

# HPC, Current State @ UW

UW/WWSU/PNNL Exploratory Workshop, University of Washington  
July 27, 2015



# Overview - Part I

- High Level View of UW Cyberinfrastructure
  - What do we do?
  - When was/is it all happening (timeline)?
  - Why do we do this?
  - How are we doing?

# Overview - Part II

- The Hyak Model
- Personal Supercomputers
- Traditional Supercomputers
- Having it all
- Does it work? Yes!

# Overview - Part III

- The Hyak Mission & A Little Vision
- The Long Tail Escalator
- Partnership & Opportunity
- Advanced User Support
- Outreach & Curation

**Q: What?**

**A: Cyberinfrastructure**

**Cyberinfrastructure is the coordinated aggregate of software, hardware and other technologies, as well as human expertise, required to support current and future discoveries in science and engineering.**

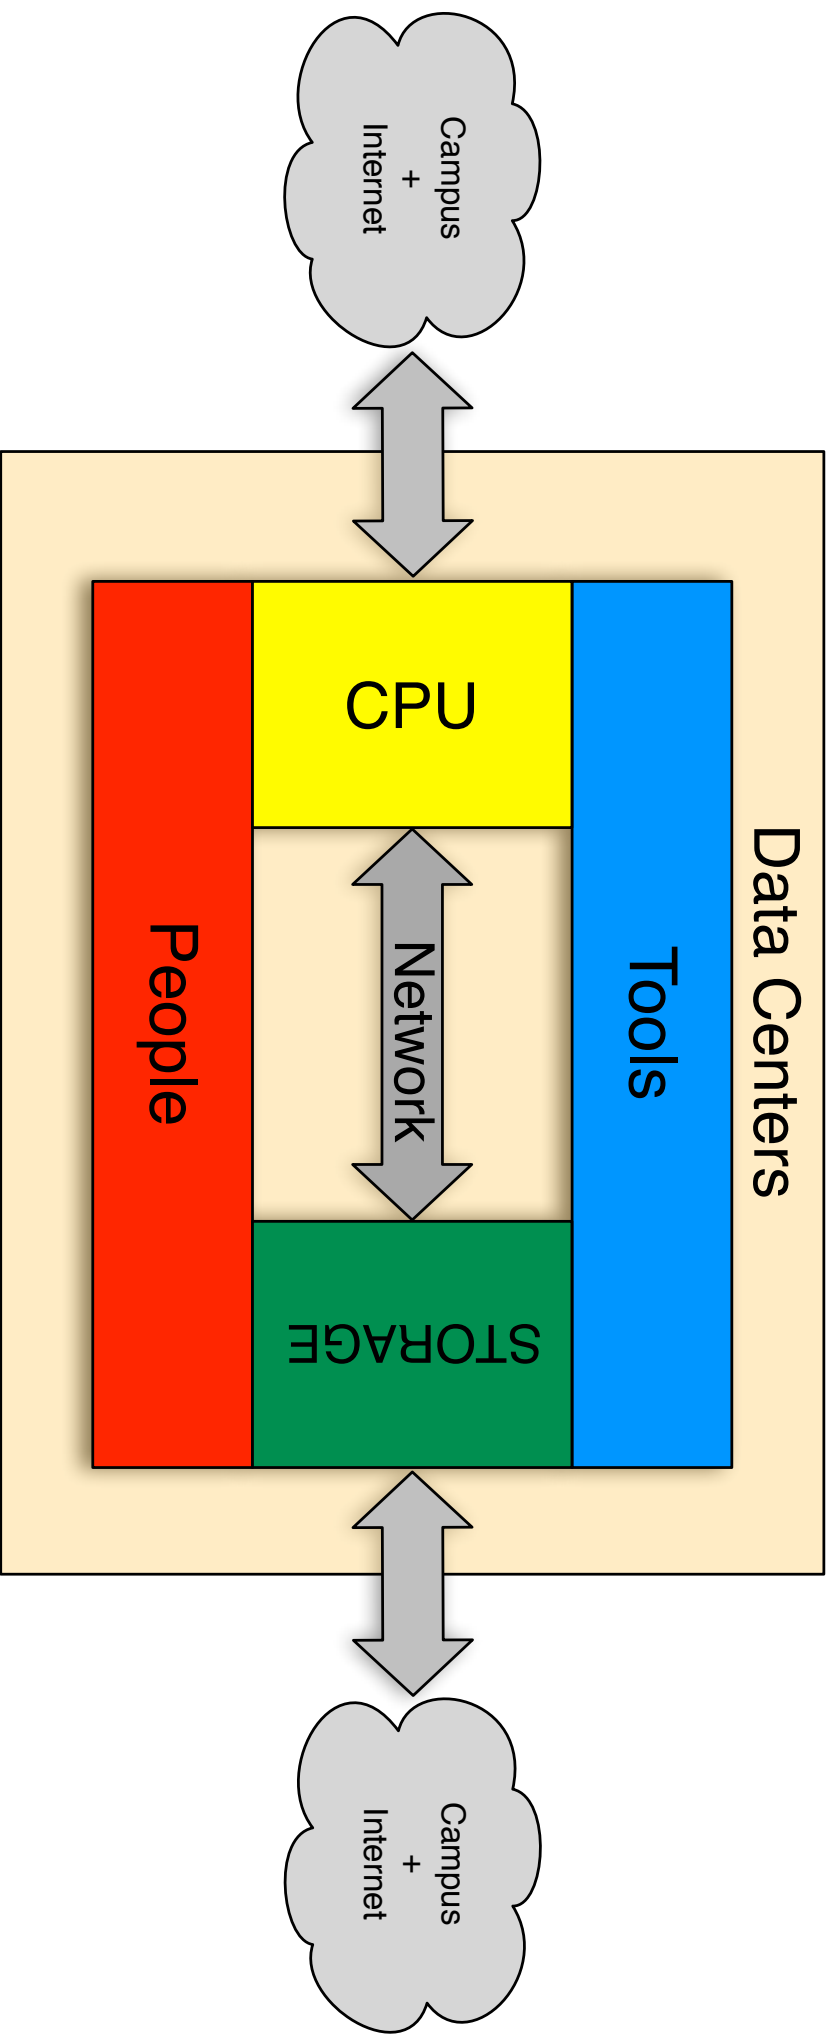
**The challenge of Cyberinfrastructure is to integrate relevant and often disparate resources to provide a useful, usable, and enabling framework for research and discovery characterized by broad access and “end-to-end” coordination.**

\* Francine Berman (February 18, 2005). "SBE/CISE Workshop on Cyberinfrastructure for the Social Sciences". San Diego Supercomputer Center.



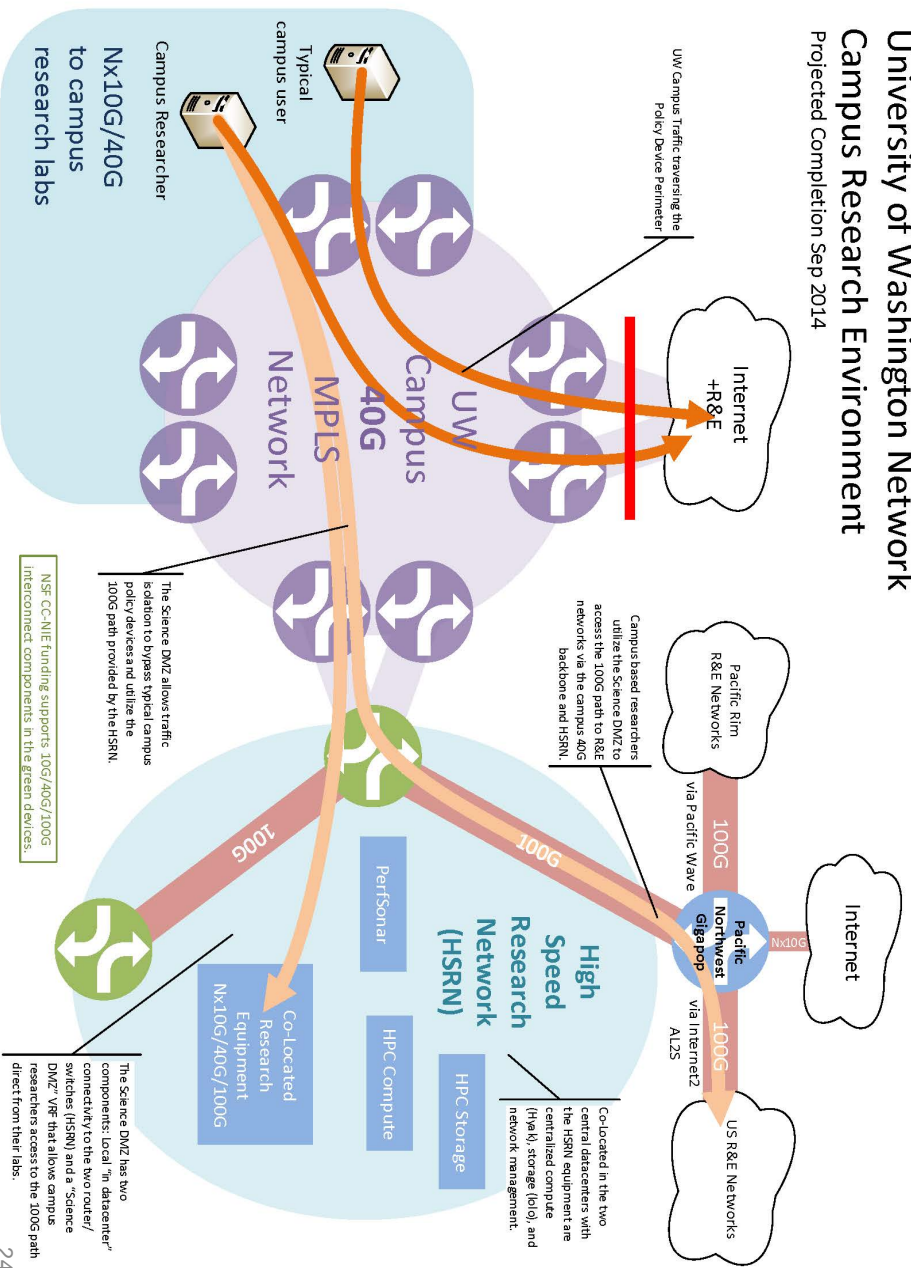
Q: What?

A: Cyberinfrastructure



# HSRN & Science DMZ

## University of Washington Campus Research Environment Projected Completion Sep 2014



Q: What?

A: On-site HPC @UW

- UW-IT is Honest Broker
  - Help connect users with most appropriate resources
  - Whether those resources are local or remote
- **Recognize an Important Role for On-Site HPC**
  - **Speed of Science**
  - **Prep for PreExascale**
  - **Big Data Pipelines**
- **Data Privacy for “Cloudy” Workloads**





# Q: When?

## A: Started in 2010

- 2000 steep uptick in computational research (Linux, cheap x86, Beowulf)
- 2005 data centers overwhelmed
- 2005 “data science” discussions begin
- 2007 VPR convenes forums to discuss solutions
- 2007 - 08 Conversations with UW Research Leaders
- 2008 VPR hires first eScientist (me)
- 2008 eScience Rollout event
- **2010 Hyak / Iolo / primitive Science DMZ (MRI + UW “club” — domain scientists/engineers solved for themselves)**
- 2014 High Speed Research Network / Science DMZ



Q: When?

A: Set Through 2024 and Beyond

2010 2011 2012 2013 2014 2015 2016 2017 2018 2019 2020 2021 2022 2023 2024 2025 2026

Hyak Phase-1

Hyak Phase-1 Refresh

Hyak Phase-2

Hyak Phase-3

????



Q: Why (Leadership POV)?

A: Competitiveness & CO2

- Faculty recruitment & retention
- Data center space crisis
- Climate Action Plan
- The scaling problem

**W**

Q: Why (user POV)?

A: Fill the Gap


- Midscale Capacity/Capability Computing
  - Speed of Science / Thought
  - Prep for PreExascale
  - Big Data Pipelines
- Data Privacy for “Cloudy” Workloads

**W**

# Q: How Happy?

# A: Very!

**Hyak Core-Hours Used in Standard Queues  
by job core-count & runtime**



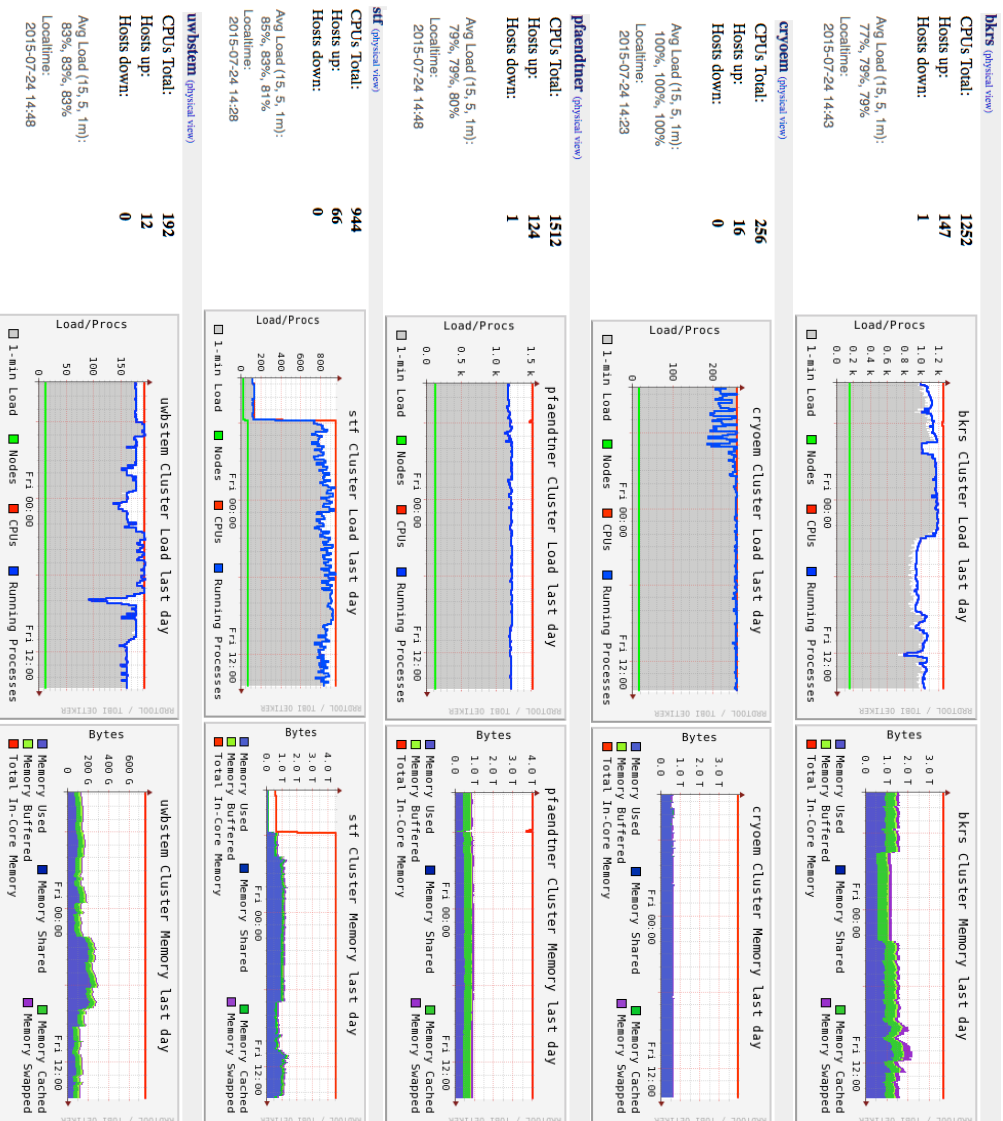
# Hyak - Campus Condo Cluster

A Little Background



**W**

# Hyak - Personal Supercomputers



# Hyak - A Traditional Supercomputer FREE!



Compute Cluster Report for Fri, 24 Jul 2015 14:22:59 -0700






Get Fresh Data

Metric  Last  Sorted

Hyak Grid > Compute > --Choose a Node


## Overview of Compute

CPU's Total: **10396**  
 Hosts up: **851**  
 Hosts down: **2**  
 Avg Load (15, 5, 1m):  
**77%, 77%, 78%**  
 Localtime:  
**2015-07-24 14:22**






# Elastic Supercomputing




# Hyak - Campus Condo Cluster

Hyak Utilization by Domain




Hyak Allocation by Domain



Q: How Well Does it Work?  
A: Very!

Hyak Core-Hours Used in Standard Queues  
by job core-count & runtime



# Hyak - a **SUPERCLOUD**

Your Personal Supercomputer

+


In a UW-IT managed “cloud”

=

*Elastic Supercomputing*

**W**

# Hyak - A Large & Growing Success




~250TF


~12,000 cores  
~900 nodes



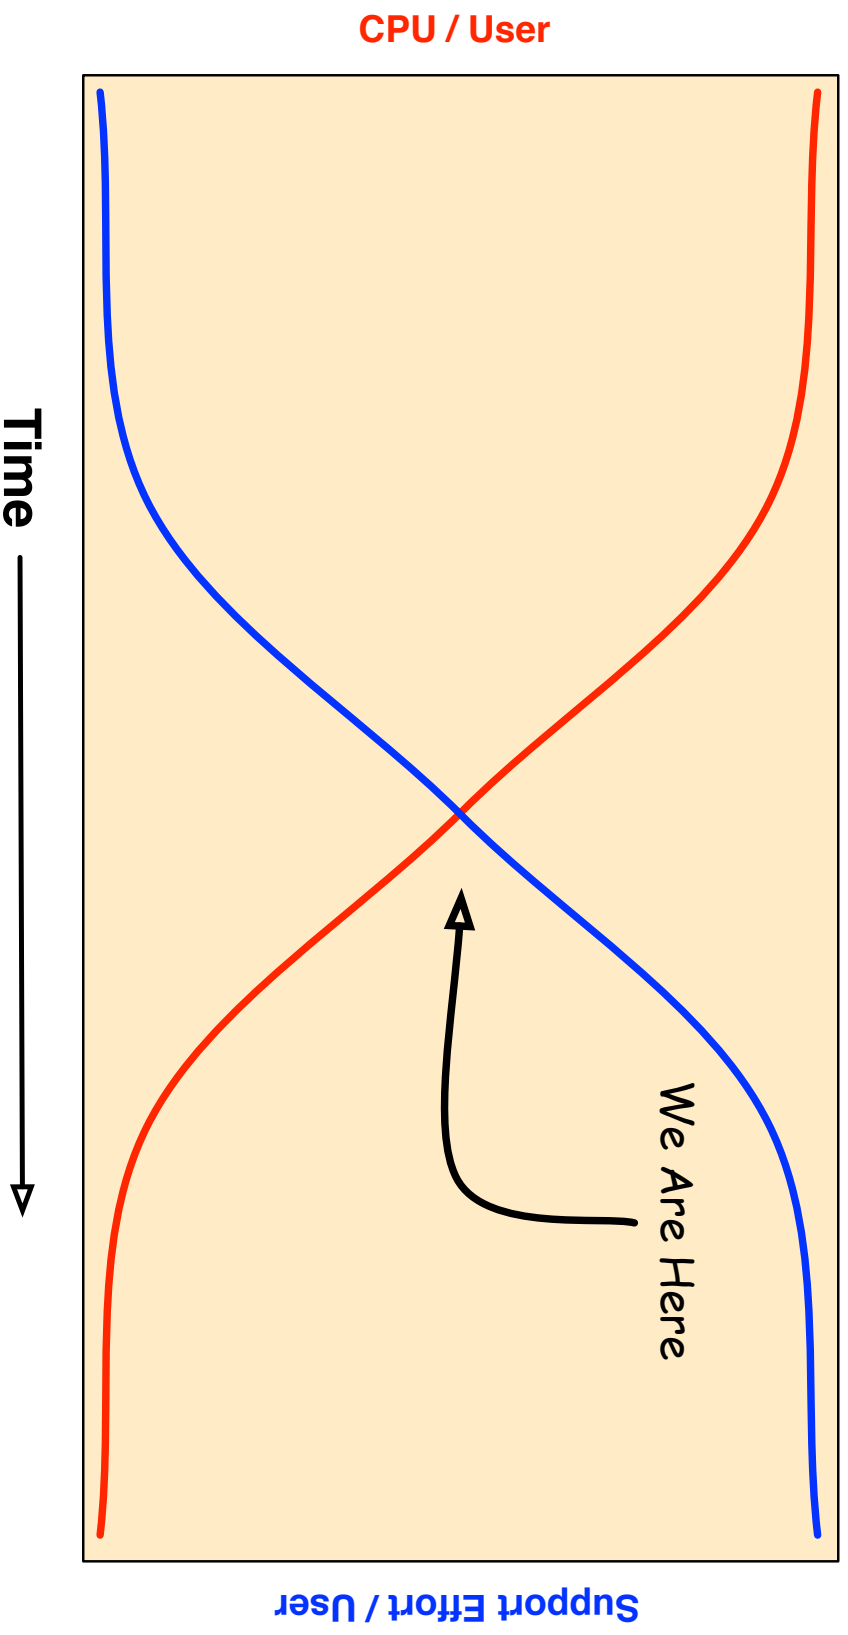
Q: What's Our Mission?  
A: The Long Tail Escalator



Q: What's Our Mission?  
A: The Long Tail Escalator



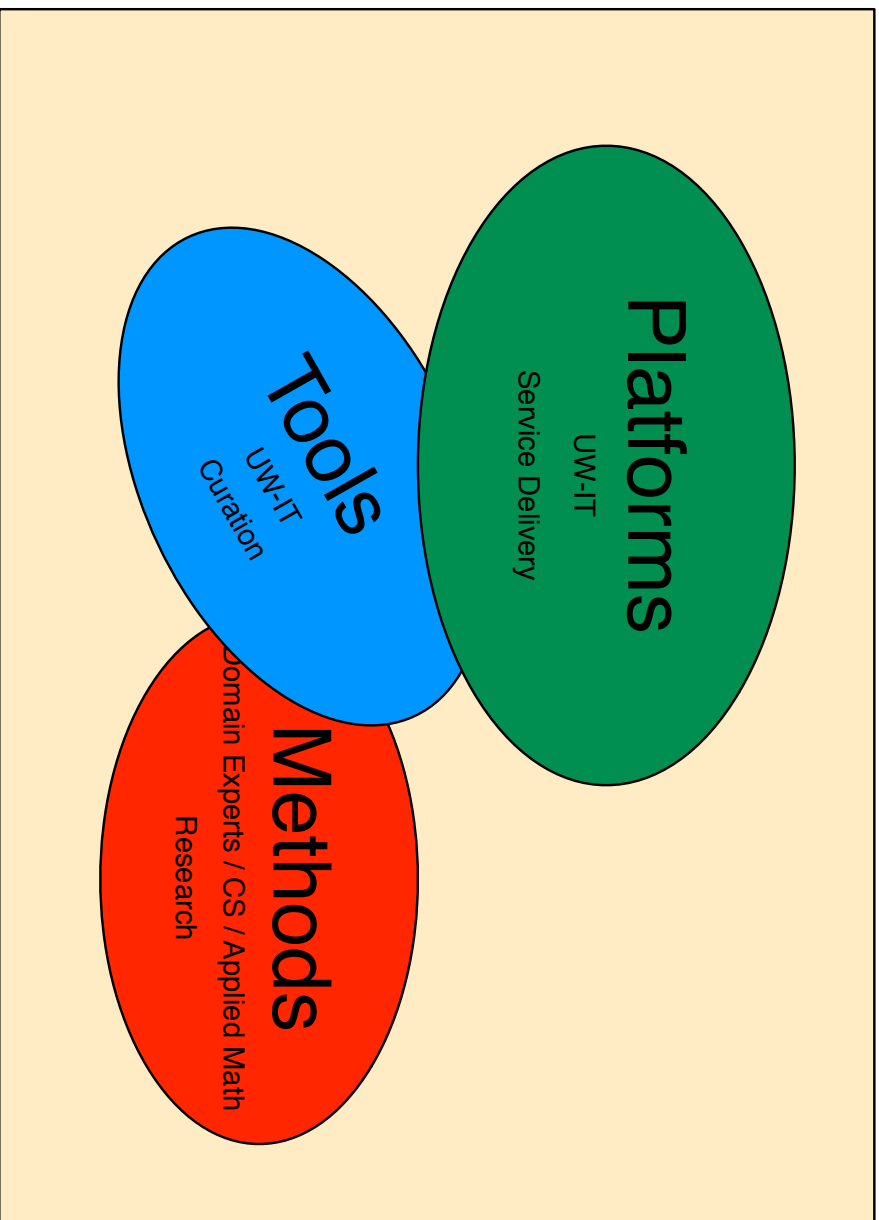
Q: Where (are we in the mission)?  
A: The Long Tail





Q: HOW?

A: Partnership & Coordination



# Q: Where (are we in the mission)? A: The Long Tail

**Ganglia**  
sourceforge.net


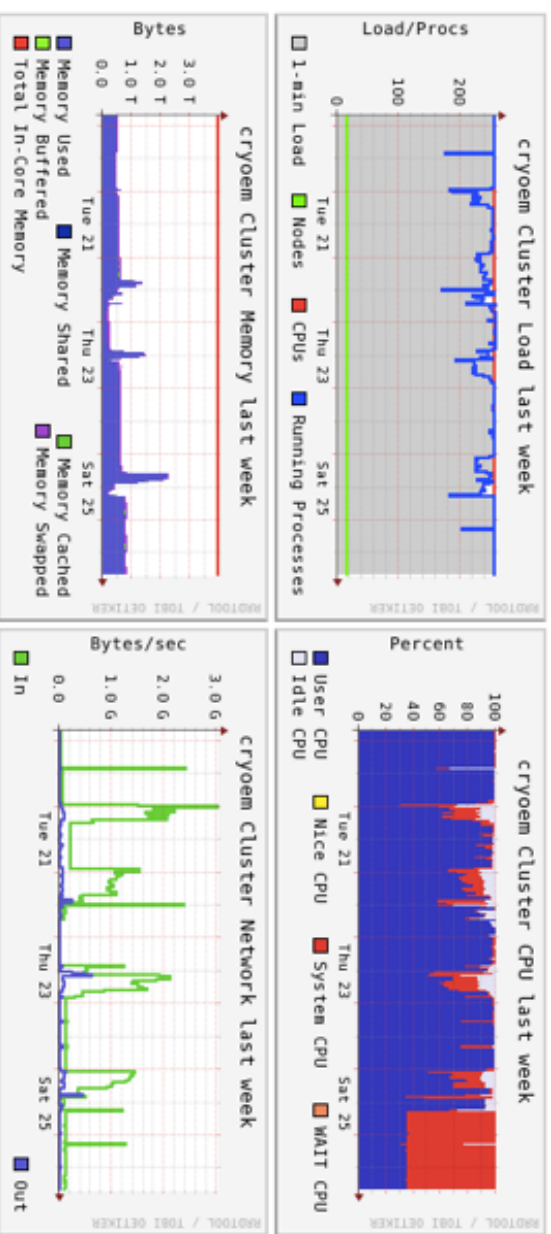
**cryoem Cluster Report for Sun, 26 Jul 2015 19:50:33 -0700**

Metric:  Last:  Sorted:

[Hyak Grid > cryoem > --Choose a Node](#)

## Overview of cryoem

CPU's Total: **256**  
Hosts up: **16**  
Hosts down: **0**  
Avg Load (1.5, 5, 1m): **100%, 100%, 100%**  
Localtime: **2015-07-26 19:50**



# HPC, Current State @ UW

UW/USW/PNNL Exploratory Workshop, University of Washington  
July 27, 2015

## Questions?

Chance Reschke <[reschke@uw.edu](mailto:reschke@uw.edu)>

