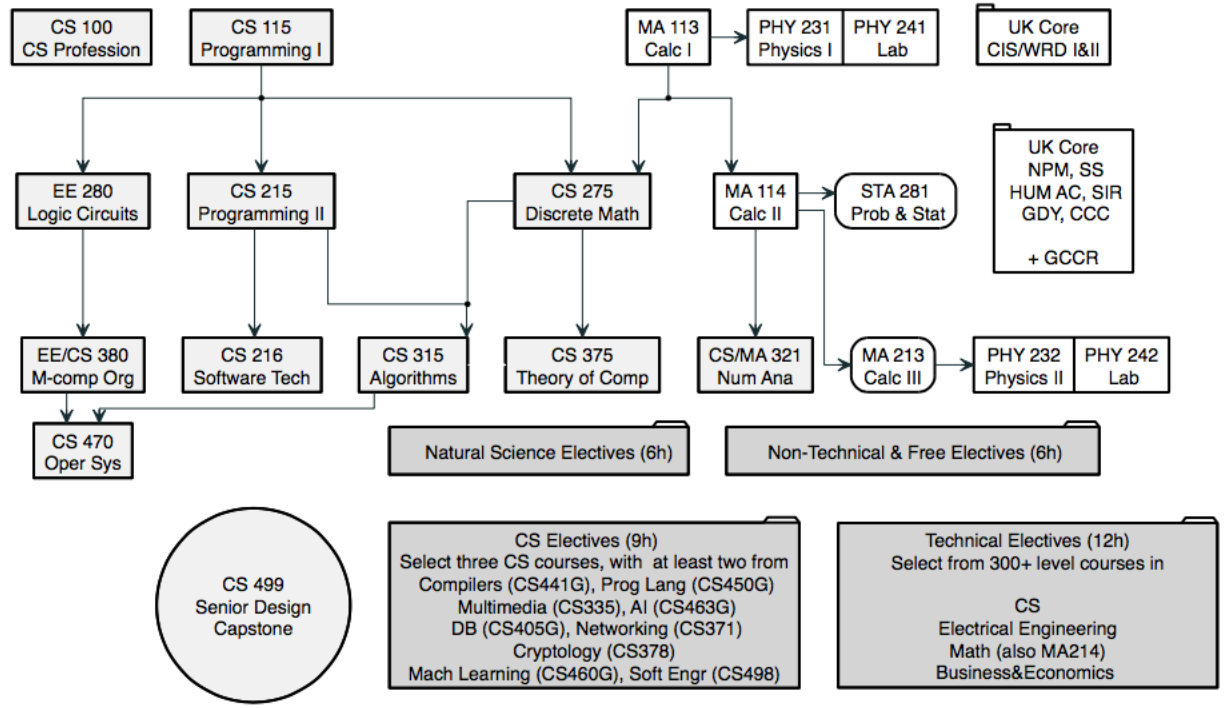

New Paradigm

- Mobility: "Old Man Shouts at Cloud"
  - We stayed all night in the lab…
  - …and we LIKED it
- Cross platform
  - Students have invested in their own hardware
- Scalable admin costs
- Automatic upgrades
- Security, privacy, authentication

Practical Issues

- Authentication with campus network
- Performance choices
- Provided “hardware” platforms and software needs for individual professors and courses
- Terminal and Desktop access (ssh; NoMachine)
- Admin
  - Provisioning for the curriculum
  - Provisioning for performance
  - Monitoring of performance (automated)

Agenda

- OpenStack services selection
- Keystone: Integrating with campus authentication
- Neutron networks choices
- Glance image and provisioning
- Monitoring and Analytics
- Scaling and performance
- Security
- Student pampering
- Cost comparison with public cloud

UNIVERSITY OF KENTUCKY

Open source software for creating private and public cloud

OpenStack choices: mature and stable services, manual installation

- Identity service - “Keystone”
- Image service – “Glance”
- Compute service – “Nova”
- Networking service – “Neutron”
- Block storage service – “Cinder”
- Dashboard service – “Horizon”

UK
UNIVERSITY OF KENTUCKY

Keystone: Integrating with Campus Authentication

- Hybrid authentication – try local, then ldap
- keystone/identity/backends/sql.py

```
if not password_hashing.check_password(password,user_ref.password):
    try:
        conn = ldap.initialize('LDAP://<SERVER>/')
        try:
            conn.simple_bind_s("   <DOMAIN>\\{}".format(unicode(user_ref.name)), unico
                return True
            except ldap.INVALID_CREDENTIALS:
                return False
        except ldap.SERVER_DOWN:
            return False
    return True
```

UNIVERSITY OF KENTUCKY

Neutron networks

- Linux bridge over OVS for simplicity
- Provider over self service for performance
- Networks: private, campus, campus jumbo, campus IPv6, campus SDN research network

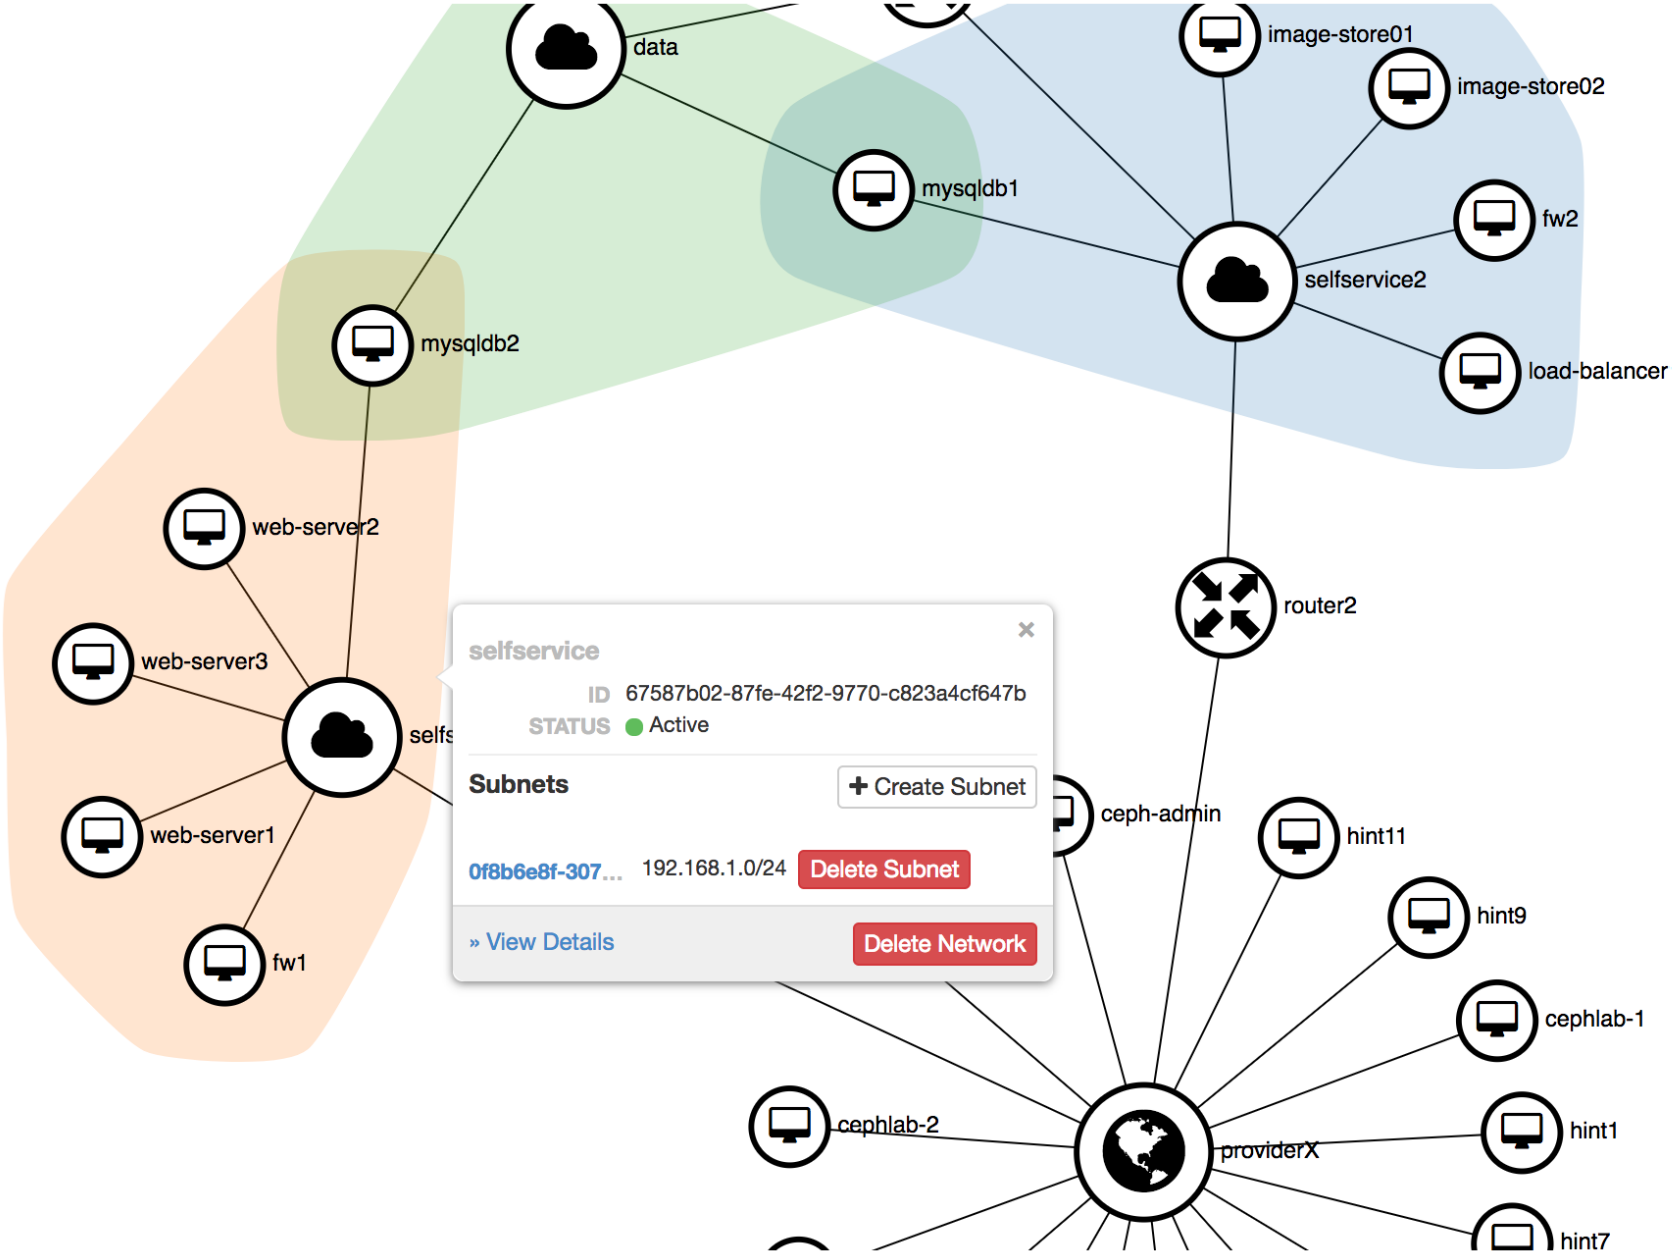

UNIVERSITY OF KENTUCKY

Glance: VM Images

- Linux desktop, Windows 10 desktop, and generic server images (CentOS, Ubuntu, Debian, openSUSE, FreeBSD)
- Cloud-init for Linux and Cloudbase-Init for early initialization of VMs
- Course packages and IDEs preinstalled on desktop images
- NoMachine – free remote desktop

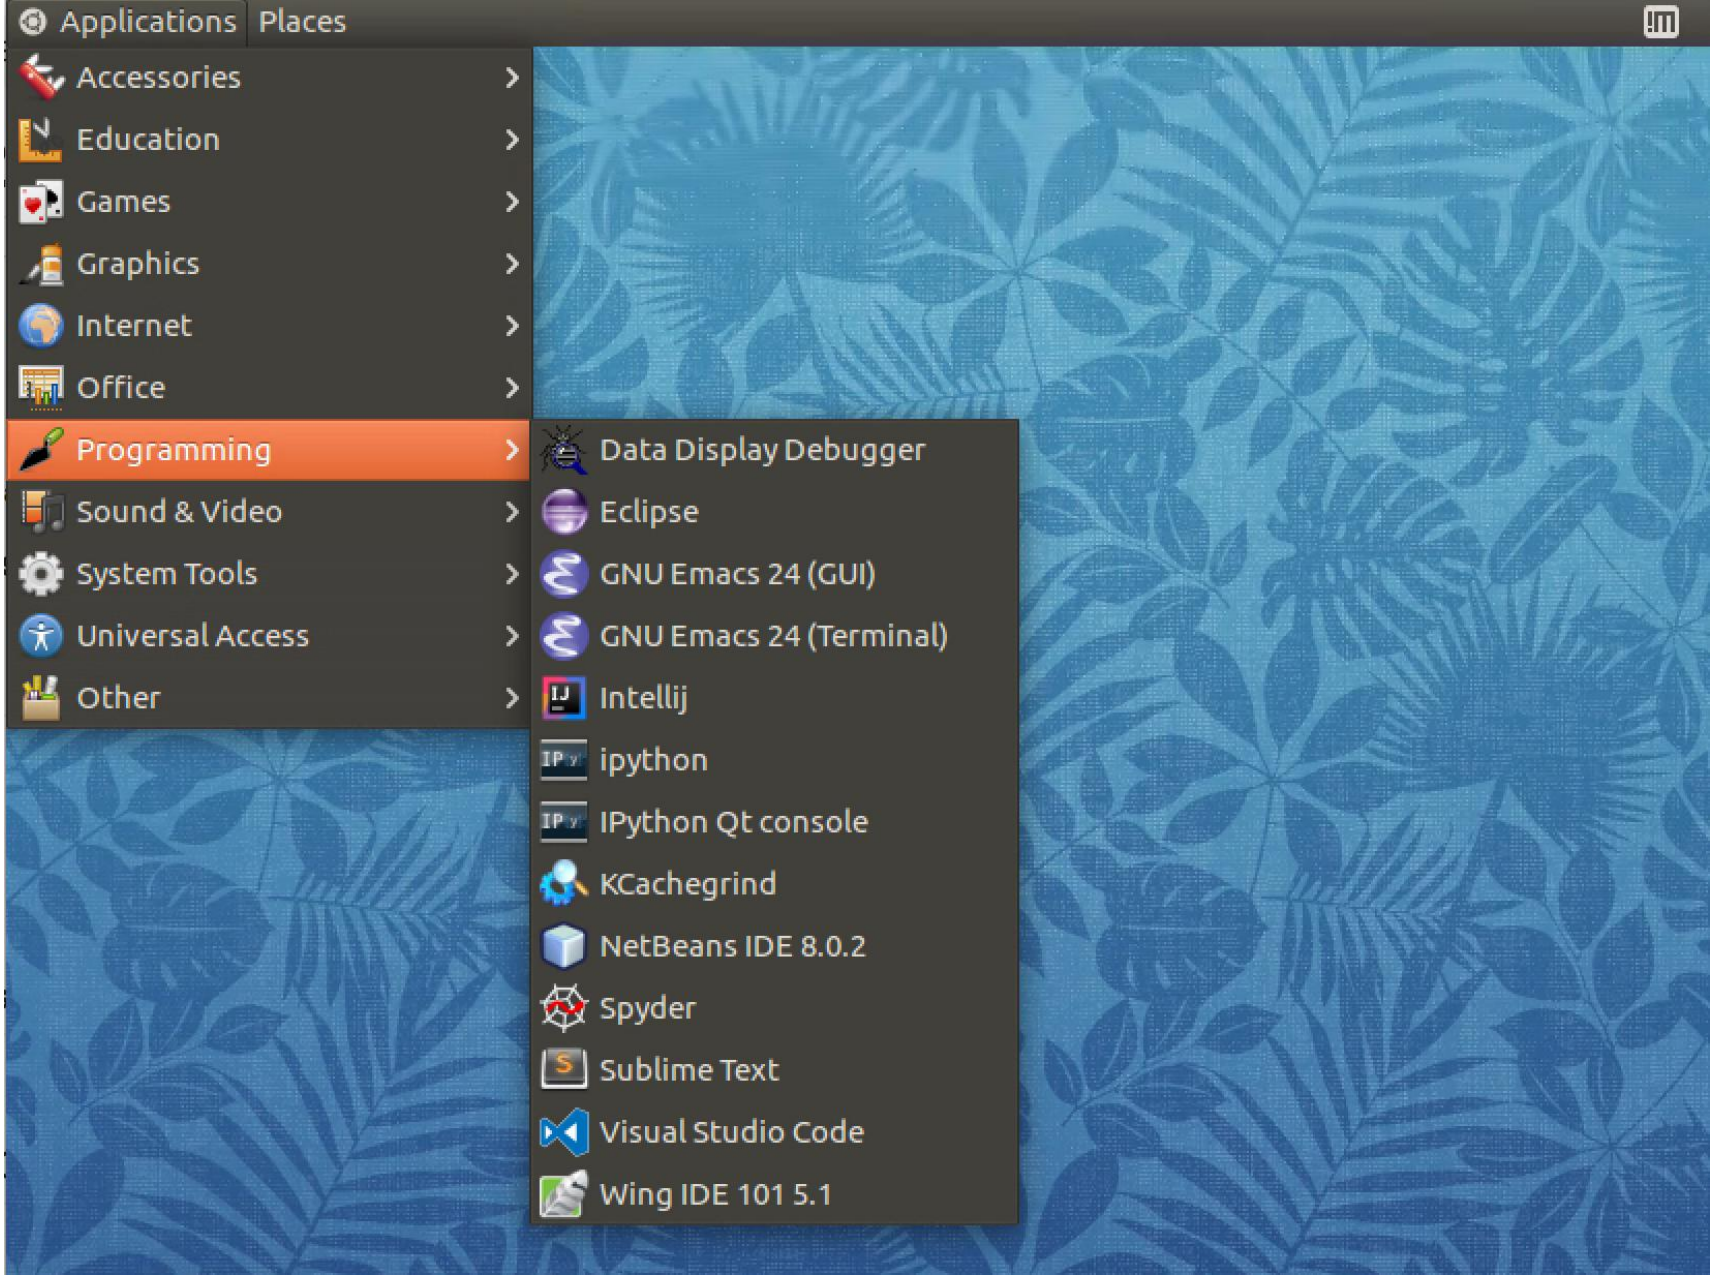
Applications
Places
Accessories
Education
Games
Graphics
Internet
Office
Programming
Sound & Video
System Tools
Universal Access
Other
Data Display Debugger
Eclipse
GNU Emacs 24 (GUI)
GNU Emacs 24 (Terminal)
Intellij
ipython
IPython Qt console
KCachegrind
NetBeans IDE 8.0.2
Spyder
Sublime Text
Visual Studio Code
Wing IDE 101 5.1


Programming
Sound & Video
System Tools
Universal Access
Other
Data Display Debugger
Eclipse
GNU Emacs 24 (GUI)
GNU Emacs 24 (Terminal)
Intellij
ipython
IPython Qt console
KCachegrind
NetBeans IDE 8.0.2
Spyder
Sublime Text
Visual Studio Code
Wing IDE 101 5.1

NoMachine Remote Destop

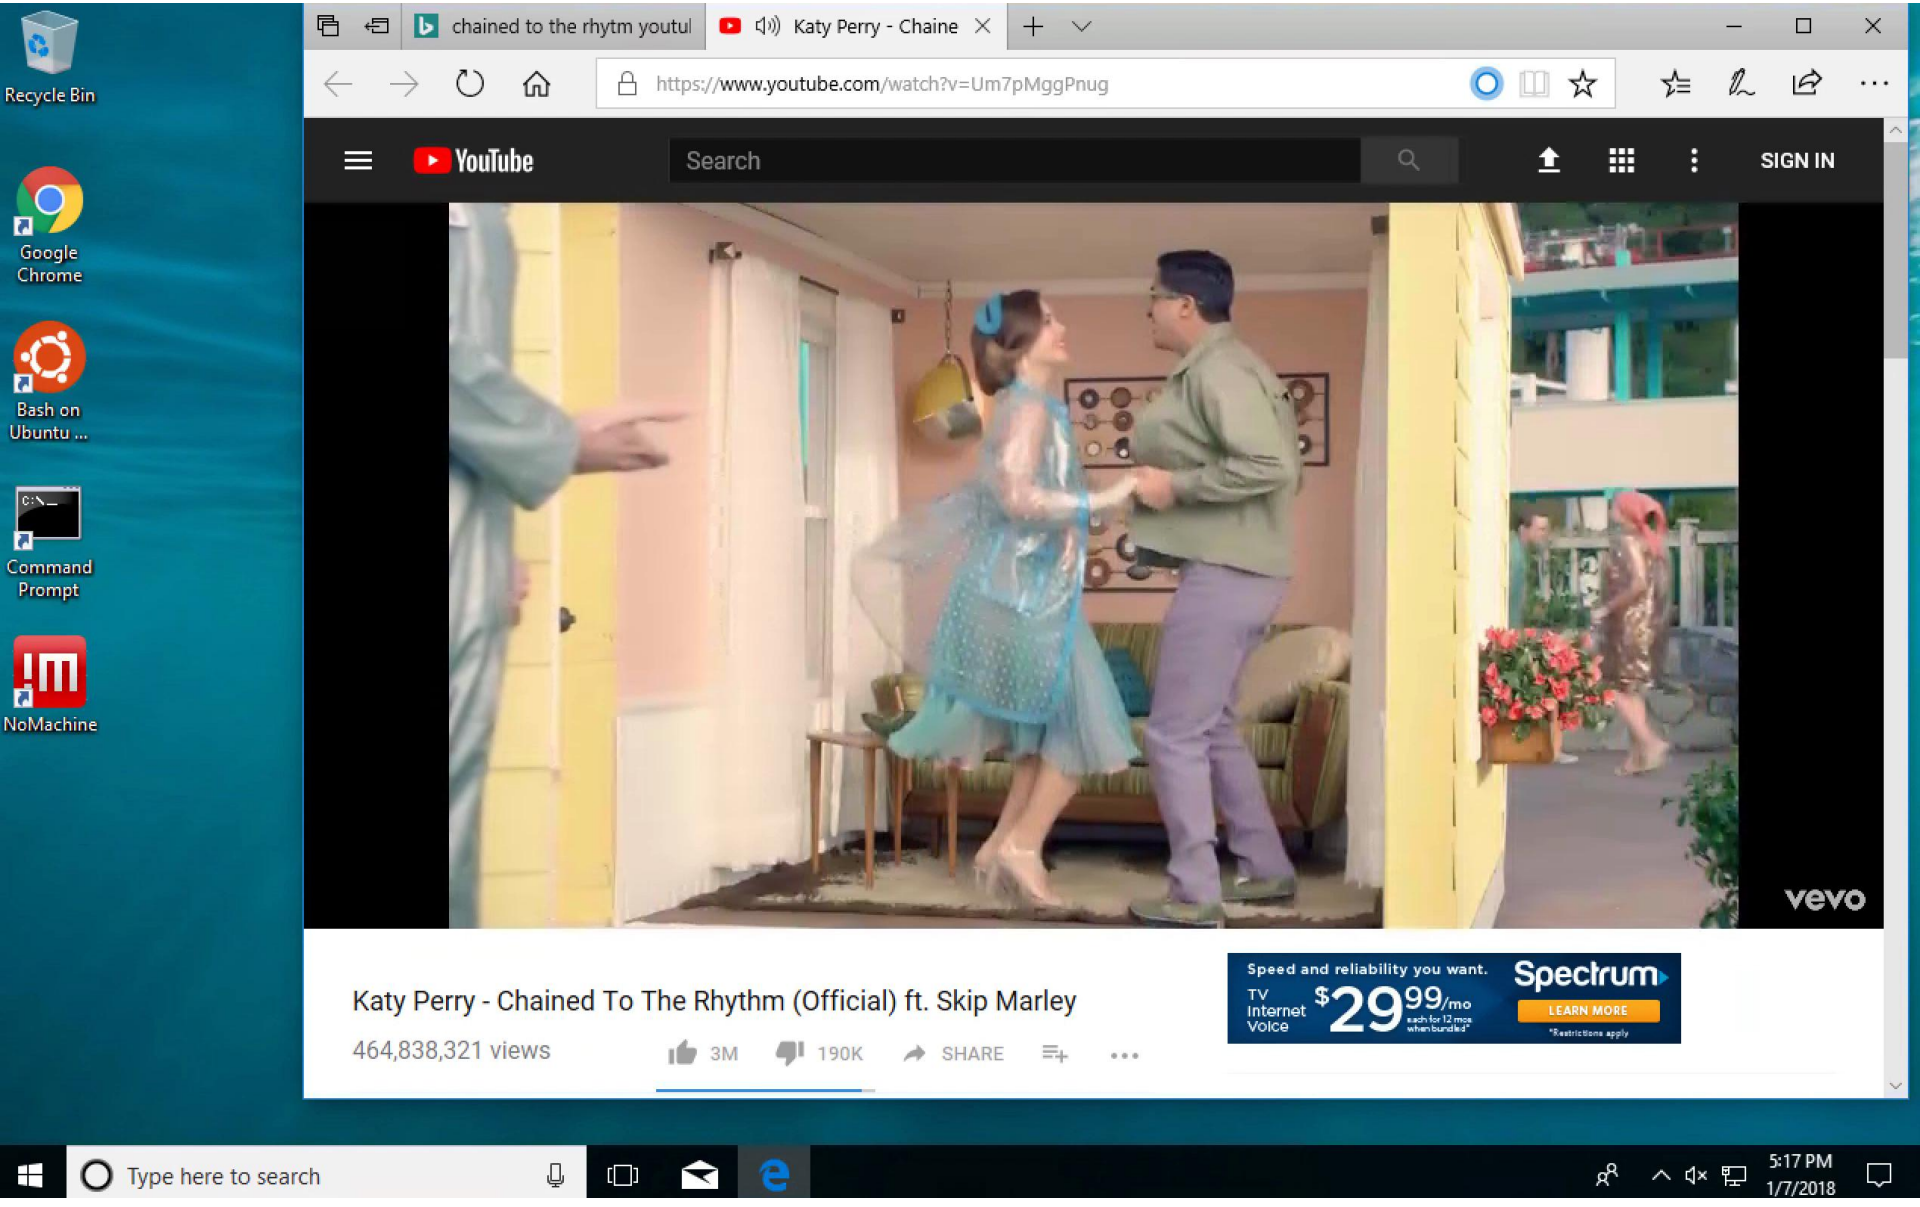

OpenStack CLI: Automating VM Provisioning

- Class roster -> grep awk manipulation -> openstack cli script

```
openstack project create
openstack user create
openstack quota set
openstack security group rule
openstack server create
openstack volume create
```

Monitoring and Analytics

- Server agents -> Time series database -> Visualization & Alerting
- Open Source
- Server load
- Network usage

Grafana

Visualizes and graphs the time-series data stored in InfluxDB.

influxdb

Delivers high performance writes and efficiently stores time-series data.

telegraf

Collects time-series data from a variety of sources

UK
UNIVERSITY OF KENTUCKY

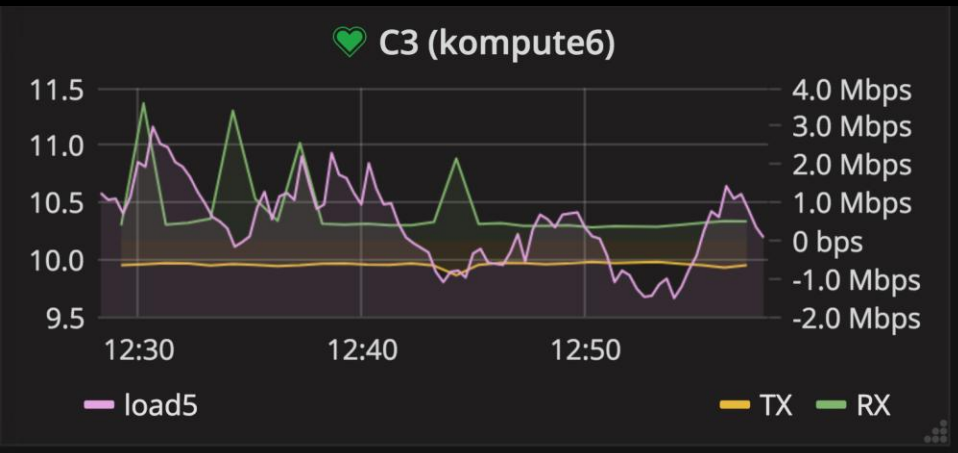

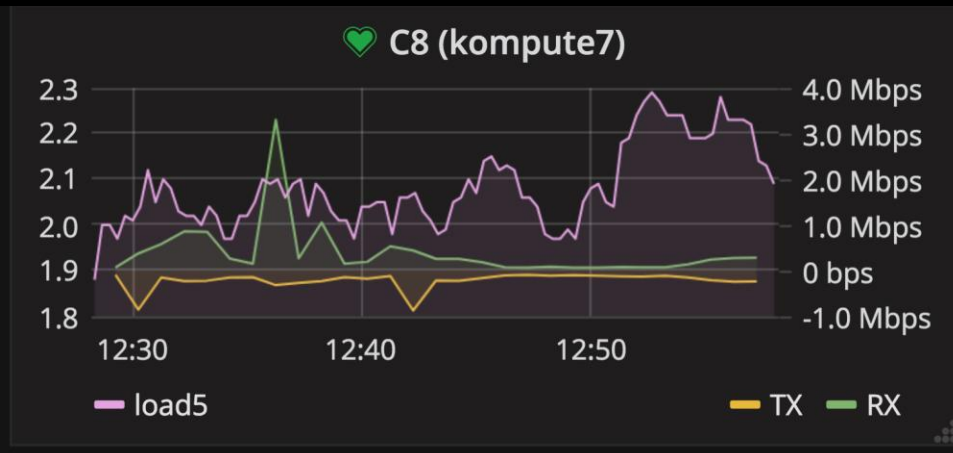

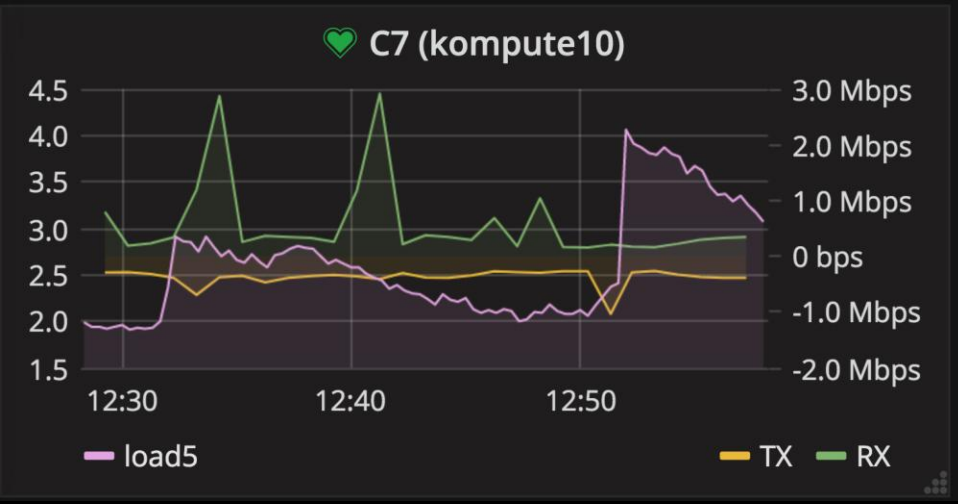

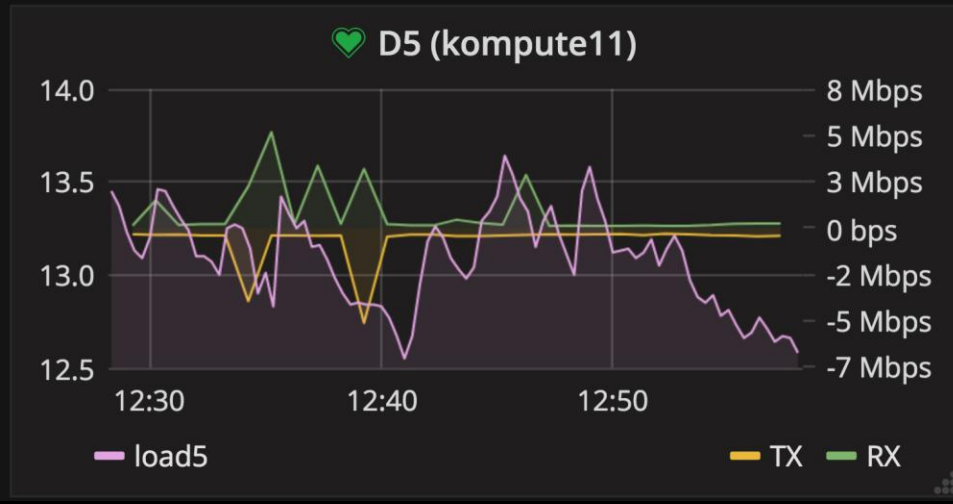

UNIVERSITY OF KENTUCKY

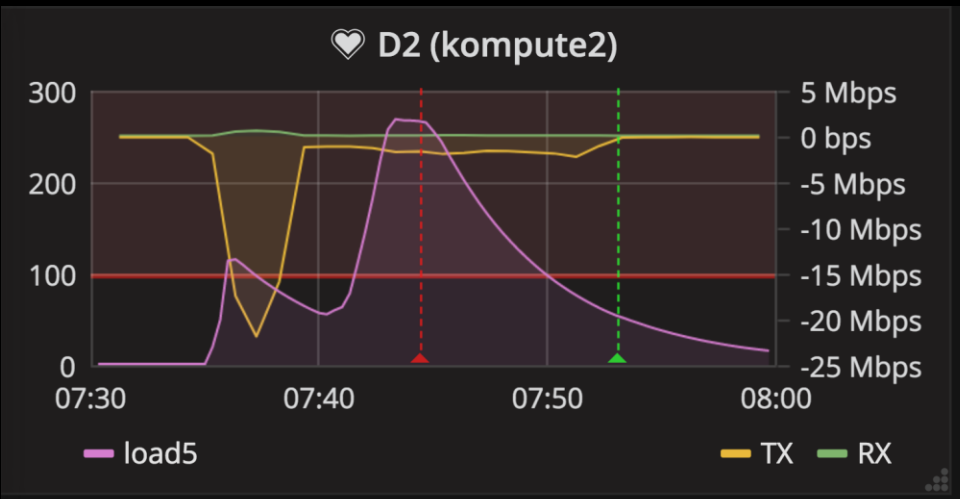

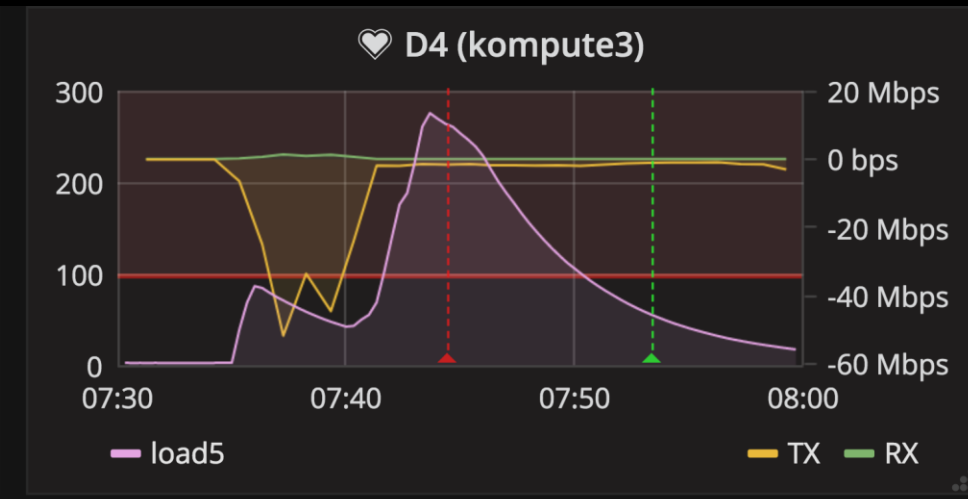

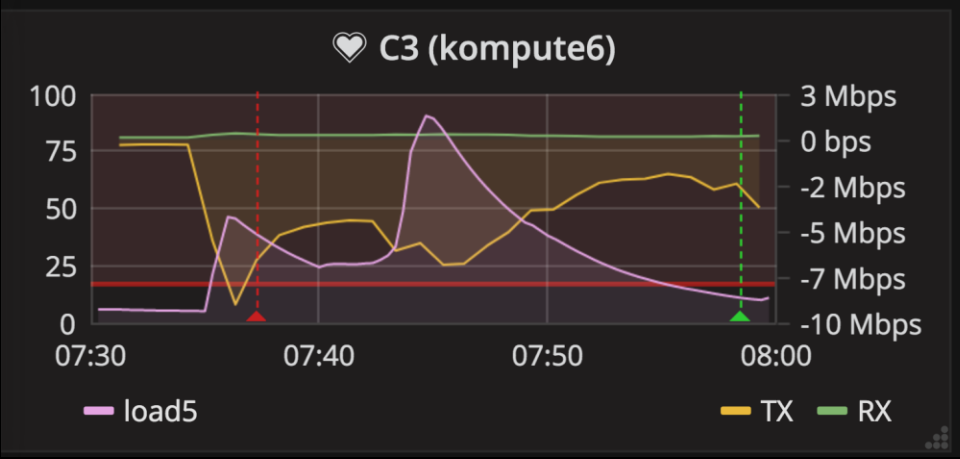

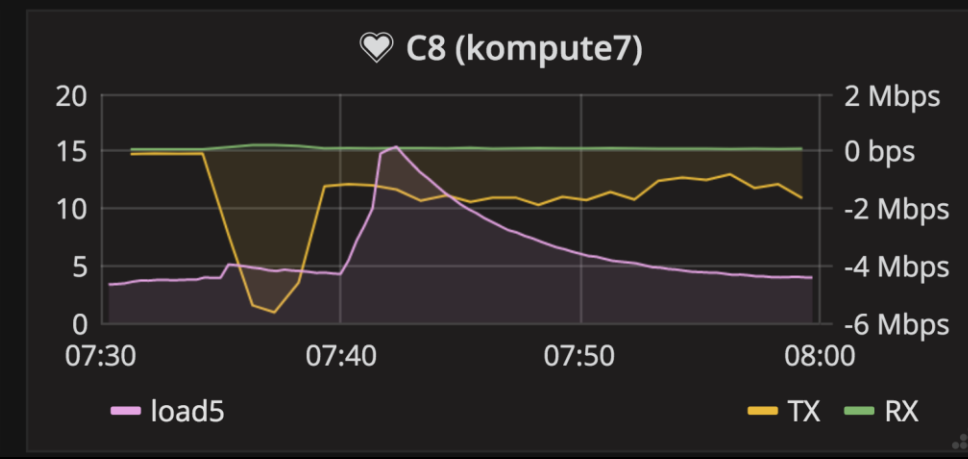

UNIVERSITY OF KENTUCKY

OS not always virtualization optimized

/etc/cron.d/anacron

30 7 * * * root test -x /etc/init.d/anacron && /usr/sbin/invoke-rc.d anacron start >/dev/null

5 3 * * * root perl -le 'sleep rand 7200' ; test -x /etc/init.d/anacron && /usr/sbin/invoke-rc.d anacron start >/dev/null

Scaling and Performance

- Stress test resource oversubscription
- Defaults: CPU 16/1, RAM 1.5/1, Disk 1/1
- We are memory bound
- Windows 10 requires more resources
- Compute node local disks for VMs
- Provider networks

Hypervisor Summary

VCPU Usage

Used 4,864 of 304

Memory Usage

Used 645GB of 2.2TB

Local Disk Usage

Used 1.6TB of 239.2TB

Hypervisor | Compute Host

Displaying 11 items

Hostname	Type	VCPUs (used)	VCPUs (total)	RAM (used)	RAM (total)	Local Storage (used)	Local Storage (total)	Instances
kompute1.netlab.uky.edu	QEMU	768	48	60.1GB	251.8GB	882GB	21.7TB	764
kompute2.netlab.uky.edu	QEMU	768	48	109GB	251.8GB	176GB	21.7TB	766
kompute3.netlab.uky.edu	QEMU	768	48	117.9GB	251.8GB	176GB	21.7TB	764
kompute4.netlab.uky.edu	QEMU	512	32	69.4GB	188.7GB	176GB	21.7TB	512
kompute5.netlab.uky.edu	QEMU	128	8	18.6GB	188.9GB	176GB	21.7TB	128
kompute6.netlab.uky.edu	QEMU	128	8	18.7GB	188.9GB	4GB	21.7TB	128
kompute7.netlab.uky.edu	QEMU	512	32	71.2GB	188.7GB	4GB	21.7TB	512
kompute8.netlab.uky.edu	QEMU	512	32	71.3GB	188.7GB	4GB	21.7TB	512
kompute9.netlab.uky.edu	QEMU	512	32	71GB	188.7GB	4GB	21.7TB	512
kompute10.netlab.uky.edu	QEMU	128	8	18.8GB	188.9GB	4GB	21.7TB	128
kompute11.netlab.uky.edu	QEMU	128	8	19GB	188.9GB	4GB	21.7TB	128

Displaying 11 items

UNIVERSITY OF KENTUCKY

Hypervisor Summary

VCPU Usage

Used 4,864 of 304

Hypervisor | Compute Host

Displaying 11 items

Hostname	Type	VCPUs (used)	VCPUs (t
kompute1.netlab.uky.edu	QEMU	768	48
kompute2.netlab.uky.edu	QEMU	768	48
kompute3.netlab.uky.edu	QEMU	768	48

UNIVERSITY OF KENTUCKY

Admin / System / All Hypervisors

All Hypervisors

Hypervisor Summary

VCPU Usage

Used 829 of 312

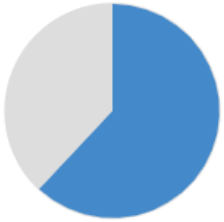

Memory Usage

Used 1.5TB of 2.4TB

Local Disk Usage

Used 4.9TB of 40.5TB

UNIVERSITY OF KENTUCKY

Security

- OpenStack servers not exposed to network
- Only the dashboard is exposed
- All VMs have dedicated security groups (ACL)
- Incoming: allow ssh, ping, remote desktop
- Outgoing: allow all

Student Pampering

- Ready to work, exclusive use, VM ☺
- “I deleted my VM” ☺
- “I deleted my project files” ☺
- Hourly (randomized) backups to persistent volume
- `perl -le 'sleep rand 3000' && rsbackup.sh`
- “I need my work from last semester” ☹

Cost Comparison

- Always on VMs 24x7
- Generous memory VMs
- OpenStack less costly than public cloud
- Public cloud estimate $300K per year
- OpenStack first year estimate $125K (equipment plus ¼ engineer)
- OpenStack next year(s) $25K (¼ engineer)

Compute Engine

145 x w10

105,850 total hours per month

VM class: regular

Instance type: n1-highcpu-4

Region: Iowa

Commitment term: 1 Year

Estimated Component Cost: $113,388.72 per 1 year

450 x Instructional labs

328,500 total hours per month

VM class: regular

Instance type: n1-highcpu-2

Region: Iowa

Commitment term: 1 Year

Estimated Component Cost: $175,948.02 per 1 year

Persistent Disk

Iowa

Storage: 30,720 GB

$14,745.60

Total Estimated Cost: $304,082.34 per 1 year

UNIVERSITY OF KENTUCKY

Compute Engine

145 x w10

105,850 total hours per month

VM class: regular

450 x Instructional labs

328,500 total hours per month

VM class: regular

304,082.34 US Dollar equals

8,923,600.35 New Taiwan Dollar

304082.34	US Dollar
8923600.35	New Taiwan Dollar

Storage: 30,720 GB

$14,745.60

Total Estimated Cost: $304,082.34 per 1 year

UNIVERSITY OF KENTUCKY

Lessons Learned

- Scalability
  - How many VMs? How much memory?
- Security
  - Only dashboard service is exposed to network
- Student Experience
  - Can modify VM configurations
  - Can “rescue” deleted VMs, deleted work
- Effective Cost Control
- Continues to makes sense to build private cloud

Thanks

Questions?