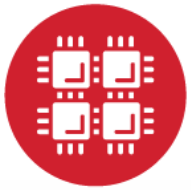

Ohio Supercomputer Center

An OH·TECH Consortium Member

Statewide Users Group

Aravind Asthagiri, Chair

Bale Theater
Aug. 8, 2013

Ohio Supercomputer Center

An OH·TECH Consortium Member

Lisa M. Hall, HC Slip Slider Professorship, Chemical and Biomolecular Engineering, The Ohio State University

Coarse-Grained Modeling of Ionomers and Salt-Doped Block Copolymers

Ohio Supercomputer Center

An OH·TECH Consortium Member

Pankaj Shah, Executive Director
Ohio Supercomputer Center & OARnet

Executive Update

2013 Focus Areas for OSC and OARnet

Ohio Supercomputer Center

- Strategic hires
- Funding opportunities and business model development
- Faculty-joint appointments!
- IntelSim - Beyond NDEMC (National Digital Engineering and Manufacturing Consortium)
- Big data management
- Research compute and storage infrastructure
- Campus outreach, regional alliances

OARnet

- Federation and Identity management
- Applications/Net+ services
- Business models and new pricing schema
- Succession planning
- Strategic hires
- Data centers
- State of Ohio transformation project/process
- Fiber IRUs

Evolving Business Models

OSC Financial Model Realignment

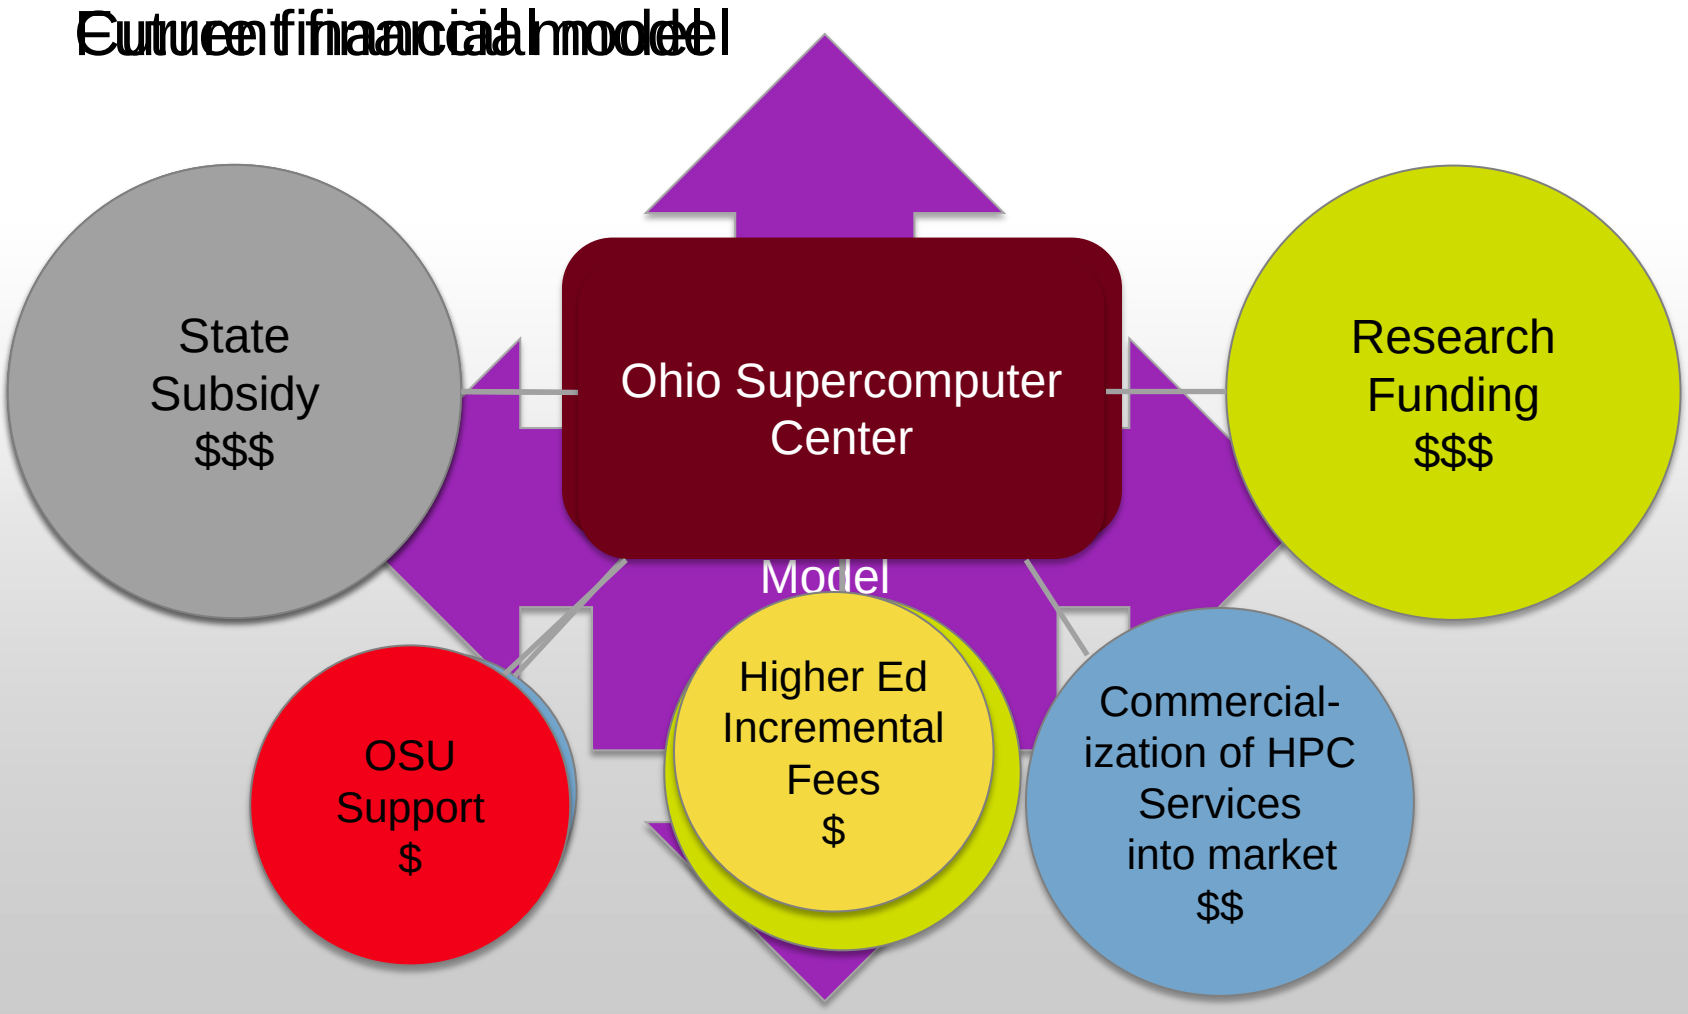

OSC Financial Update

	Budget	Forecast	Variance
Sources of Revenue			
User Fees	200,000	651,572	451,572
Grant/Contracts	1,806,651	1,984,565	177,915
Other Funding	0	0	0
Regents Funding (Higher Ed Subsidy)	3,347,412	3,347,412	0
Operating/Reserve Transfers	0	0	0
Total Revenue	**5,354,063**	**5,983,550**	**629,487**
Operating Expenses			
Personnel/Benefits	3,392,056	3,474,052	81,996
Equipment Maintenance	100,000	151,838	51,838
Equipment	16,000	32,103	16,103
Other Operating Expenses	2,200,609	2,307,325	106,716
Operating/Reserve Transfers		0	0
Total Operating Expenses	**5,708,665**	**5,965,319**	**256,653**
Net Income (Loss)	**(354,603)**	**18,231**	**372,834**

Synergy Between OSC and OARnet

Ohio's 100 Gigabit per Second Network

OARnet: 1,850 Miles of High Speed Fiber

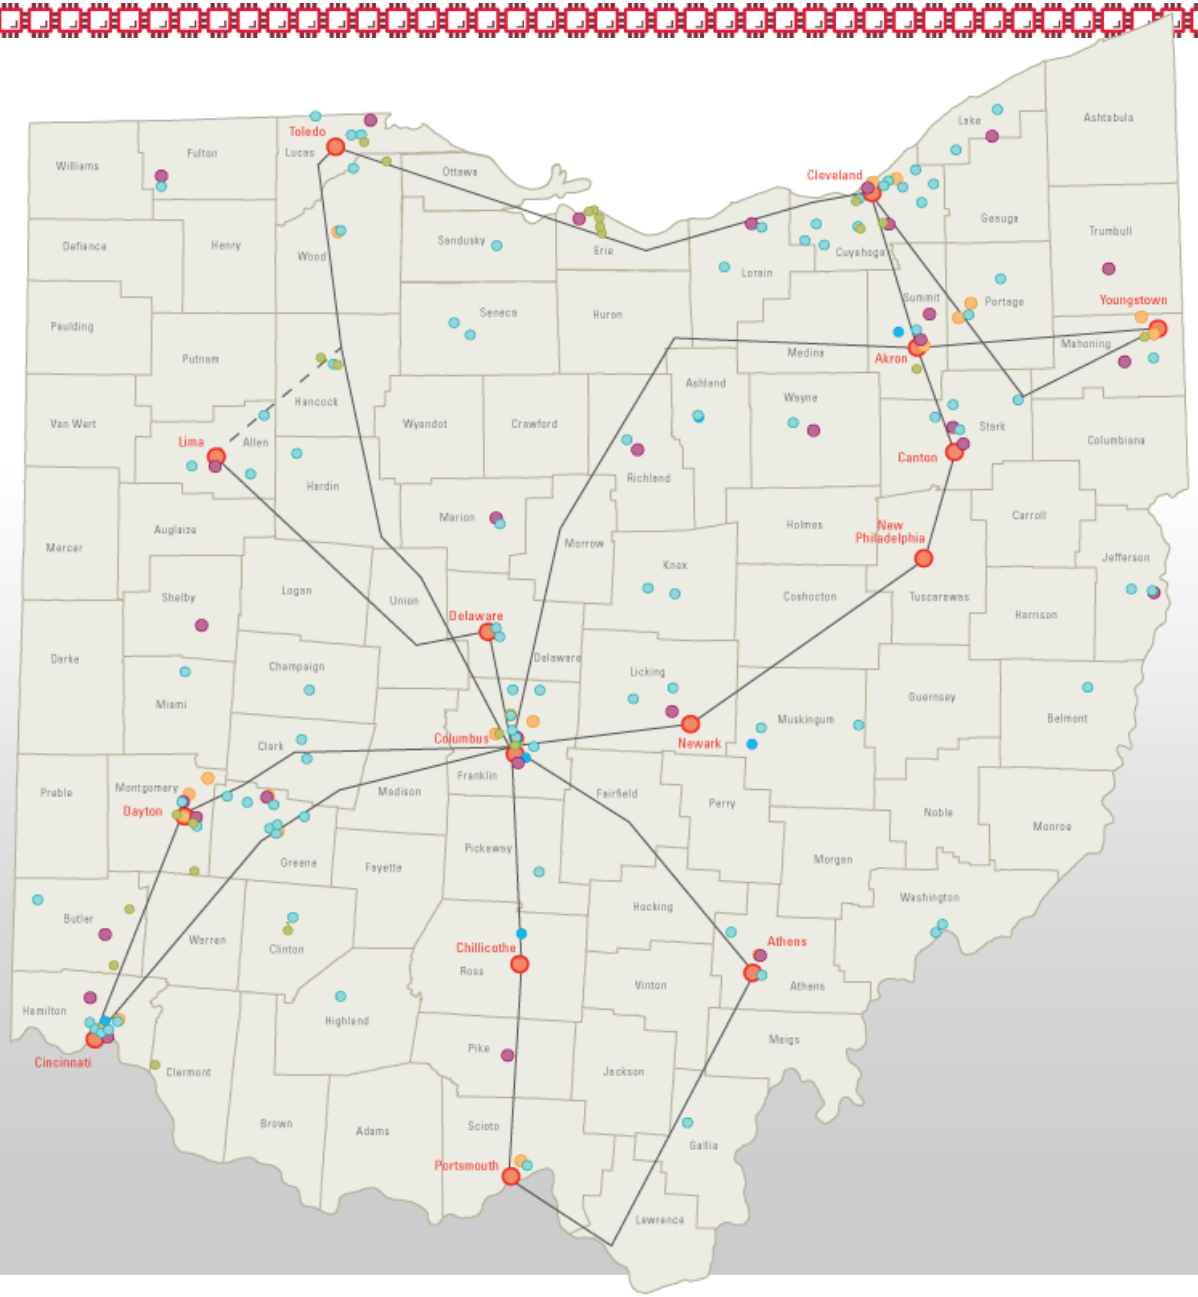

Legend

- Higher Education
- Local Governments and State Agencies
- K-12 Informational Technology Centers & Large Urban Schools
  eTech Affiliate
- Public Broadcast Stations and Radio Reading Services
  eTech Affiliate
- Health Care Facilities
- **OARnet** Pops
- **OARnet** Backbone

Internet2 Network

Wright Patterson Air Force Base Research Initiative

OARnet & DREN Peering

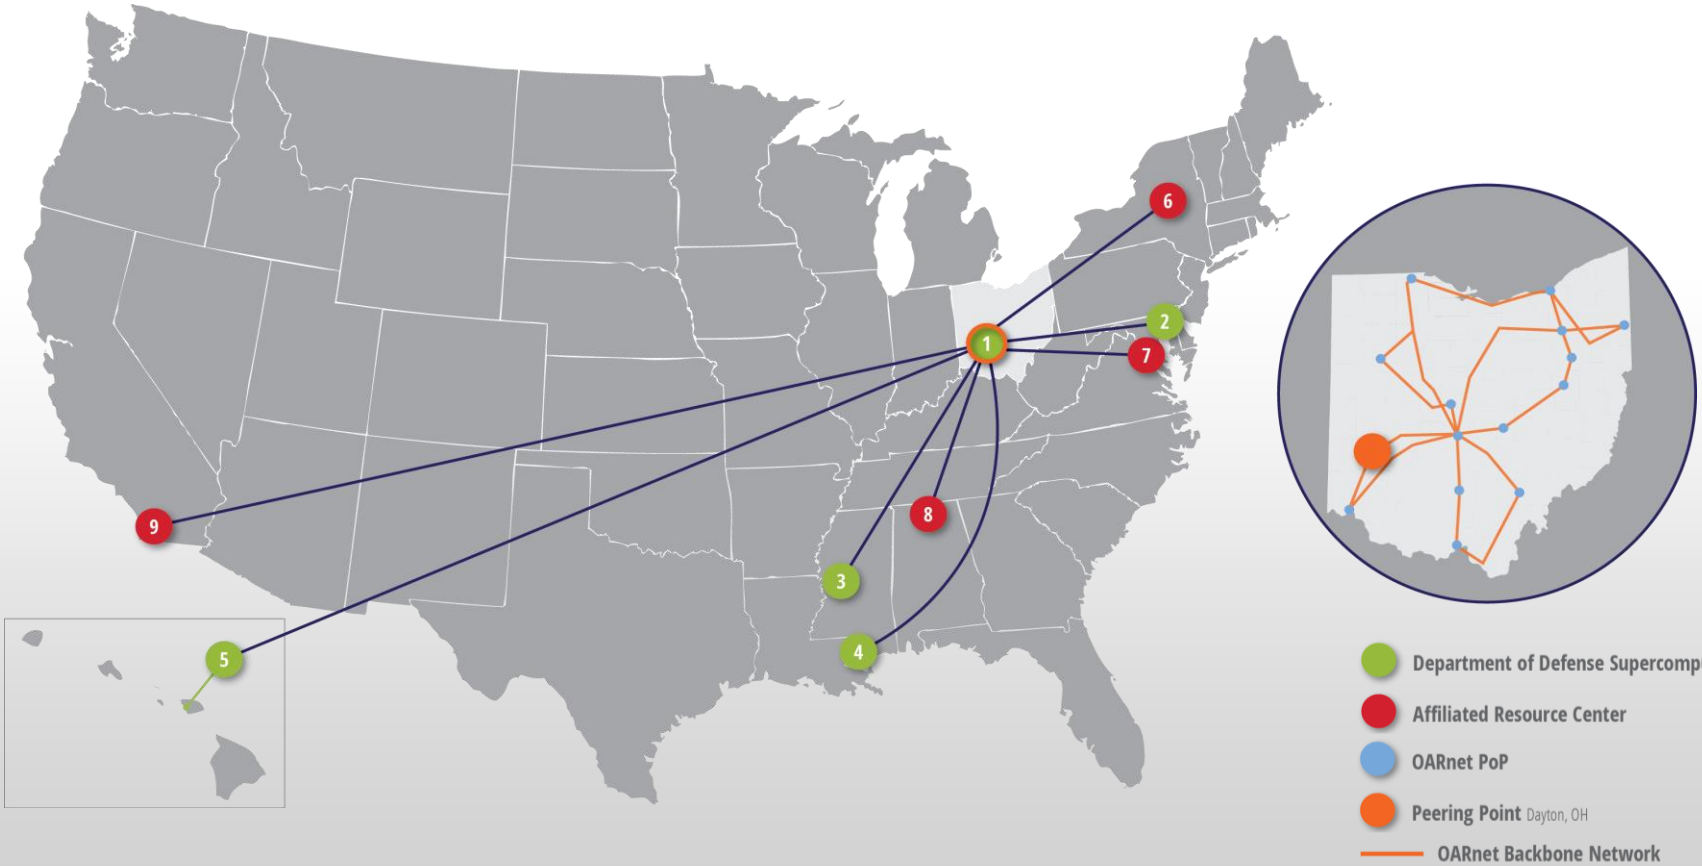

1	US Air Force Research Laboratory	Wright-Patterson AFB, OH	6	Air Force Research Laboratory/Information Directorate	Rome, NY
2	US Army Research Laboratory	Aberdeen Proving Ground, MD	7	Naval Research Laboratory	Washington, DC
3	US Army Engineer Research and Development Center	Vicksburg, MS	8	Space & Missile Defense Command	Huntsville, AL
4	Navy DoD Supercomputing Resource Center	Stennis Space Center, MS	9	Space and Naval Warfare Systems Center Pacific	San Diego, CA
5	Maui High Performance Computing Center	Kihei, Maui, HI			

OARnet & NASA Peering

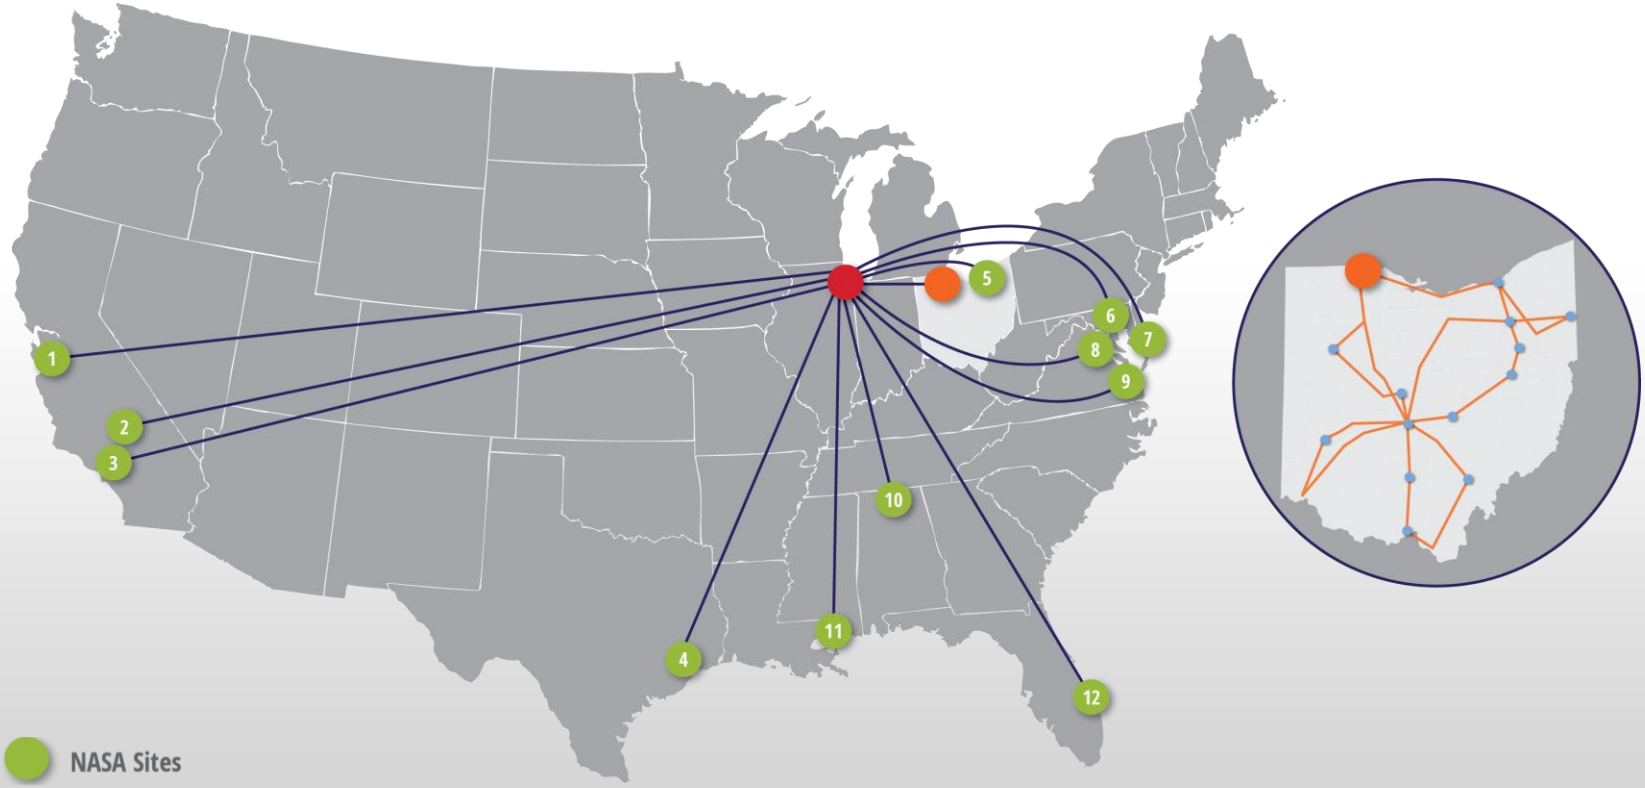

- NASA Sites
- NASA Trusted Internet Connection
- OARnet PoP
- OARnet Peering Point Toledo, Ohio
- OARnet Backbone Network

1	Ames Research Center	Moffett Field, California
2	Dryden Flight Research Center	Edwards, California
3	JET Propulsion Laboratory	Pasadena, California
4	Johnson Space Center	Houston, Texas
5	Glenn Research Center	Cleveland, Ohio
6	Goddard Space Flight Center	Greenbelt, Maryland
7	Wallops Flight Facility	Wallops Island, Virginia
8	NASA Headquarters	Washington, District of Columbia
9	Langley Research Center	Hampton, Virginia
10	Marshall Space Flight Center	Huntsville, Alabama
11	Stennis Space Center	Stennis Space Center, Mississippi
12	Kennedy Space Center	Kennedy Space Center, Florida

Unmanned Aerial Systems

Photo from http://www.aviationnews.eu/2010

Services Beyond the Network

IntraOhio and Intranet

Co-location

Emergency Web Hosting

Service Desk

Identity Management

Virtualization with VMware

Future: Needs for Growth

Factors for Continued OSC Growth

- To be competitive for National Science Foundation high performance computing grants, we need to be prepared to support the Extreme Science and Engineering Discovery Environment (XSEDE)
  - Meet XSEDE Level 1 software and services baseline, which will require 2 FTEs to comply
  - Demonstrate progress towards meeting these capabilities
  - Need a principal investigator with strong NSF connections and national HPC reputation

Condo-of-Condos Program

- Collaborative effort to integrate networking, database, and HPC related activities to improve research productivity
- 11 institutions, led by Clemson University
  - Requirements: Science DMZ, 100-Gig network, local HPC and storage
- Proposal to NSF submitted
- Would fund 1-2 FTE specialists at OSC to assist with large-scale data projects and HPC analytics

Make It In America Challenge

- Challenge will fund up to 15 projects, at $2 - $4 million per project over 3 years
- Funded by:
  - Economic Development Administration
  - National Institute for Standards & Technology Manufacturing Extension Partnership
  - Dept. of Labor's Employment and Training Administration
- OSC proposal builds on IPP grant; focuses on 65 economically disadvantaged Ohio counties
- If awarded, would leverage federal funding to improve results we expect to achieve through Third Frontier grant

National Network for Manufacturing Innovation

- National Network for Manufacturing Innovation (NNMI) RFPS
  - 2 related RFPS
  - Whitepapers due early Aug, proposals mid Oct, award mid Dec
  - $70M federal funding over 5 years, required 1:1 cost share
- Digital Manufacturing and Design Innovation (DMDI) Institute
  - Goal: create a resource to focus on complex issues in manufacturing and develop solutions
  - OSC included in U. Illinois Labs-led team, as part of NDEMC
  - OSC included in Colorado Advanced Manufacturing Alliance-led effort
- Lightweight and Modern Metals Manufacturing Innovation (LM3I)
  - Goal: advance processing and fabrication technologies for lightweight and modern metals
  - OSC included in EWI-led team, as part of OSU involvement

Capital Proposal Planning

- Major growth in compute and storage demand
  - Higher Education
    - Health and Biosciences
    - Energy
    - Food production
    - Materials
  - Manufacturing
    - Modeling and simulation
  - Technology
    - IntelSim

Capital Proposal Planning

- Multiple infrastructure projects to support larger scales
  - Infrastructure servers
  - Storage servers
  - SOCC remediation
- Goal 1: Compute cluster proposal
  - 0.5 to 1 PF
  - Deployed early 2015
- Goal 2: Storage expansion proposal
  - 4 PB spinning disk, 10+ PB tape

Detailed Proposal Items

- User-facing hardware and software
  - Compute cluster
  - Project storage upgrades
  - LTO6 tape library upgrade
  - Software capital purchases/upgrades
- Non-user facing hardware and software
  - SOCC UPS/PDU upgrades
  - Cooling infrastructure
- Infrastructure
  - Other, broader upgrades
    - Interface lab/BALE upgrades
    - Fishbowl/showroom
    - Conference room improvements
  - Staff system upgrades
  - Furniture, work environment

Workforce Education

OH·TECH | Ohio Technology Consortium
A Division of the **Ohio Board of Regents**

TABLE OF ORGANIZATION & STEM Student Employment

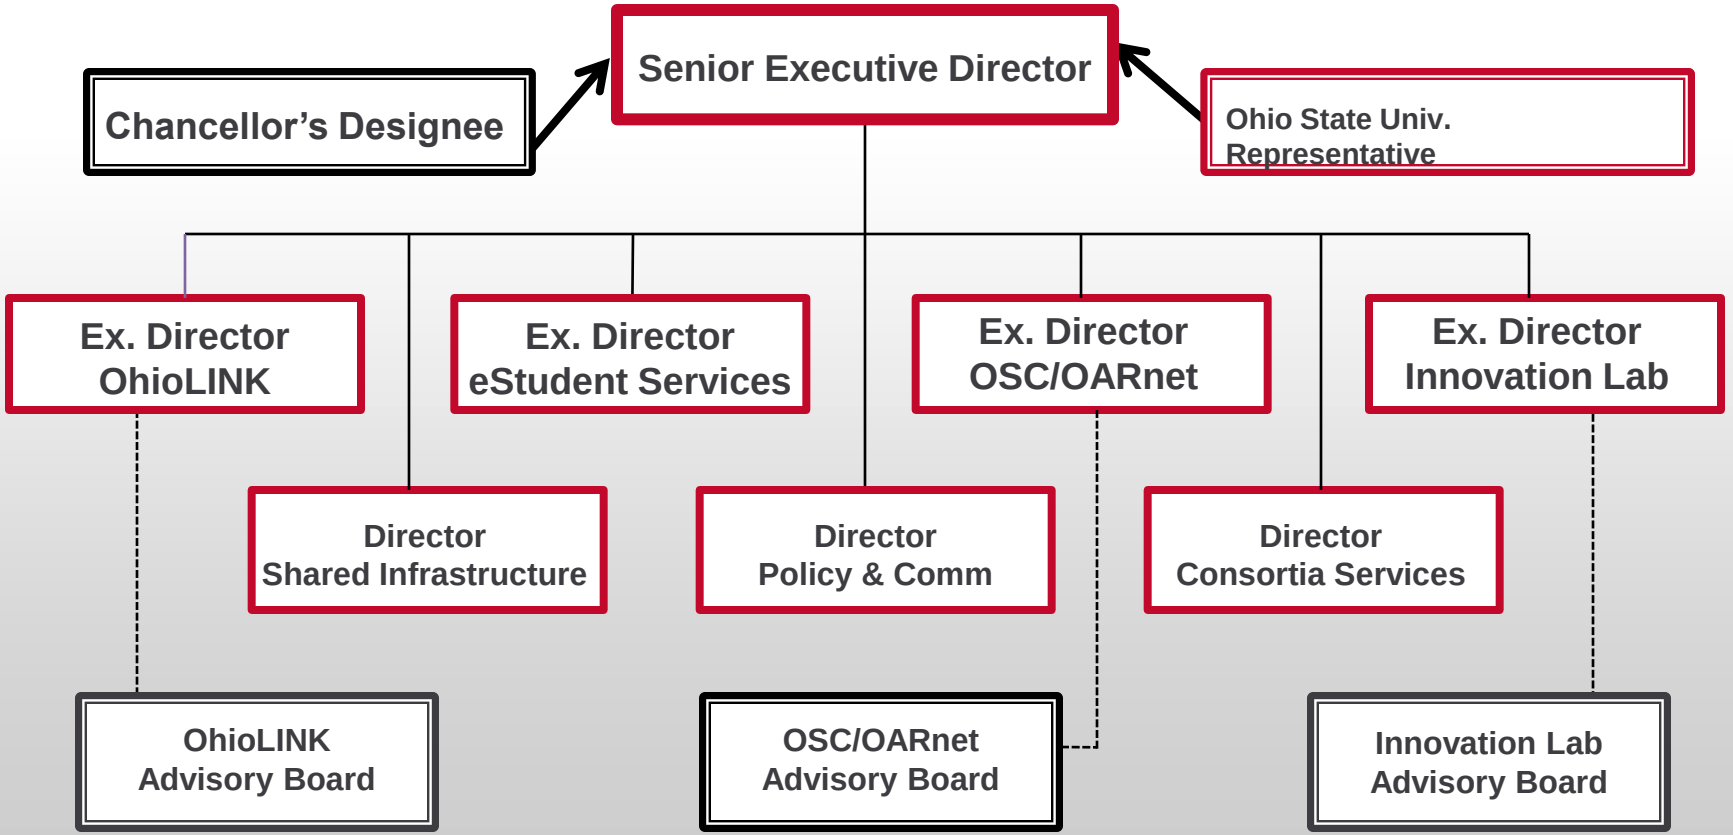

Number of Graduate/Undergraduate Students Employed at:

OSC: 15
OARnet: 9
Shared Infrastructure: 3
Policy & Comm: 3
eStudent Services: 2

KEY:
Black = OBR
Red = OSU
——— = Reporting

Ohio Supercomputer Center

Types of Student Employment Opportunities

OARnet

Engineering
- Undergraduate/graduate students
  - Scripting for router configs and servers
  - Programing
  - Configurations of servers
  - Clean up of configs
  - Building/testing new network applications
  - Develop and maintain network maps

Service Desk Agent
- Undergraduate/graduate students
  - Answer calls from clients
  - Submit tickets into ServiceNow
  - Call vendors for updates
  - Monitor network Update client incidents

Shared Infrastructure

Desktop Support
- Undergraduate students
  - Set up new systems
  - Deploy hardware deployment
  - Provide Tier I/Tier II support

eStudent Services

eTutoring Collaborative
- Undergraduate students
  - Operate eTutoring platform by participating in virtual meetings, webinars, and summer workshops
  - Assist with quality assurance reviews, analysis and discussions

Ohio Supercomputer Center

Research/Programming
- Graduate research assistants
  - Help create middleware software, web portals, etc.

Technical Support
- Undergraduate students
  - Support supercomputers at the State of Ohio Computer Center
  - Activities include: replacing parts, assisting with hardware troubleshooting, placing or replacing cable or network fiber

User Support
- Undergraduate/graduate students
  - Take first line calls from users
  - Assist in resolving problems
  - Create support tickets and escalate as needed

Interface Lab
- Undergraduate research assistants
  - Software development
  - Assists director in laboratory organization, including demonstrations, equipment specification, ordering, and integration
  - Prepares systems(hardware/software) for conferences and remote studies.
  - Currently conducting bone force analysis for Temporal Bone Project
- Graduate research assistant (PhD)
  - Developing fluid representation for Temporal Bone Project

Policy & Communication

Public Policy
- Undergraduate/graduate students
  - Assists with research and issue tracking for technology and higher education
  - Review and analyze state and federal legislation
  - Assists with maintenance and cultivation of key contacts

Creative Services & Comm
- Undergraduate students
  - Assists with design of brochures, maps, videos, photography
  - Assists with creation of web pages and upkeep
  - Write press releases, marketing material, research reports, etc.

Ohio Supercomputer Center

An OH·TECH Consortium Member

Christopher Hadad

Allocations Committee Update

Projects awarded between meetings (1)

- Standard Requests (10,000 RUs)
  - Peter C Tandy, Physics, KSU
  - Thomas Statler, Physics and Astronomy, OSU
  - Dong Qian, Mechanical, Industrial, and Nuclear Engineering, UCN
  - Ksenija Glusac, Chemistry, BGS
  - Joel Johnson, Electrical and Computer Engineering, OSU
- Classroom Requests (5,000 RUs)
  - Walter Schilling (with David Hudak), Electrical Engineering and Computer Science, UTL
- Total of 55,000 RUs

Projects awarded between meetings (2a)

- Startup Requests (5,000 RUs each) – total: 120,000 RUs
  1. Thomas Santner, Statistics, OSU
  2. David Cole, Earth Sciences, OSU
  3. Janet Steffensmeier, Political Science, OSU
  4. Ian Howat, Earth Sciences, OSU
  5. Brian Baugh, Finance, OSU
  6. Daniel Turner, Chemistry, UDA
     - new PI, but not new to OSC (PhD degree from OSU)
  7. Soheil Soghrati, Mechanical and Aero Engineering, OSU
  8. Utku Solpuker, Earth Sciences, OSU
     - Former postdoctoral researcher under Dr. Ibaraki, now a Research Associate I
  9. Timothy Hewett, Physiology and Cell Biology, OSU
  10. Pavel Anzenbacher, Jr, Chemistry, BGS
  11. Jay Hollick, Molecular Genetics, OSU
  12. Terrence Walker, Physics, OSU

Projects awarded between meetings (2b)

- Startup Requests (5,000 RUs each) – total: 120,000 Rus

13. John Saldanha, Marketing and Logistics, Fisher College of Business, OSU
14. Toshiaki Shinoda, Geography, OSU
15. Ryan Yoder, Chemistry and Biochemistry, OSU-Marion
    - OSU PhD (August 2013), recently appointed as Assistant Professor at OSU-Marion
16. Jeffrey Campbell, Arts and Sciences, OSU
17. Michael Strickland, Physics, KSU
    - Former post-doctoral researcher with Dr. Furnstahl (OSU Physics), recently started at Kent State Physics
18. Jason Halloran, Biomedical Engineering, CCF, start-up, 5,000
19. Timothy Huerta, Family Medicine, OSU, start-up, 5,000
20. Hannah Shafaat, Chemistry and Biochemistry, OSU, start-up, 5,000
21. Cameron Thraen, AEDE, OSU, start-up, 5,000
22. Daniel Kinnamon, Internal Medicine, OSU, start-up, 5,000
23. Joseph Houpt, Psychology, WSU, start-up, 5,000
24. Annika Peter, Physics, OSU, start-up, 5,000

Major (6) Requests – 30,000 RUs each

- Major Requests
  1. KT Arasu, Mathematics and Statistics, WSU, major, 30,000
  2. Shili Lin, Statistics, OSU, major, 30,000
  3. Sichun Yang, Center for Proteomics and Bioinformatics, CWR, major, 30,000
  4. David Ball, Chemistry, CLS, major, 30,000
  5. Anne B McCoy, Chemistry, OSU, major, 30,000
  6. Shaurya Prakash, Mechanical and Aero Engineering, OSU, major, 30,000

Discovery (18) Requests

1. Rafael Bruschweiler, Chemistry and Biochemistry, OSU, subsidized, 250,000
2. Christopher Hadad, Chemistry, OSU, Discovery, 500,000
3. John Herbert, Chemistry, OSU, Discovery, 375,000
4. Hendrik Heinz, Polymer Engineering, UAK, Discovery, 100,000
5. Daniel Lacks, Chemical Engineering, CWR, Discovery, 100,000
6. Pearlly Yan, Comprehensive Cancer Center, OSU, Discovery, 50,000
7. Alexei Federov, Bioinformatics Lab, UTL, Discovery, 100,000
8. Chenglong Li, Medicinal Chemistry, OSU, Discovery, 250,000
9. Ian Howat, Earth Sciences, OSU, Discovery, 50,000
10. Shaojun Wang, Computer Science and Engineering, WSU, Discovery, 150,000
11. Alexandre Sousa, Physics, UCN. Discovery, 60,000
12. Shaojun Wang, Computer Science and Engineering, WSU, Discovery, 150,000
13. Purnima Kumar, Dentistry, OSU, Discovery, 100,000
14. Peter White, Pediatrics, Nationwide Children's Hospital, Discovery, 48,000
15. Ruoming Jin, Computer Science, KSU, Discovery, 150,000
16. Sandip Mazumder, Mechanical and Aero Engineering, OSU, Discovery, 200,000
17. Thomas J Humanic, Physics, OSU, Discovery, 100,000
18. Shili Lin, Statistics, OSU, Discovery, 60,000

Total Awards for Two Allocations Meetings

- Startup (5,000 RUs): 24 requests, **120,000**
- Classroom (5,000 RUs): 1 request, **5,000**
- Standard (10,000 RUs): 5 requests, **50,000**
- Major (30,000 RUs): 5 requests, **150,000**
- Discovery requests: **2,693,000**

- Total of **3,018,000** RUs
- Approved 8TB and 25 TB storage requests

- Need OSC policy for Storage and Investment by Users

Ohio Supercomputer Center

An OH·TECH Consortium Member

Doug Johnson

Supercomputing Update

Outline

- System utilization
- Software utilization
- Storage utilization
- Operations activities in February / March
- Upcoming

System utilization summary, April - July

System	Job count	CPU hours	Utilization	Users	Groups
Glenn	580,571	7,885,884	51%	424	197
Oakley	513,354	17,839,053	73%	432	204

Glenn Lifetime Utilization

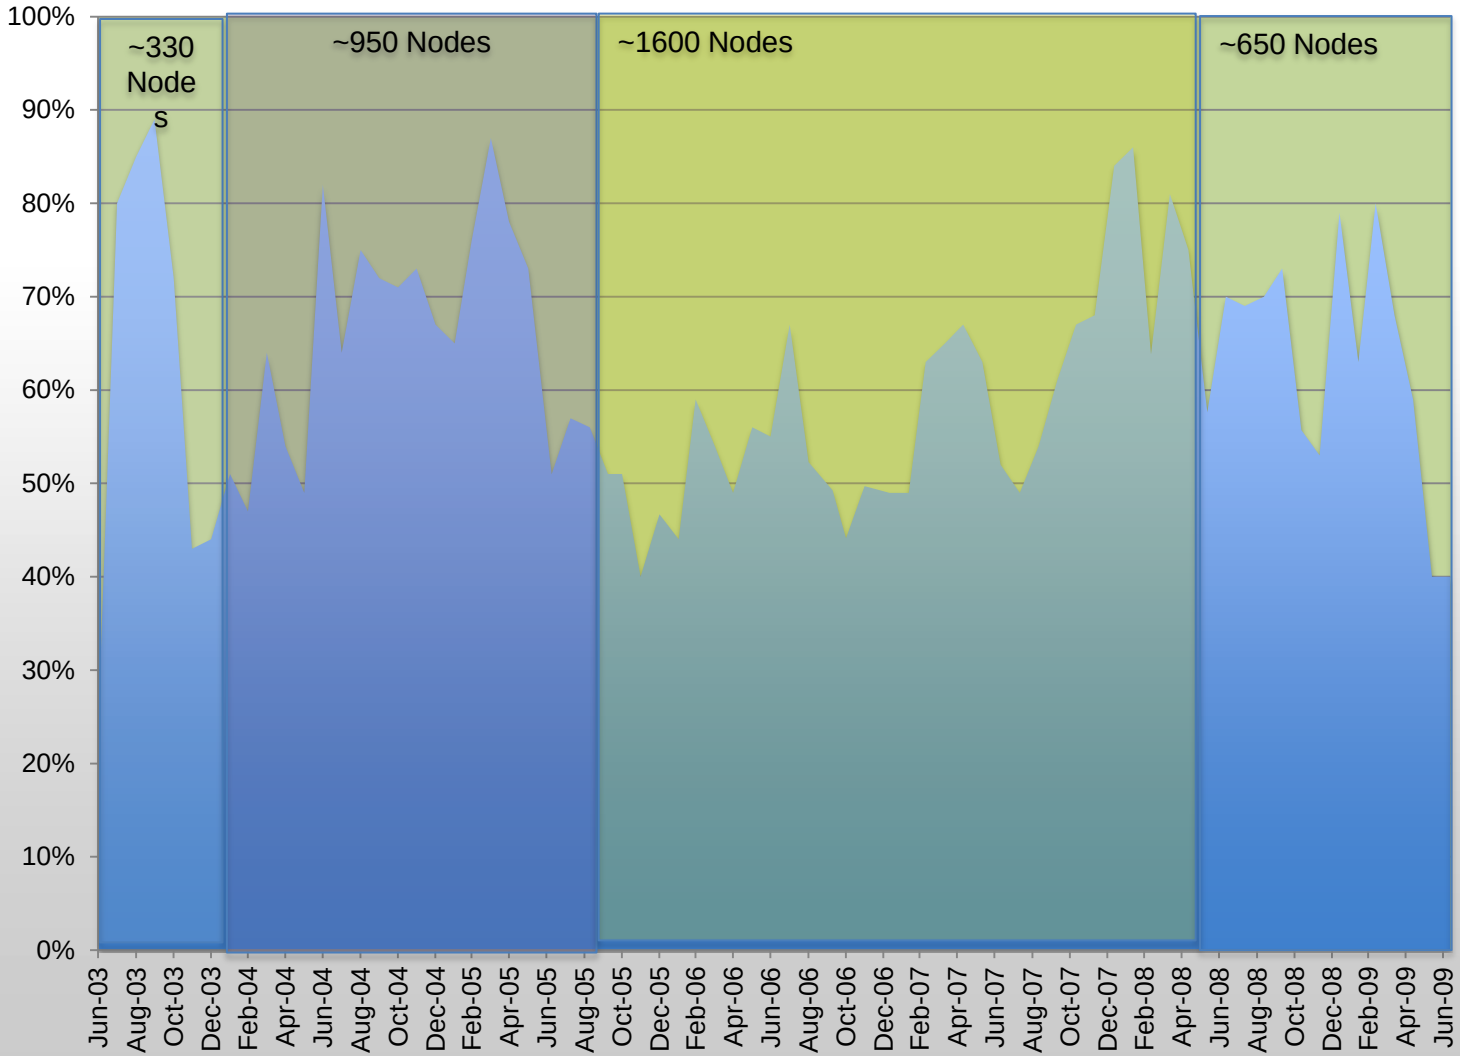

Oakley Lifetime Utilization

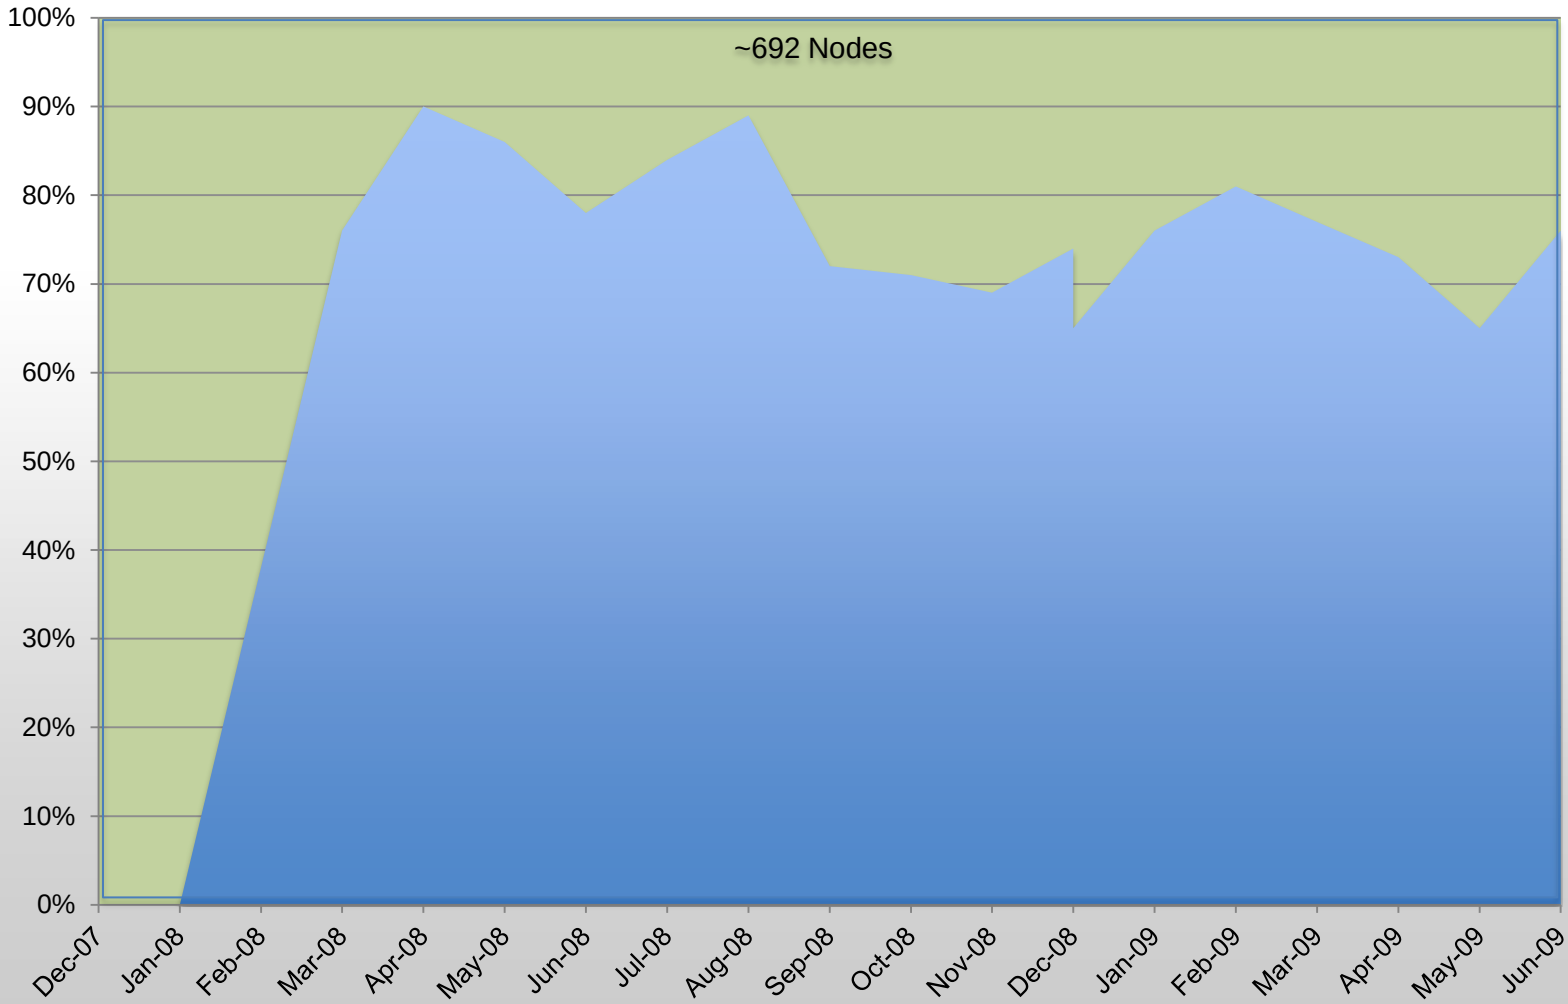

Glenn Utilization By Institution (April - July)

Institution	Job count	CPU hours	Users	Groups
Ohio State University	498,127	4,440,435	219	92
Bowling Green	6,534	980,703	25	8
University of Cincinnati	13,677	714,479	53	16
Wright State University	4,897	482,647	13	6
University of Akron	898	392,295	4	2
Case Western Reserve	11,443	318,606	6	3
Ohio University	12,380	134,443	12	10
University of Toledo	16,933	83,315	16	10
Cleveland State University	804	48,624	5	3
Youngstown State	2,076	46,681	3	3
Kent State University	184	33,280	5	5
Wittenberg	7,379	19,678	1	1
Commercial	91	16,714	6	5
Nationwide Children's Hospital	828	13,030	4	3
OSC	1,786	10,550	27	13

Oakley Utilization By Institution (April - July)

Institution	Job Count	CPU hours	Users	Groups
Ohio State University	446,020	11,165,440	254	117
Commercial	3,807	2,005,577	24	11
University of Toledo	9,758	1,130,148	14	6
Case Western Reserve	4,507	406,924	10	4
Kent State University	26,377	350,328	16	8
Wright State University	5,349	347,944	5	3
University of Cincinnati	8,027	242,453	13	8
Nationwide Children's Hospital	2,249	201,079	4	2
Nimbis Services/NDEMC	763	181,109	14	8
OSC	3,385	137,353	50	17
University of Akron	874	127,922	4	3
Cleveland State	895	116,684	7	3
Ohio University	441	32,206	2	2
Central State University	262	32,097	2	2
Bowling Green	108	29,424	5	4

Software Utilization, Glenn

Package	jobcount	cpuhours	users	groups
vasp	18,044	1,381,558	19	10
molcas	4,314	884,668	11	3
a_out	1,684	781,152	25	18
matlab	19,538	706,729	17	10
amber	24,519	669,515	13	7
gaussian	11,835	629,970	98	40
gromacs	21,522	608,040	12	8
lammps	5,355	319,716	6	4
AliEn	323,382	309,847	1	1
fluent	2,912	218,995	40	21
ansys	1,059	191,864	43	24
namd	526	180,678	4	3
charmm	130	91,046	1	1
qchem	1,160	88,907	7	2
ncbi	3,087	72,147	21	4
lsdyna	525	59,813	8	7
abaqus	1,210	56,417	20	17
aces2	108	41,747	1	1
turbomole	304	40,812	7	7
ccsm	21	35,976	1	1

Software Utilization, Oakley

package	jobcount	cpuhours	users	groups
vasp	28,665	3,276,065	28	8
OpenFOAM	1,700	1,783,728	25	15
qchem	111,846	1,728,091	18	5
matlab	16,004	1,139,578	38	27
lammps	736	1,070,431	5	3
gaussian	12,088	960,414	59	22
wrf	14,729	650,863	6	5
amber	128,632	527,950	15	10
namd	762	401,222	2	2
cpmd	715	391,903	1	1
gromacs	4,357	273,926	21	10
cuda	3,120	167,998	29	23
fluent	434	145,476	18	12
comsol	145	141,788	5	4
autodock	1,565	94,412	7	3
a_out	1,011	78,315	35	31
starccm	387	48,517	3	2
lmf	4,700	45,741	6	3
dlpoly	136	33,966	2	1
abaqus	881	23,328	15	14

Home directory usage trends (GB)

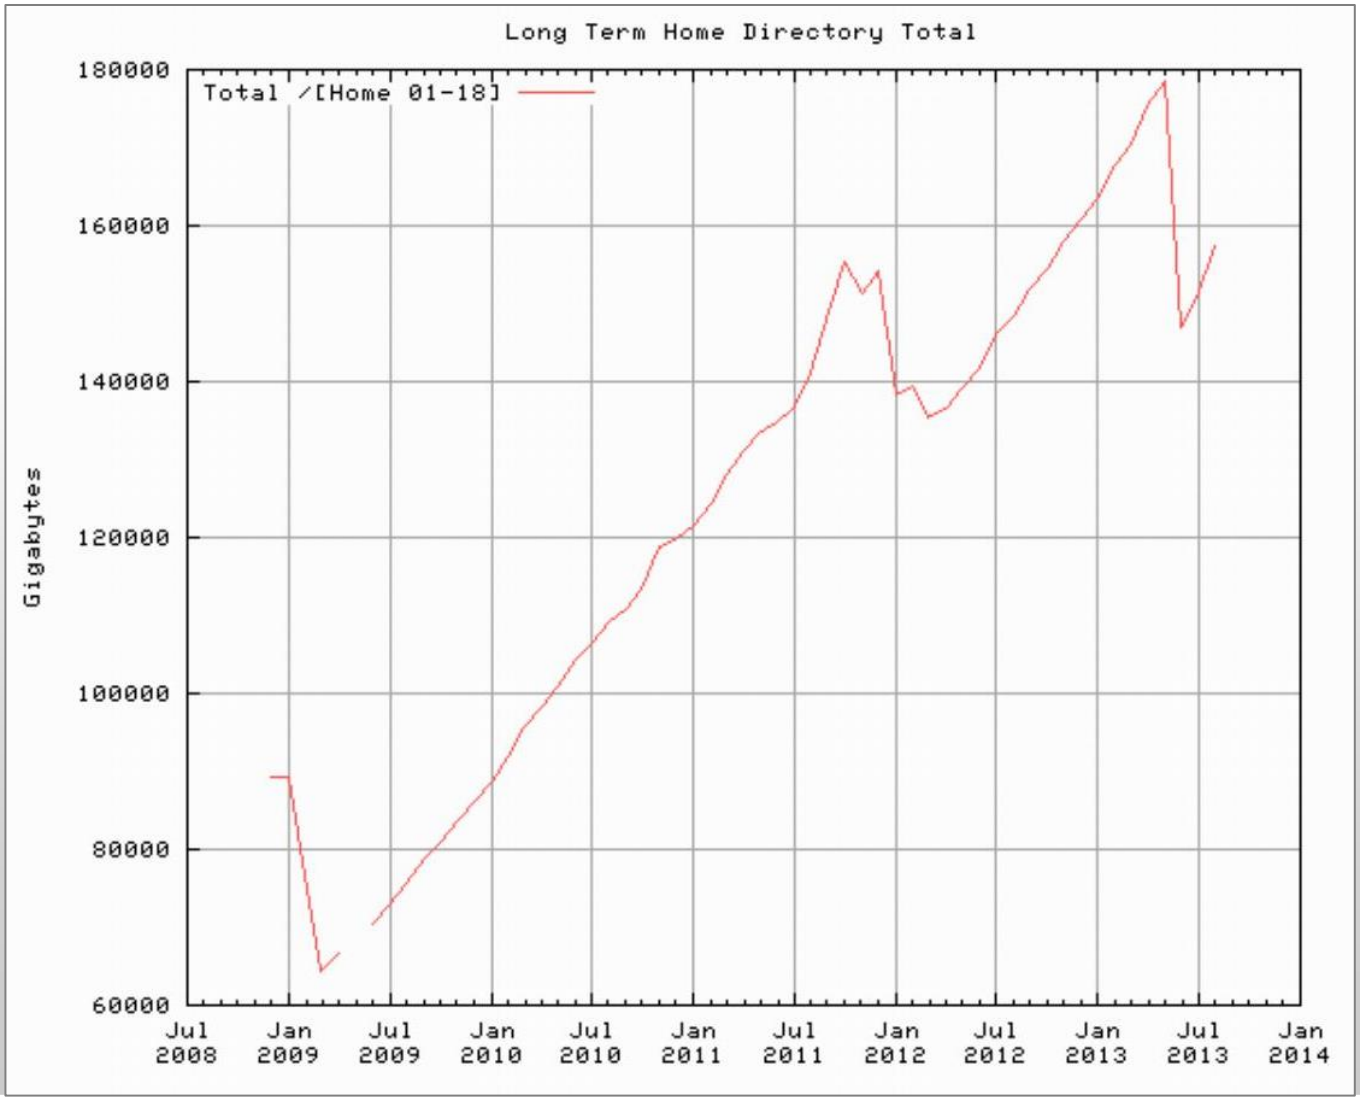

Home, Project, and GPFS Storage Utilization Summary

File systems	Total usable (TB)	Used (TB)	Free (TB)	Files (millions)
Home	290	165	125	124
Project	630	450	180	82
GPFS	400	123	278	29

Tape Library Utilization

Tape type	Total tapes	Used tapes‡	Free tapes	Free capacity*
LTO-3	634	571	63	37TB
LTO-4	2100	1591	509	595TB
LTO-5	10	0	10	22TB

Tape library slots	Used slots	Free slots	TSM reported used capacity	Free capacity†
3238	2734	~500	1.4PB	1.7PB

‡Fragmented tape use increases used tapes count, planning merge.

*Assuming 1.5:1 compression ratio.

†Assuming LTO-5 tape media is added, single copy, and all tapes used, 1.5:1 compression ratio, and a tape merge to reclaim tapes.

Operations Activities in April - July

- Center downtime June 4
  - Upgraded GPFS software to version 3.5.0-6
  - Changed older group names from the form G-[GID] where GID is the Unix group ID to the project name (PAS1234 for example)
  - Replaced faulty controller in the storage array used for the project file systems
  - Software security updates, including the Linux kernel on all Oakley nodes
  - PCI slot change for network card in home directory servers
  - Add redundant disk to web service node used for OnDemand services

Operations Activities in April - July (cont.)

- IBM LTFS EE beta evaluation
  - Two LTO-5 tape drives and 10 tape cartridges installed
  - Open format for import and export of data via tape
  - Alternative to IBM Tivoli TSM for Hierarchical Storage Management (HSM) services for GPFS
  - Evaluation ongoing
- Ruby cluster update
  - Cloned root file system from Oakley cluster
    - Important for supporting hardware condo model
  - Installed new Intel Xeon Phi cards
  - Efforts to make this a multi-user system still ongoing

Operations Activities in April - July (cont.)

- GPFS hardware upgrade update
  - New GPFS servers physically installed
    - OS installation is separate from cluster environments
  - Metadata storage hardware raid array installed
  - Additional Infiniband network switches installed
  - Testing of hardware ongoing
- Changes to how backups of GPFS file system are performed
  - Preparation for migration, and larger file system
- Scheduler reservation improvements for GPU nodes

Upcoming

- Power changes at the SOCC
  - Downtime on September 29 or 30 for cutover
  - Potential for additional capacity
- Beta test of IBM LTFS EE software, and LTO5 donated tape drives (ongoing)
- OSC software infrastructure improvements
  - SSO and Grid services
- GPFS server replacement
  - Migration of project file systems to GPFS

Upcoming (cont.)

- Merge project and current GPFS storage using existing storage arrays
  - Approximately 1.1PB file system (~500TB free if we migrated today.)
- 8 GPFS Network Shared Disk (NSD) and NFS servers
  - Intel Xeon Processor E5-2620 processors
  - 128GB memory
  - 4 QDR Infiniband ports
  - 2 10GbE ports
  - 2 8Gb Fibre-Channel ports (NSD nodes)
  - Replaces 26 servers, but increased performance
- New storage array for metadata
  - Support for several billion files

Upcoming (cont.)

- HIPAA data support at OSC
- Deploy ‘Ruby’ research cluster for MIC work
  - 8 HP SL250 nodes
    - Dual Intel Xeon E5 2670 (Sandy Bridge) 16 cores total
    - 128 GB RAM
    - 1 TB HDD
    - 200 GB SSD node
    - Intel Xeon Phi 5110p card
  - Two NVIDIA K20X Kepler GPUs
  - FDR Infiniband (56 Gbps), QDR cards in nodes
  - Clone and diverge Oakley root file system (test of configuration management software)

Upcoming (cont.)

- Home directory file system and server replacement
  - Determine criteria (performance, capacity, etc.)
  - RFP in late 2013
  - Migrate in early 2014.
- Open Science Grid (OSG) services
- Cost recovery calculations
  - Cost recovery will be recalculated yearly
- Condo model
  - Allow OSC academic users to expand GPFS storage, compute resources

Ohio Supercomputer Center

An OH·TECH Consortium Member

Alan Chalker, Ph.D.

Director of Client Engineering

IntelSim

Ongoing Efforts to Support Industry

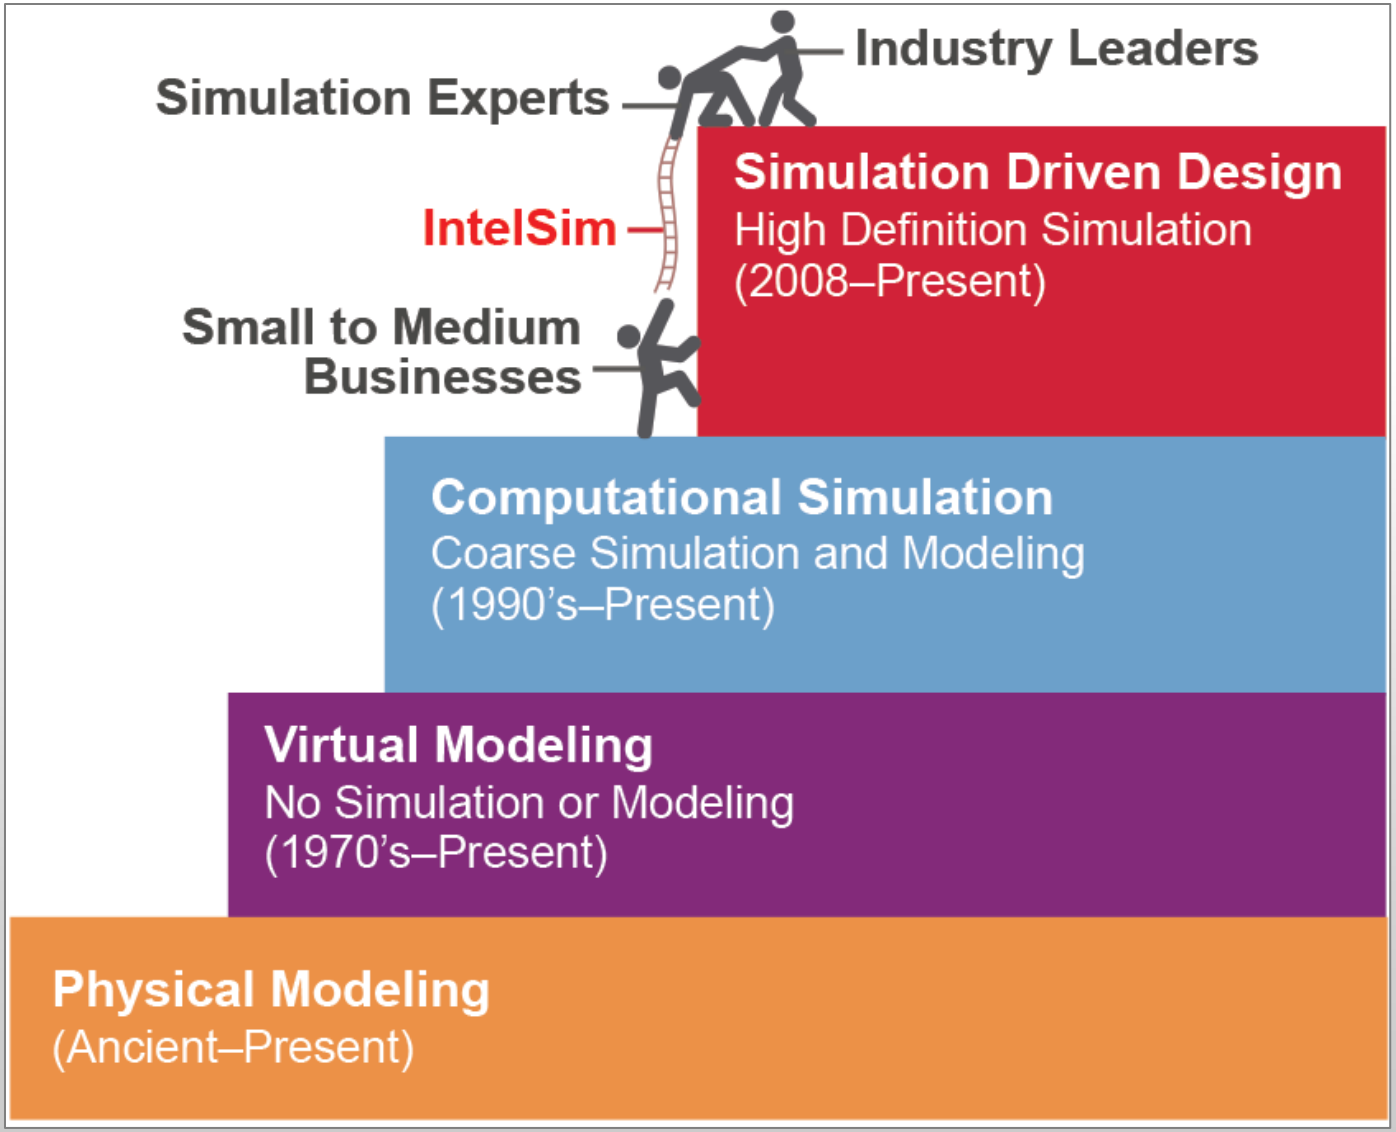

NDEMC
Leading 21st Century Manufacturing

OSC NDEMC projects include:

- **Greenlight Optics**
  Loveland, Ohio
- **Plastipak**
  Medina &
  Columbus, Ohio
- **Jeco Plastics**
  Plainfield, Indiana
- **Replex Plastics**
  Mount Vernon, Ohio
- **Applied Sciences**
  Cedarville, Ohio
- *Eastlake, Ohio*
- **Technology Management Inc.**
  Highland Heights, Ohio
- **Engendren**
  Kenosha, Wisconsin

INTELSIM

an i

IntelSim (Intelligent Simulation Platform)

Who

- Ohio Supercomputer Center, TotalSim, AltaSim, Kinetic Vision, Nimbis Services, P&G, and Intel

What

- Will develop a platform to provide cloud-based manufacturing simulation apps, sold through an e-commerce marketplace

Why

- Untapped market that will reduce barriers to simulation-driven design for smaller manufacturers
- Estimate apps will enable businesses to:
  - Reduce time to market by 50%
  - Decrease product development costs by 30%

Example App: WeldPredictor

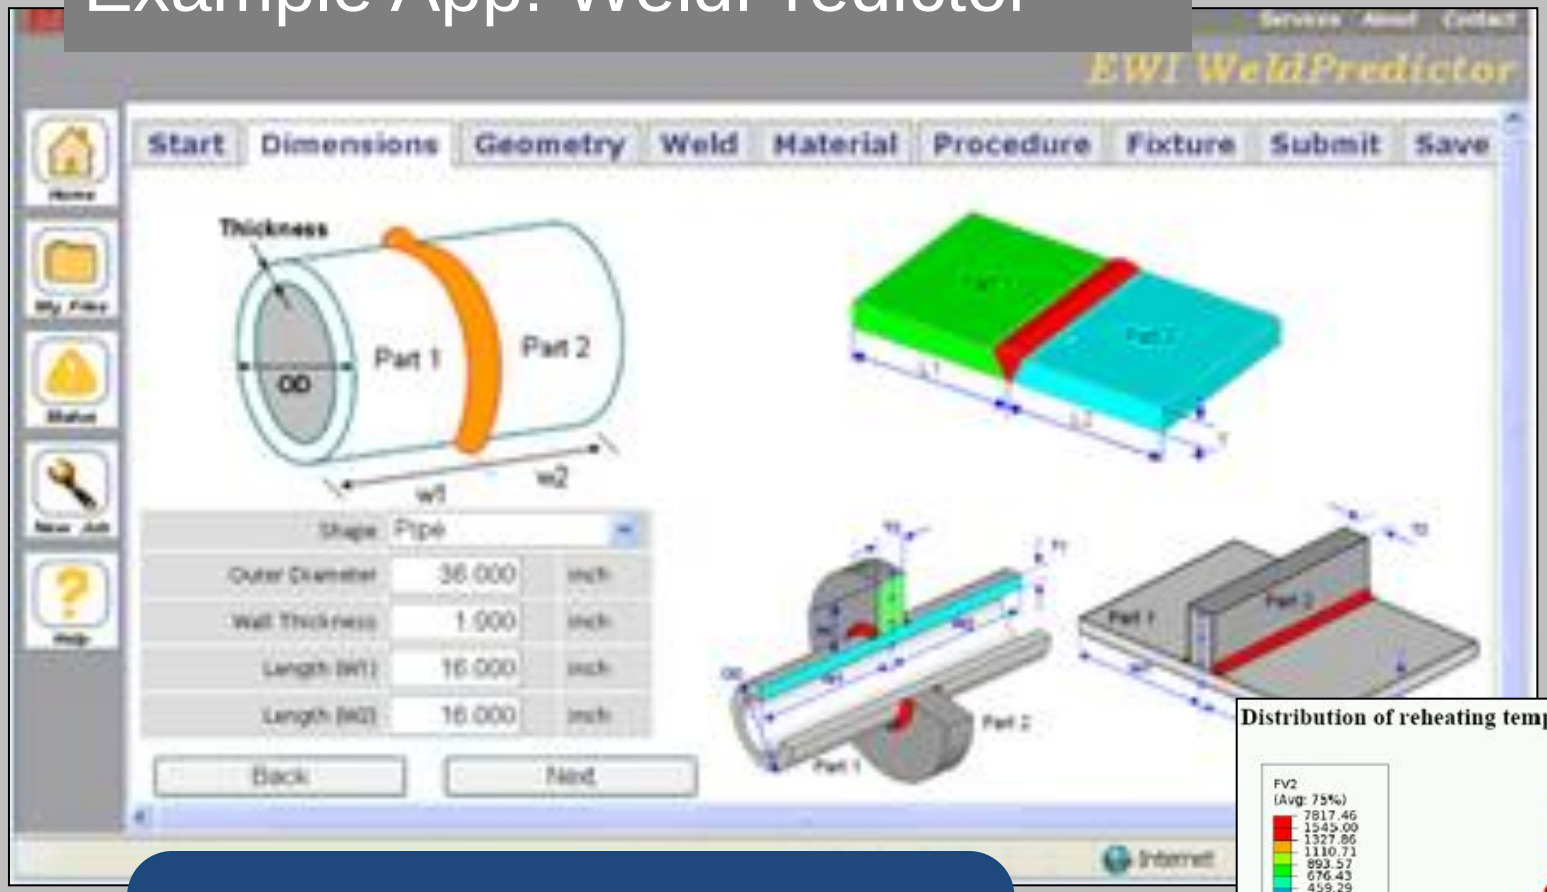

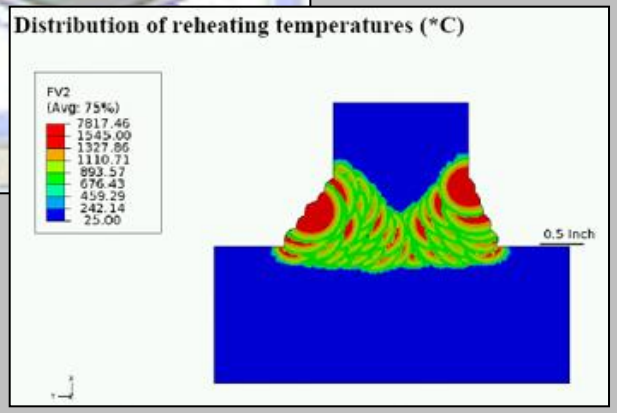

WeldPredictor: Award-winning Portal

- **2011 IDC Innovation in HPC Award**
- **2009 International Institute of Welding Sossenheimer Award**

Example App: Heat Sink Predictor

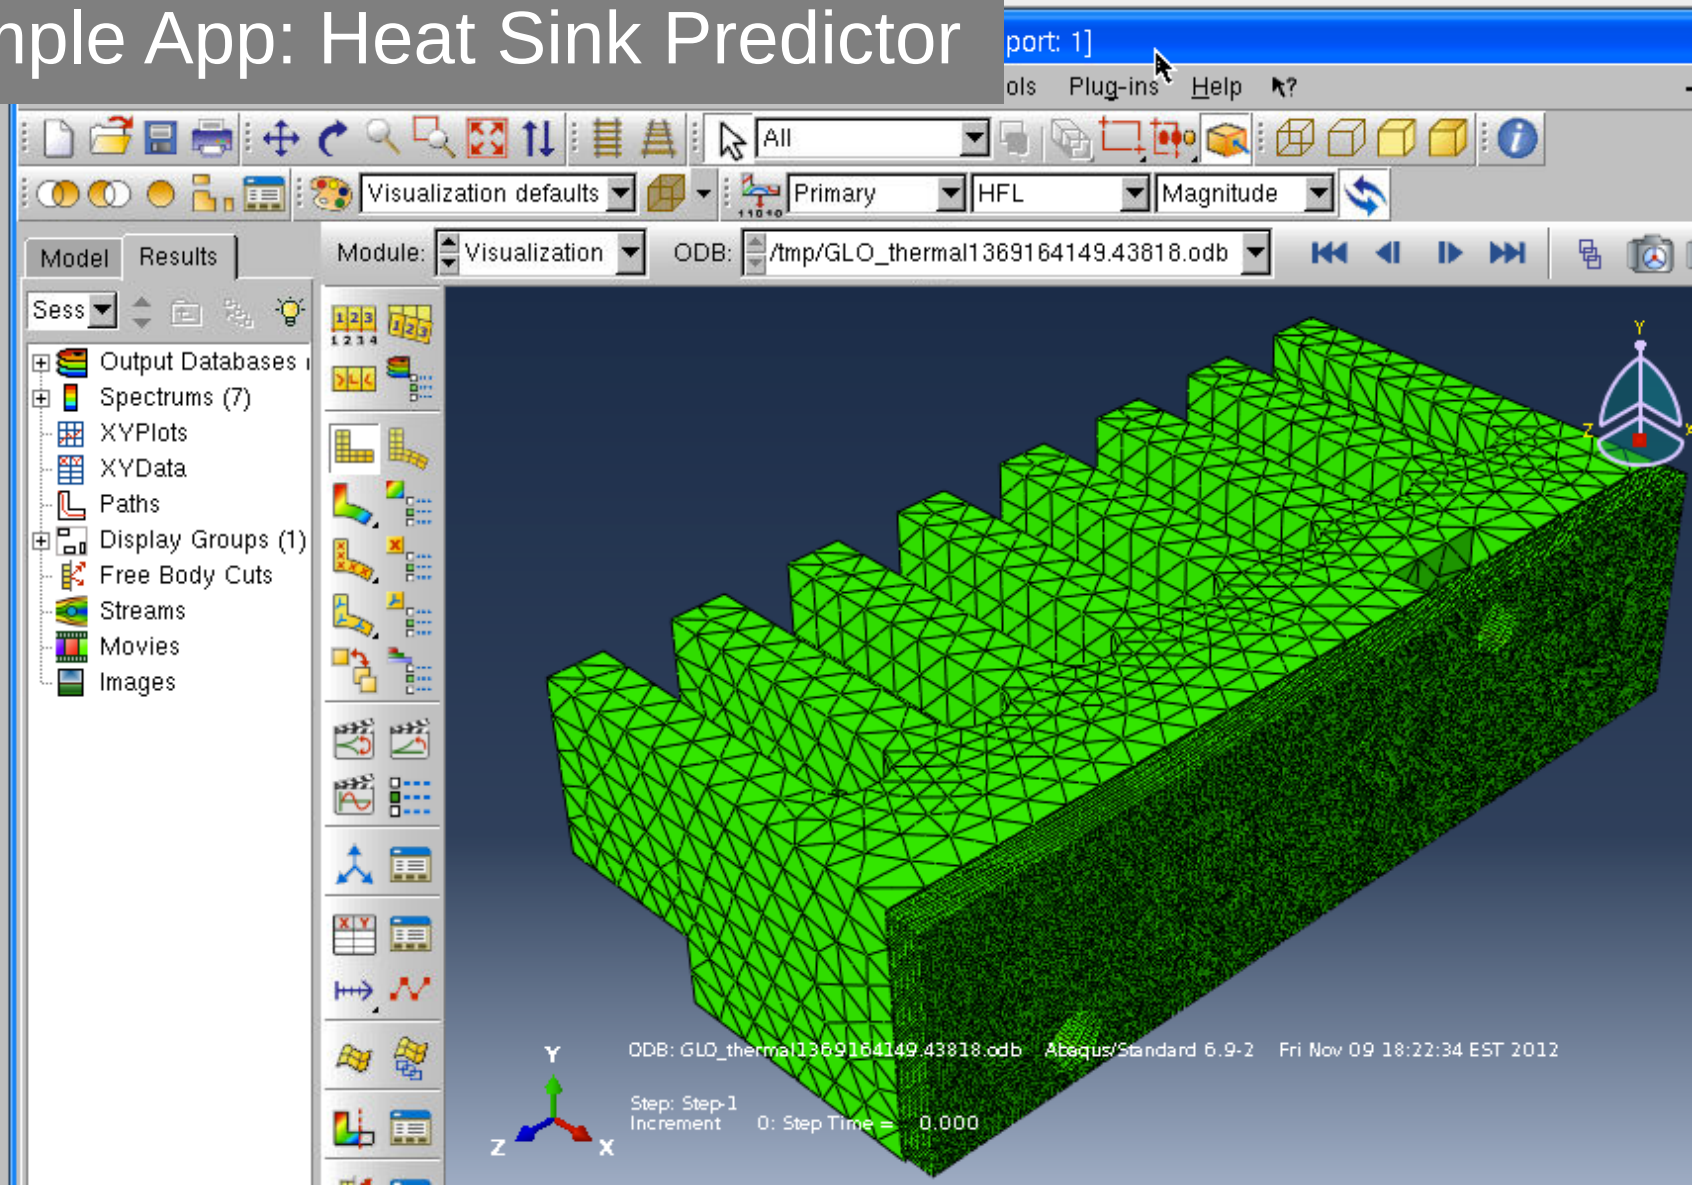

Example App: Truck Add-On Predictor

TRUCK ADD-ON PREDICTOR

My Trucks

tauW_rightSide_001.png

Run 6

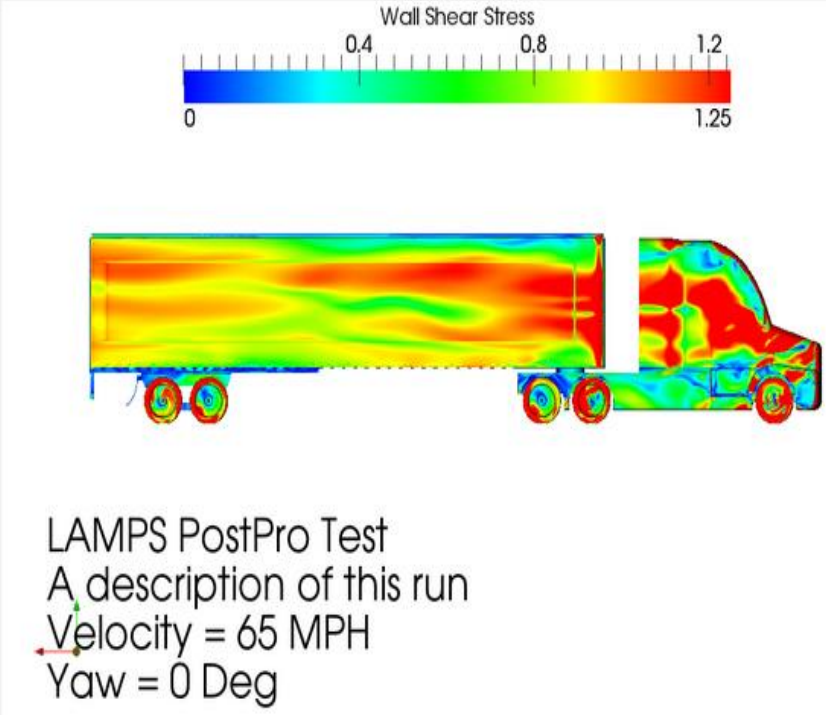

LAMPS PostPro Test
A description of this run
Velocity = 65 MPH
Yaw = 0 Deg

Run 5

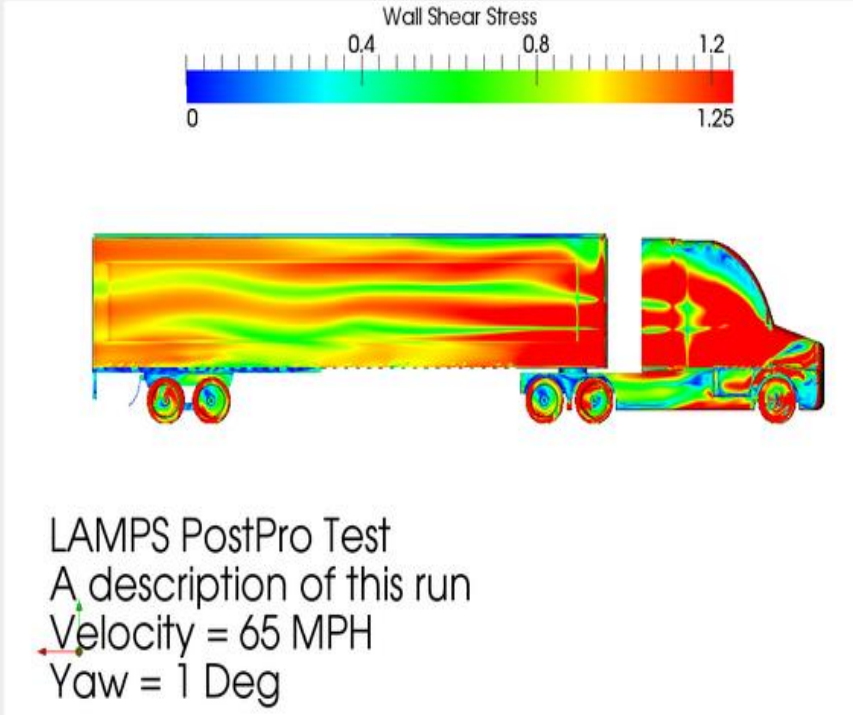

LAMPS PostPro Test
A description of this run
Velocity = 65 MPH
Yaw = 1 Deg

IntelSim Deliverables

- **Commercially ready platform** consisting of:
  - Web-based App Store
  - Supporting infrastructure
  - App development tools
- **Six initial apps** of broad marketability, such as:
  - Virtual wind tunnel
  - Crush test
- **Piloting and marketing** to Ohio-based small- and medium-sized manufacturers
- **Training programs** for employees to leverage the technology
  - Sinclair Community College
  - Lorain County Community College

Costs and Economic Benefits to Ohio

Costs

- $3 million from the Ohio Third Frontier Commission
- $3.4 million in cost share from our partners

Benefits

- **Revenue** for Nimbis, app developers and OSC
- **Additional investment** of $6 million (expected) to create new apps in four to six years
- **Low-cost access** to simulation-driven design will benefit Ohio manufacturers
- **Enhanced employment opportunities,** via training to Ohio workers.

Ohio Supercomputer Center

An OH·TECH Consortium Member

Dave Hudak, Ph.D.

Senior Research Scientist

Research Update

Agenda

- Program opportunities
  - Proposal status
  - National collaborations
- Research updates
  - XSEDE13 conference
- OnDemand update

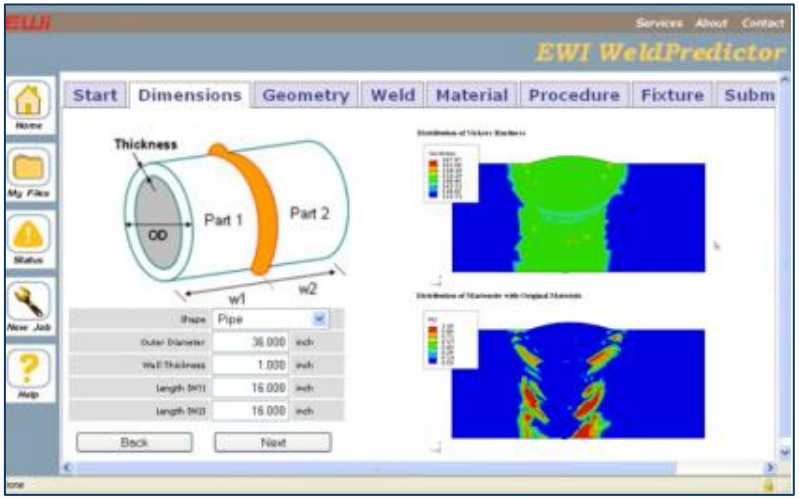

OSC and EWI partnered to create an HPC portal for welding solutions.

Pending OSC Proposals

Funding Opportunity	Proposal Submission Date
National Science Foundation – unsolicited - under Clemson University **"The Condo of Condos"**	08/31/17
São Paulo Research Foundation (FAPESP) Brazil and Ohio State University Partnership 2013 Call **"Building a Digital Bridge for Research Collaboration between Ohio State University and São Paulo Campuses"**	05/20/13
National Science Foundation – NSF 13-529 under RENCI EarthCube: Developing a Community-Driven Data and Knowledge Environment for the Geosciences (Building Blocks) **"EarthCube Building Blocks: An Interoperability Infrastructure for Geoscience Communities"**	05/22/13
Department of Commerce EDA/NIST/ETA – miia2013 – Make it in America Challenge **"IntelSim: Boosting Rural Ohio Manufacturing through Intelligent Simulation"**	05/31/13
National Science Foundation – unsolicited **"SC13, SC14 Little Fe Workshop"**	06/27/13

Awarded OSC Proposals
as of 08/01/13

Funding Opportunity	OSC Award Amount
Ohio Third Frontier Innovation Platform Program **"Intelligent Simulation Platform for Product Commercialization"**	$2,499,936
National Science Foundation – NSF 10-599 – Supplemental Opportunities for Enhancing Diversity in the Geosciences (OEDG) **"OEDG Track 1: Expanding Geoscience Diversity through Simulated Field Environments for Students with Physical Disabilities"**	$19,391

Additional Collaborations in Discussion

- Expanded XSEDE collaboration
- Department of Defense HITS Program
  - HPCMP IT Services
- Condo of Condos Program
  - Collaborative effort to integrate networking, database, and HPC related activities to improve research productivity
  - 11 institutions, led by Clemson University
  - NSF proposal status pending
  - Initial projects
    - WAN performance testing between OSC and Clemson
    - Campus Bridging Workshop at TACC
      - http://citi.clemson.edu/cbw2013

XSEDE Activities

- XSEDE'13 Conference in San Diego, July 22-25
  - Papers
    - "Providing a Supported Online Course on Parallel Computing," Steve Gordon, et. al.
    - "OSC OnDemand: A Web Platform Integrating Access to HPC Systems, Web and VNC Applications," David Hudak, et. al.
  - Student Presentations
    - "Performance Modeling of GPU-based High Energy Physics Analysis Software," Dion Paul (Dr. Tomko)
    - "Hybrid Parallel Application Development on Xeon Phi," James Currie (Dr. Hudak)
- New XSEDE Allocations
  - XRAC, SDSC Gordon – Nationwide Children's Hospital
  - Startup, TACC Stampede – OSU CSE

Summer Research Interns

- Funded by NSF, in conjunction with Wittenberg University
- Two undergraduate students on-site at OSC for summer
- “Investigating Options for a More Portable GooFit Package”
  - Cole Taylor, working with Dr. Tomko
- “Multilevel Parallel Programming Best Practices: Examples and Tutorial”
  - Eric Mann, working with Dr. Hudak

Collaboration with the MVAPICH Group

- MVAPICH User Group (MUG) Meeting ’13
  - Aug 26-27
  - OSC Bale Theater
  - http://mug.mvapich.cse.ohio-state.edu/
- Publications
  - K. C. Kandalla, H. Subramoni, K. Tomko, D. Pekurovsky and D. K. Panda, A Novel Functional Partitioning Approach to Design High-Performance MPI-3 Non-Blocking Alltoallv Collective on Multi-core Systems, Int'l Conference on Parallel Processing (ICPP '13), October 2013
  - J. Jose, S. Potluri, K. Tomko and D. K. Panda, Designing Scalable Graph500 Benchmark with Hybrid MPI+OpenSHMEM Programming Models, Int'l Supercomputing Conference (ISC '13), June 2013

OnDemand Usage

- “Soft launch” January 2013
- Friendly users
- Training classes
- SUG announcement in March
- Steady increase in interested users

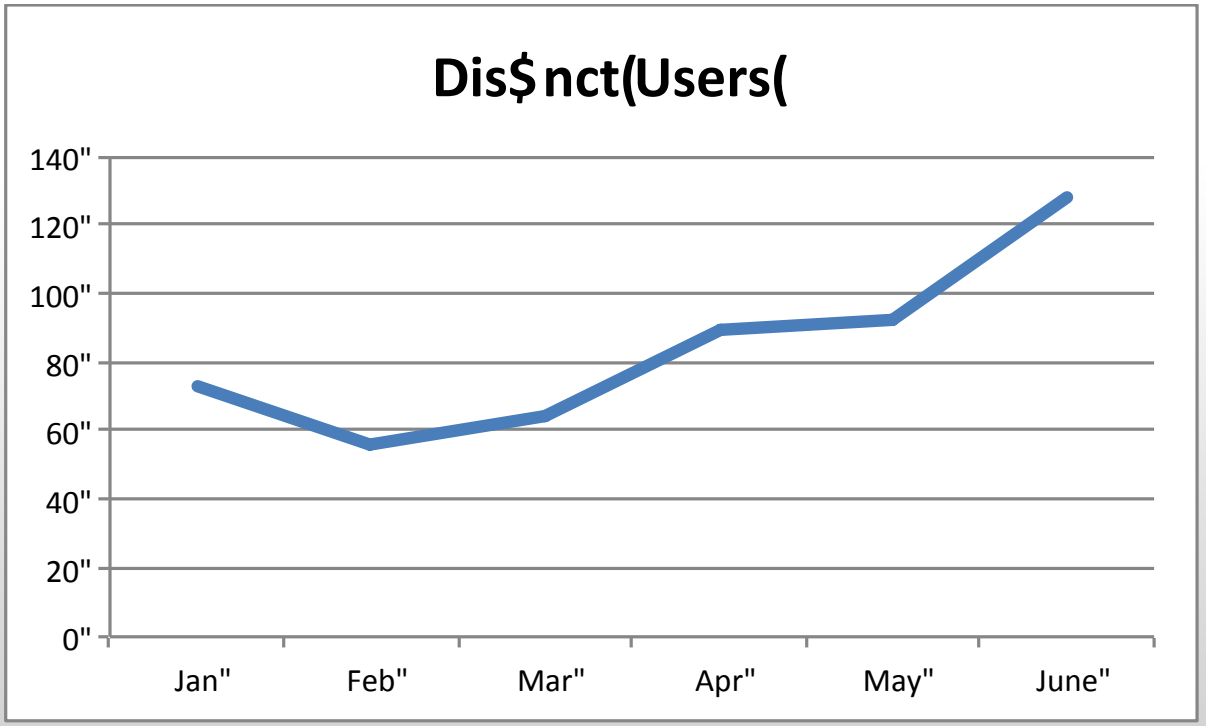

OnDemand Usage

- Distinct users up from 60 to 120 (2X)
- App launches up from 500 to 3000 (6X)
- Indicates users are doing more with OnDemand

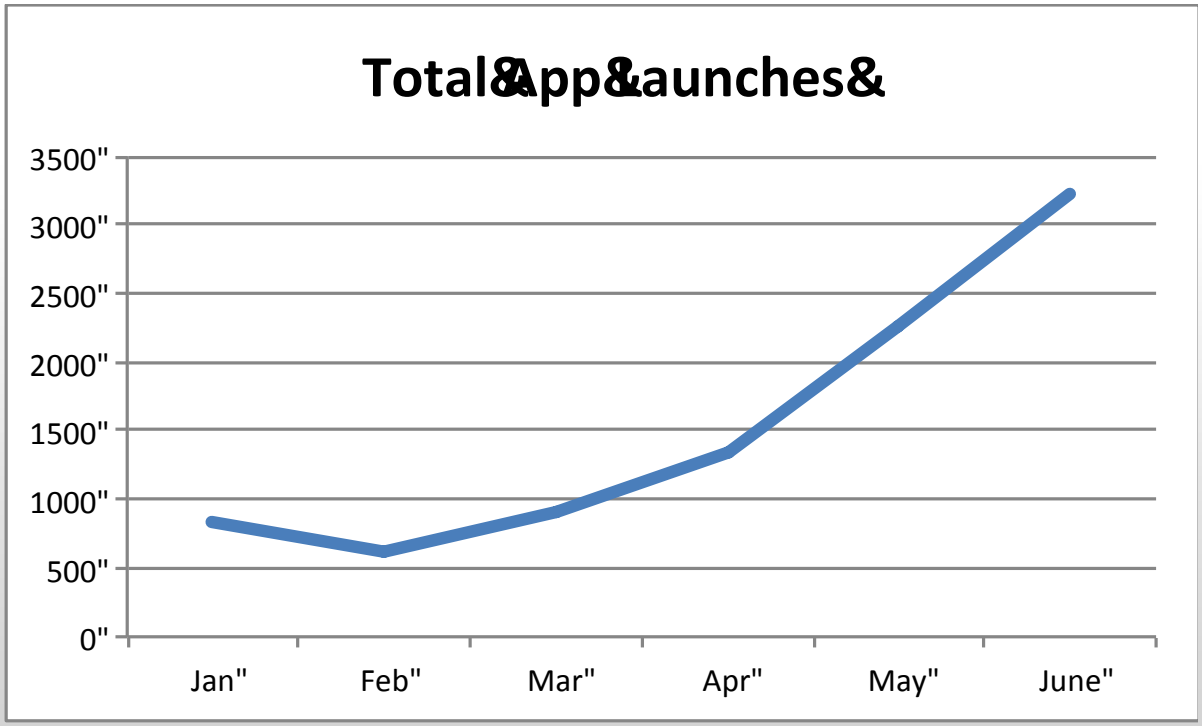

OnDemand Usage

- Four app categories shown
  - File
  - Job
  - Terminal
  - VNC

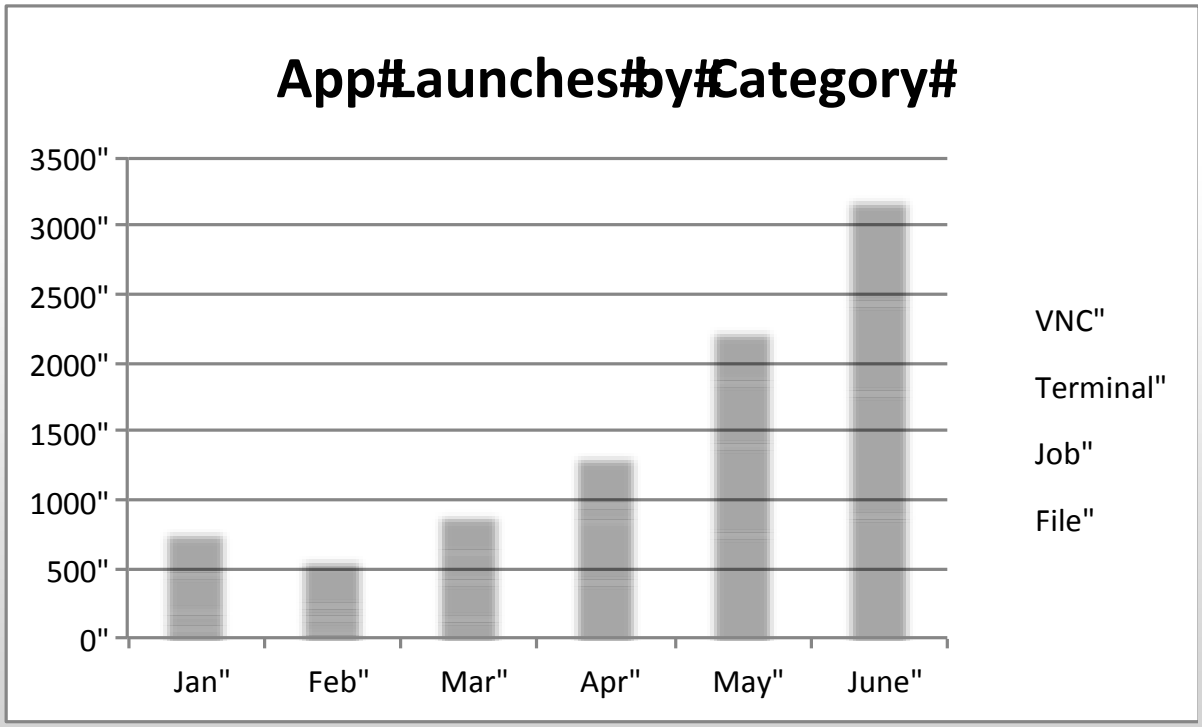

OnDemand Roadmap

- Oakley desktops
- Improved data transfer options
  - Sftp and Globus Online
- Improved job interface
  - “My Jobs” to have all user’s jobs, not just ones created by app
  - Show performance statistics for completed jobs
  - Git for job templates and sharing user-created templates
- VNC solutions not browser-based, not firewall friendly
  - Interested in HTML5 VNC client, like Guacamole
- Support authentication via shibboleth from home institutions

Ohio Supercomputer Center

An OH·TECH Consortium Member

Brian Guilfoos, Science & Technology Support Manager

User Support, Education and Training

Agenda

- Odds and Ends
- LUSTRE challenges
- Memory containers
- HPC software
- Academic user authorization
- Incidents
- Compute workload distribution
- Queue wait times
- Upcoming training

Please ask questions as we go! Lots of material today!

General

- Working towards replacement of ARMSTRONG
  - Expect to replace password/shell management functionality later this year
  - RU utilization charts, other PI functions to migrate later
  - Looking to “modernize” user authorization
- 2 summer interns from Wilberforce University
  - Helped develop database for tracking software agreements

LUSTRE Data Expiration

- LUSTRE is getting full
  - Using approximately 413T of 626T available (~67%)
  - More than 65 million files; takes several days to traverse filesystem
- As LUSTRE is temporary space, we will start expiring old data
  - Files untouched for 180 days will be removed (currently 19+ million files)
  - Removal will occur every day or once a week
  - In the event of the file system being full, we will delete the oldest files first until the filesystem is usable again.
- Eventually: 1 million file count quota per user
  - Reminder: LUSTRE performs best with big files!

Memory Containers and RU Charges

- Technical solution in testing
- Limits users to a maximum of 4GB/core* (Oakley)
- Memory requests in batch scripts will now be limits
- Two phase rollout
  - Better diagnostic information at the end of jobs
  - Memory enforcement & charging changes
    - nodes=1:ppn=3
      - Implicit limit of 12 GB of RAM
    - nodes=1:ppn=1 / mem=12GB
      - Will be assigned 1 core, 12 GB of RAM
      - Will be charged for 3 cores (12GB/(4GB/core))
- Parallel jobs unaffected; already charged for all cores

HPC Software

- Center initiative to more proactively manage user facing software on our HPC systems
- Types of questions that we have for users
  - Currently default modules (compilers, MPI,..) don't change after the system is bought online. Should there be periodic updates?
  - What is the right suite of software for our primary science domains?
- We would like to put together some domain focus groups to further this discussion
  - If you or someone from your research group would like to participate, contact ktomko@osc.edu or the help desk.

Academic User Authorization

- We need a better handle on tracking who is using OSC systems
- Coming changes may include:
  - Forced regular password changes
  - Periodic PI authorization of users
  - More aggressive enforcement of “no shared accounts” policy
  - Occasional (re)authorization to user’s home institution
  - Updated Acceptable Use Policy (replacing our code of ethics), and a requirement for all academic users to acknowledge receipt and agree to abide by it.

Incidents

- 132 new tickets in the last 120 days
  - 57 in “High Performance Computing”
  - 55 in “Statewide Software Licensing”
  - Remainder scattered across other categories
- 106 closed or resolved
- 16 active
- 10 waiting

Job Distribution – Number of Jobs by Core Count

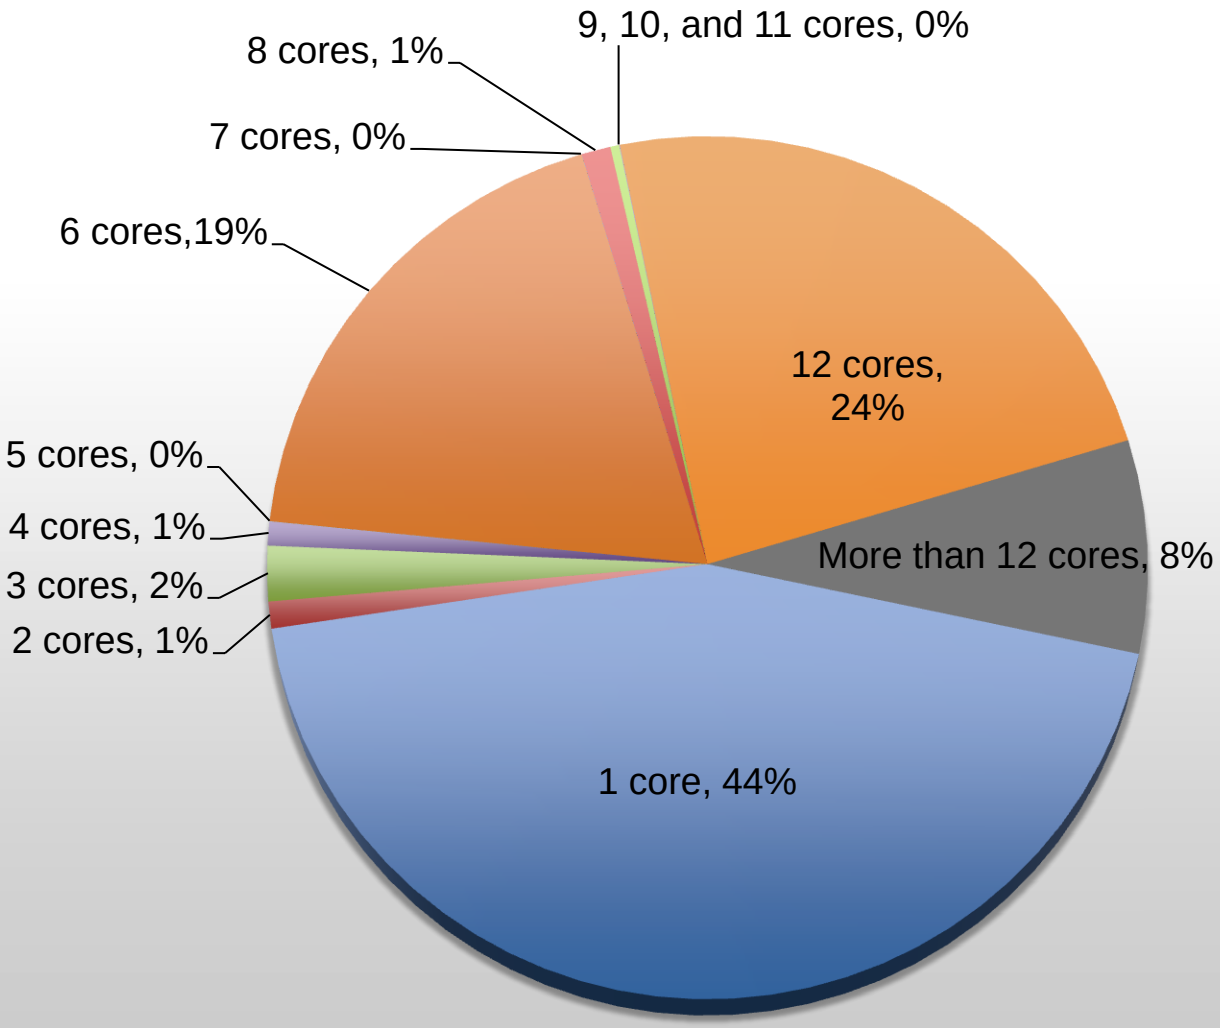

Job Distribution – Number of Jobs, 14 Cores or More

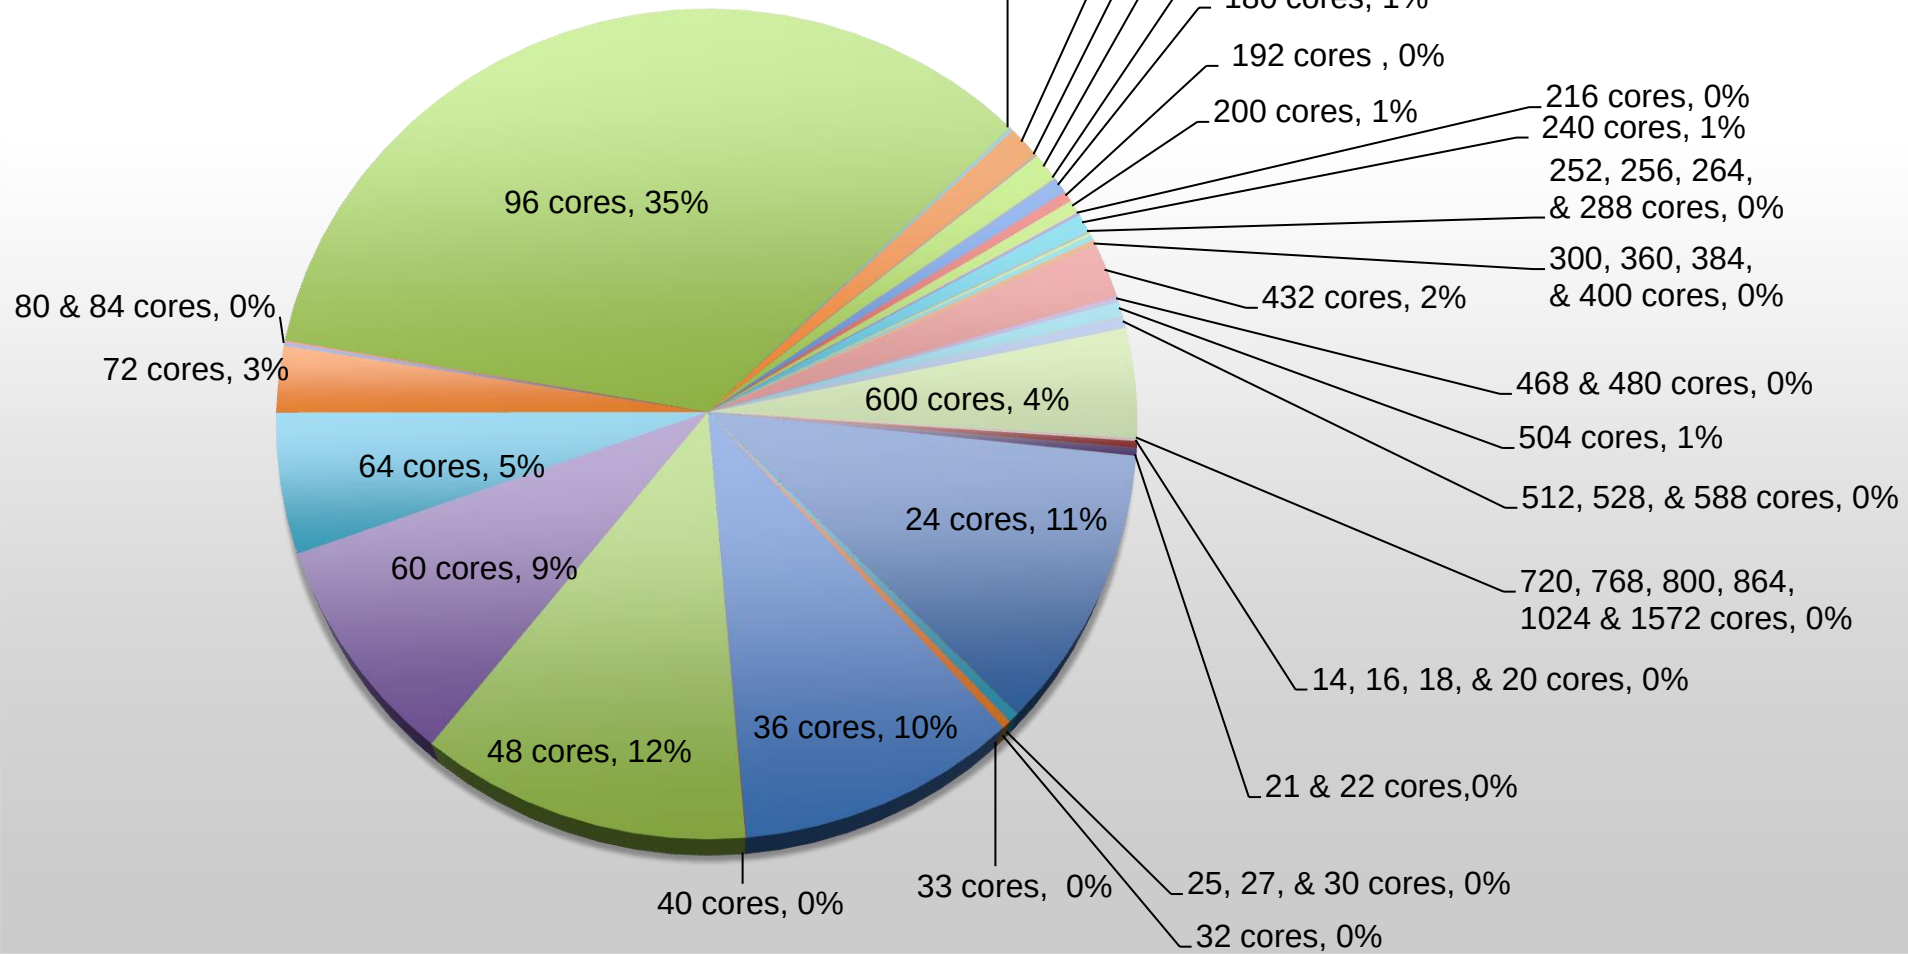

Job Distribution – Compute Hours by Core Count

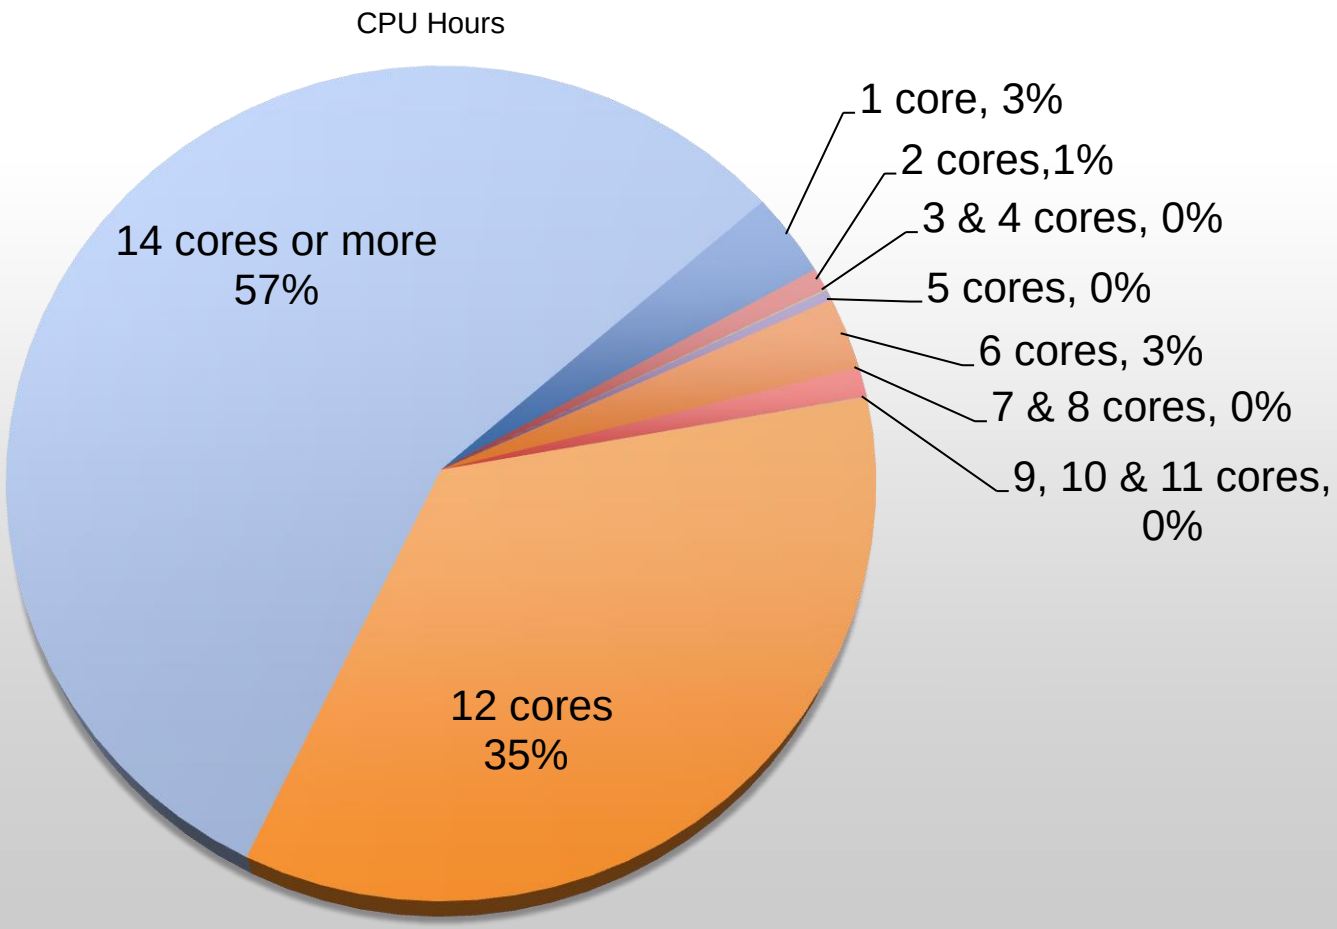

Job Distribution – Compute Hours by Core Count, 14 or More Cores

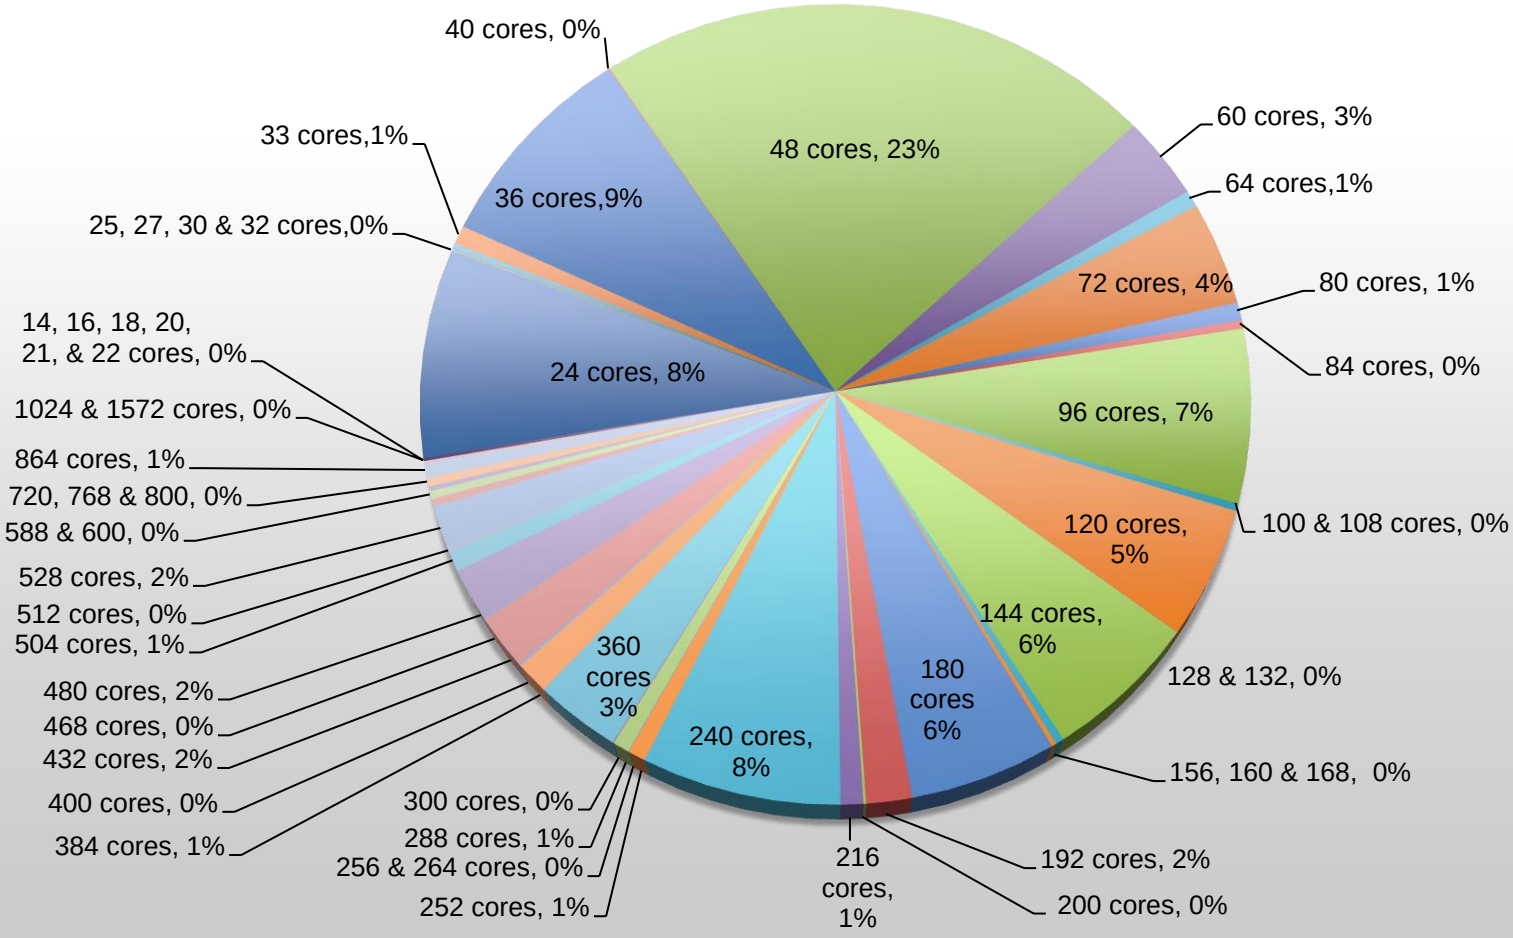

Time Spent in the Queue

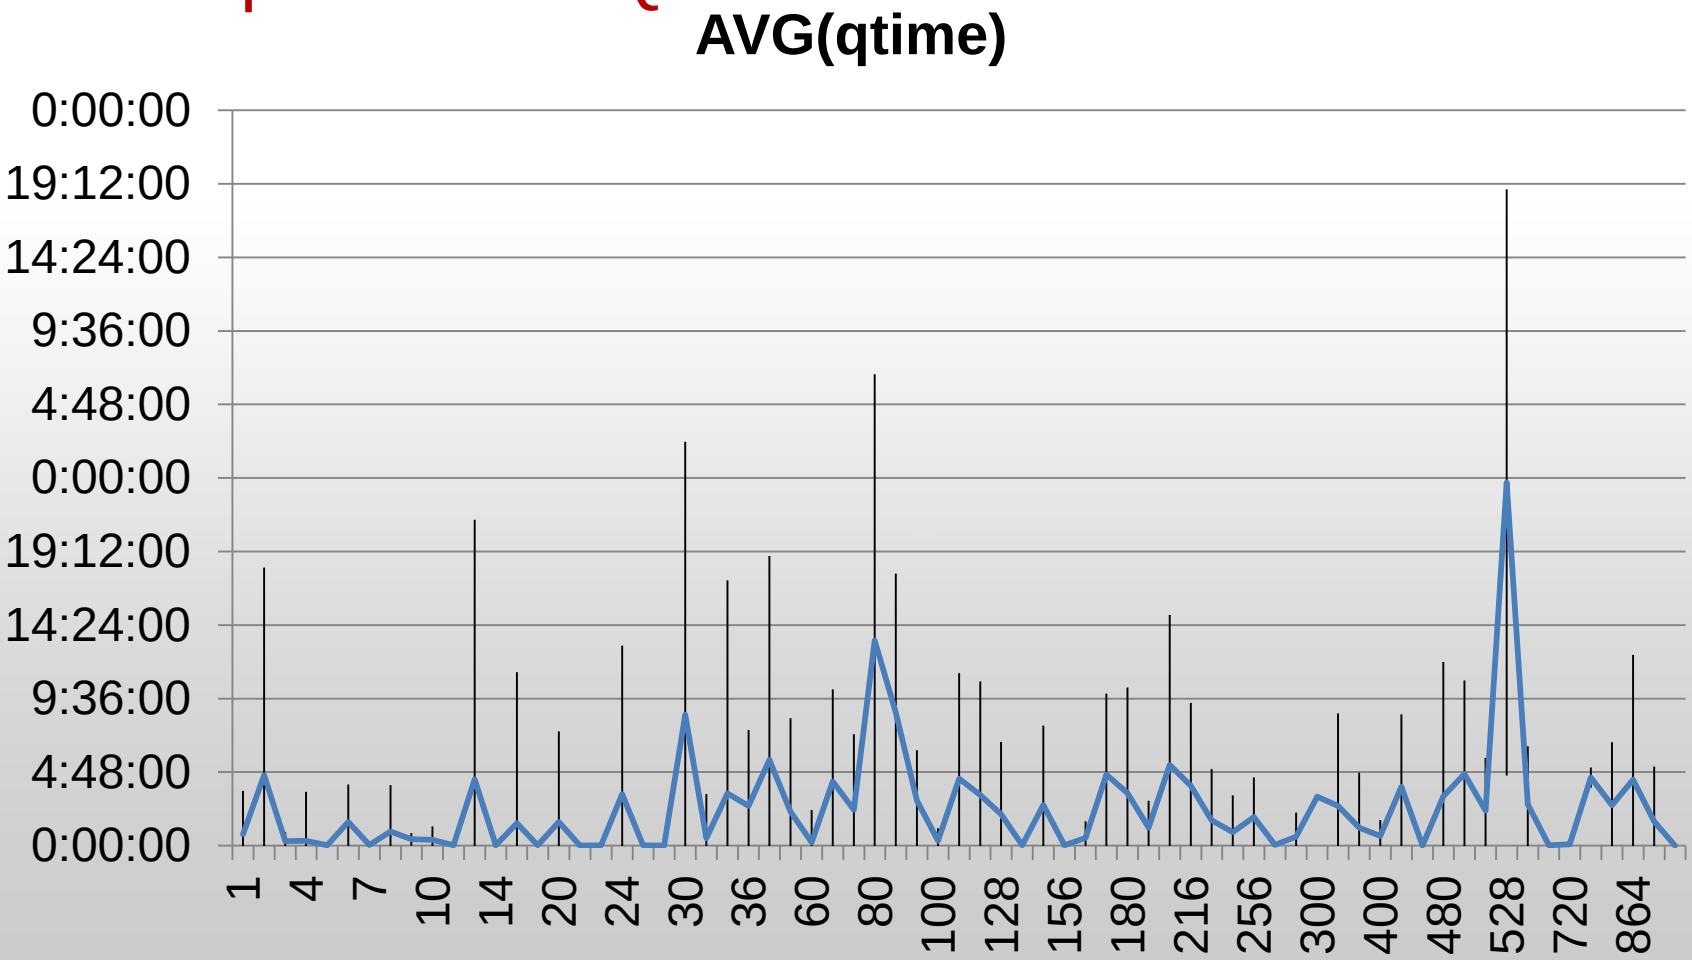

Training

- Confirmed
  - Intro to OSC, @ OSU (PRB 1060)
    - Sept 3, 3-5pm
  - Intel Optimization for Xeon & Phi @ OSC
    - Sept 25, 9am-5pm
- Still working on details
  - Intro to OSC @ Cleveland State
    - 9am-12:30pm, meeting with research groups by appt after
  - Intro to OSC @ UT/BGSU?
  - Intro to OSC @ UC?
  - Intro to OSC @ OU?

Training

- Partnering with PSC to offer remote XSEDE training (tentative, NOT CONFIRMED dates)
  - OpenMP, Oct 1
  - OpenACC, Nov 5-6
  - MPI, Dec 3-4
  - Big Data (Hadoop & SPARQL), Feb 4
- Video training for OnDemand
  - In production
  - To be posted on osc.edu in the OnDemand documentation

Stay up to date!

- www.osc.edu/supercomputing
  - Entry point for all documentation
  - Current “Known Issues” listed at the bottom of the page
- osc.edu/n
  - System notices (downtimes, etc)
- osc.edu/e
  - Upcoming events (training, etc)
- Follow us on Twitter for push notices
  - www.twitter.com/HPCNotices
- For help, contact OSC Help
  - oschelp@osc.edu

Ohio Supercomputer Center

An OH·TECH Consortium Member

Rick Prairie

Software Committee

Software Updates

- Turbomole
  - Updated to version 6.3.1
- Intel compilers (Fortran, C, C++)
  - Updated to version 13.1.3

Renewals – Calendar Year 2013

Software	Expiration (2013)	License Term (years)	Projected Cost	Renewal Status/Started
MATLAB Distributed Computing Server	March 1	1	$3,750	Completed
CSD	April 1	1	$1,500	Completed
COMSOL	June 15	1	$4,374	Completed
OSU MATLAB	Sept. 30	1	$19,464	No
LS-Dyna	Dec. 27	1	$500	No
Gaussian	Dec. 31	1	$12,000	No

Ohio Supercomputer Center

An OH·TECH Consortium Member

John Heimaster

Hardware Committee

The OH-TECH Consortium

Ohio Supercomputer Center provides high performance computing, software, storage and support services for Ohio's scientists, faculty, students, businesses and their research partners.

OARnet connects Ohio's universities, colleges, K-12, health care and state and local governments to its high-speed fiber optic network backbone. OARnet services include co-location, support desk, federated identity and virtualization.

OhioLINK serves nearly 600,000 higher education students and faculty by providing a statewide system for sharing over 50 million physical and digital library resources, while aggregating costs among its 90 member institutions.

eStudent Services provides students increased access to higher education through e-learning and technology-enhanced educational opportunities, including virtual tutoring.

Research & Innovation Center will operate, when opened, as the proving grounds for next-generation technology infrastructure innovations and a catalyst for cutting-edge research and collaboration.

Questions

Pankaj Shah
Executive Director,
Ohio Supercomputer Center and OARnet
pshah@oh-tech.org

Like Us on Facebook:
https://facebook.com/ohiosupercomputercenter
Follow Us on Twitter:
https://twitter.com/ohiosupercomputerctr

1224 Kinnear Road
Columbus, OH 43212
Phone: (614) 292-9191