

KickStart 2018

David Swanson, PhD
Director, Holland Computing Center

“HCC operates as a core facility and provides services that are critical to the research productivity of faculty and students across the University of Nebraska campuses.”

–HCC 5-year Review, 2016

“HCC exists to reduce the time to science.”

–David Swanson

HCC is good at ...

- **Big Data:** We manage over 10,000 TeraBytes (10PB)
- **Scaled Out Computing:** 35,000 cores, OSG, advanced networking, high bandwidth transfers, Globus Online
- **Scaled Up Computing:** Parallel processing, XSEDE, GPUs

Recent Study of Data Scientists*

- “Available computing power” was the second most common problem for big data analysis.
- Distributed or parallel processing was the least common solution to their big data needs.
- This could be attributed to the difficulty of processing data on distributed resources.

*K Rexer. Data miner survey-2013 survey summary report.
Rexer Analytics, Winchester, 2013.

HCC has great people

- Most are computational scientists
- 5 Full time research personnel
- 7 Sys Admins
- 5 Application Specialists
- Favorite collaborative tools: table with a whiteboard
- Collaboration, not outsourcing outfit

HCC has great people

- 118 Schorr Center
- 152 Peter Kiewit Institute
- Buffet Cancer Center 5.12.397
- hcc-support@unl.edu
- <http://hcc.unl.edu>

HCC is free*

- *Shared usage of resources:* no charge
- *Best effort support:* no charge
- *Dedicated computing, storage or support:* up-front fees

**to many users! Paid for by NRI, NSF, and NU researchers*

Storage Options

- Currently 4 types of storage
 - /home — backed up, per machine, 25 GB/user, free
 - /work — purged, per machine, 50 TB / group, free
 - **/common — not purged, mounted on all HCC machines**
 - Attic — near-line archive, not mounted, \$25/TB/year

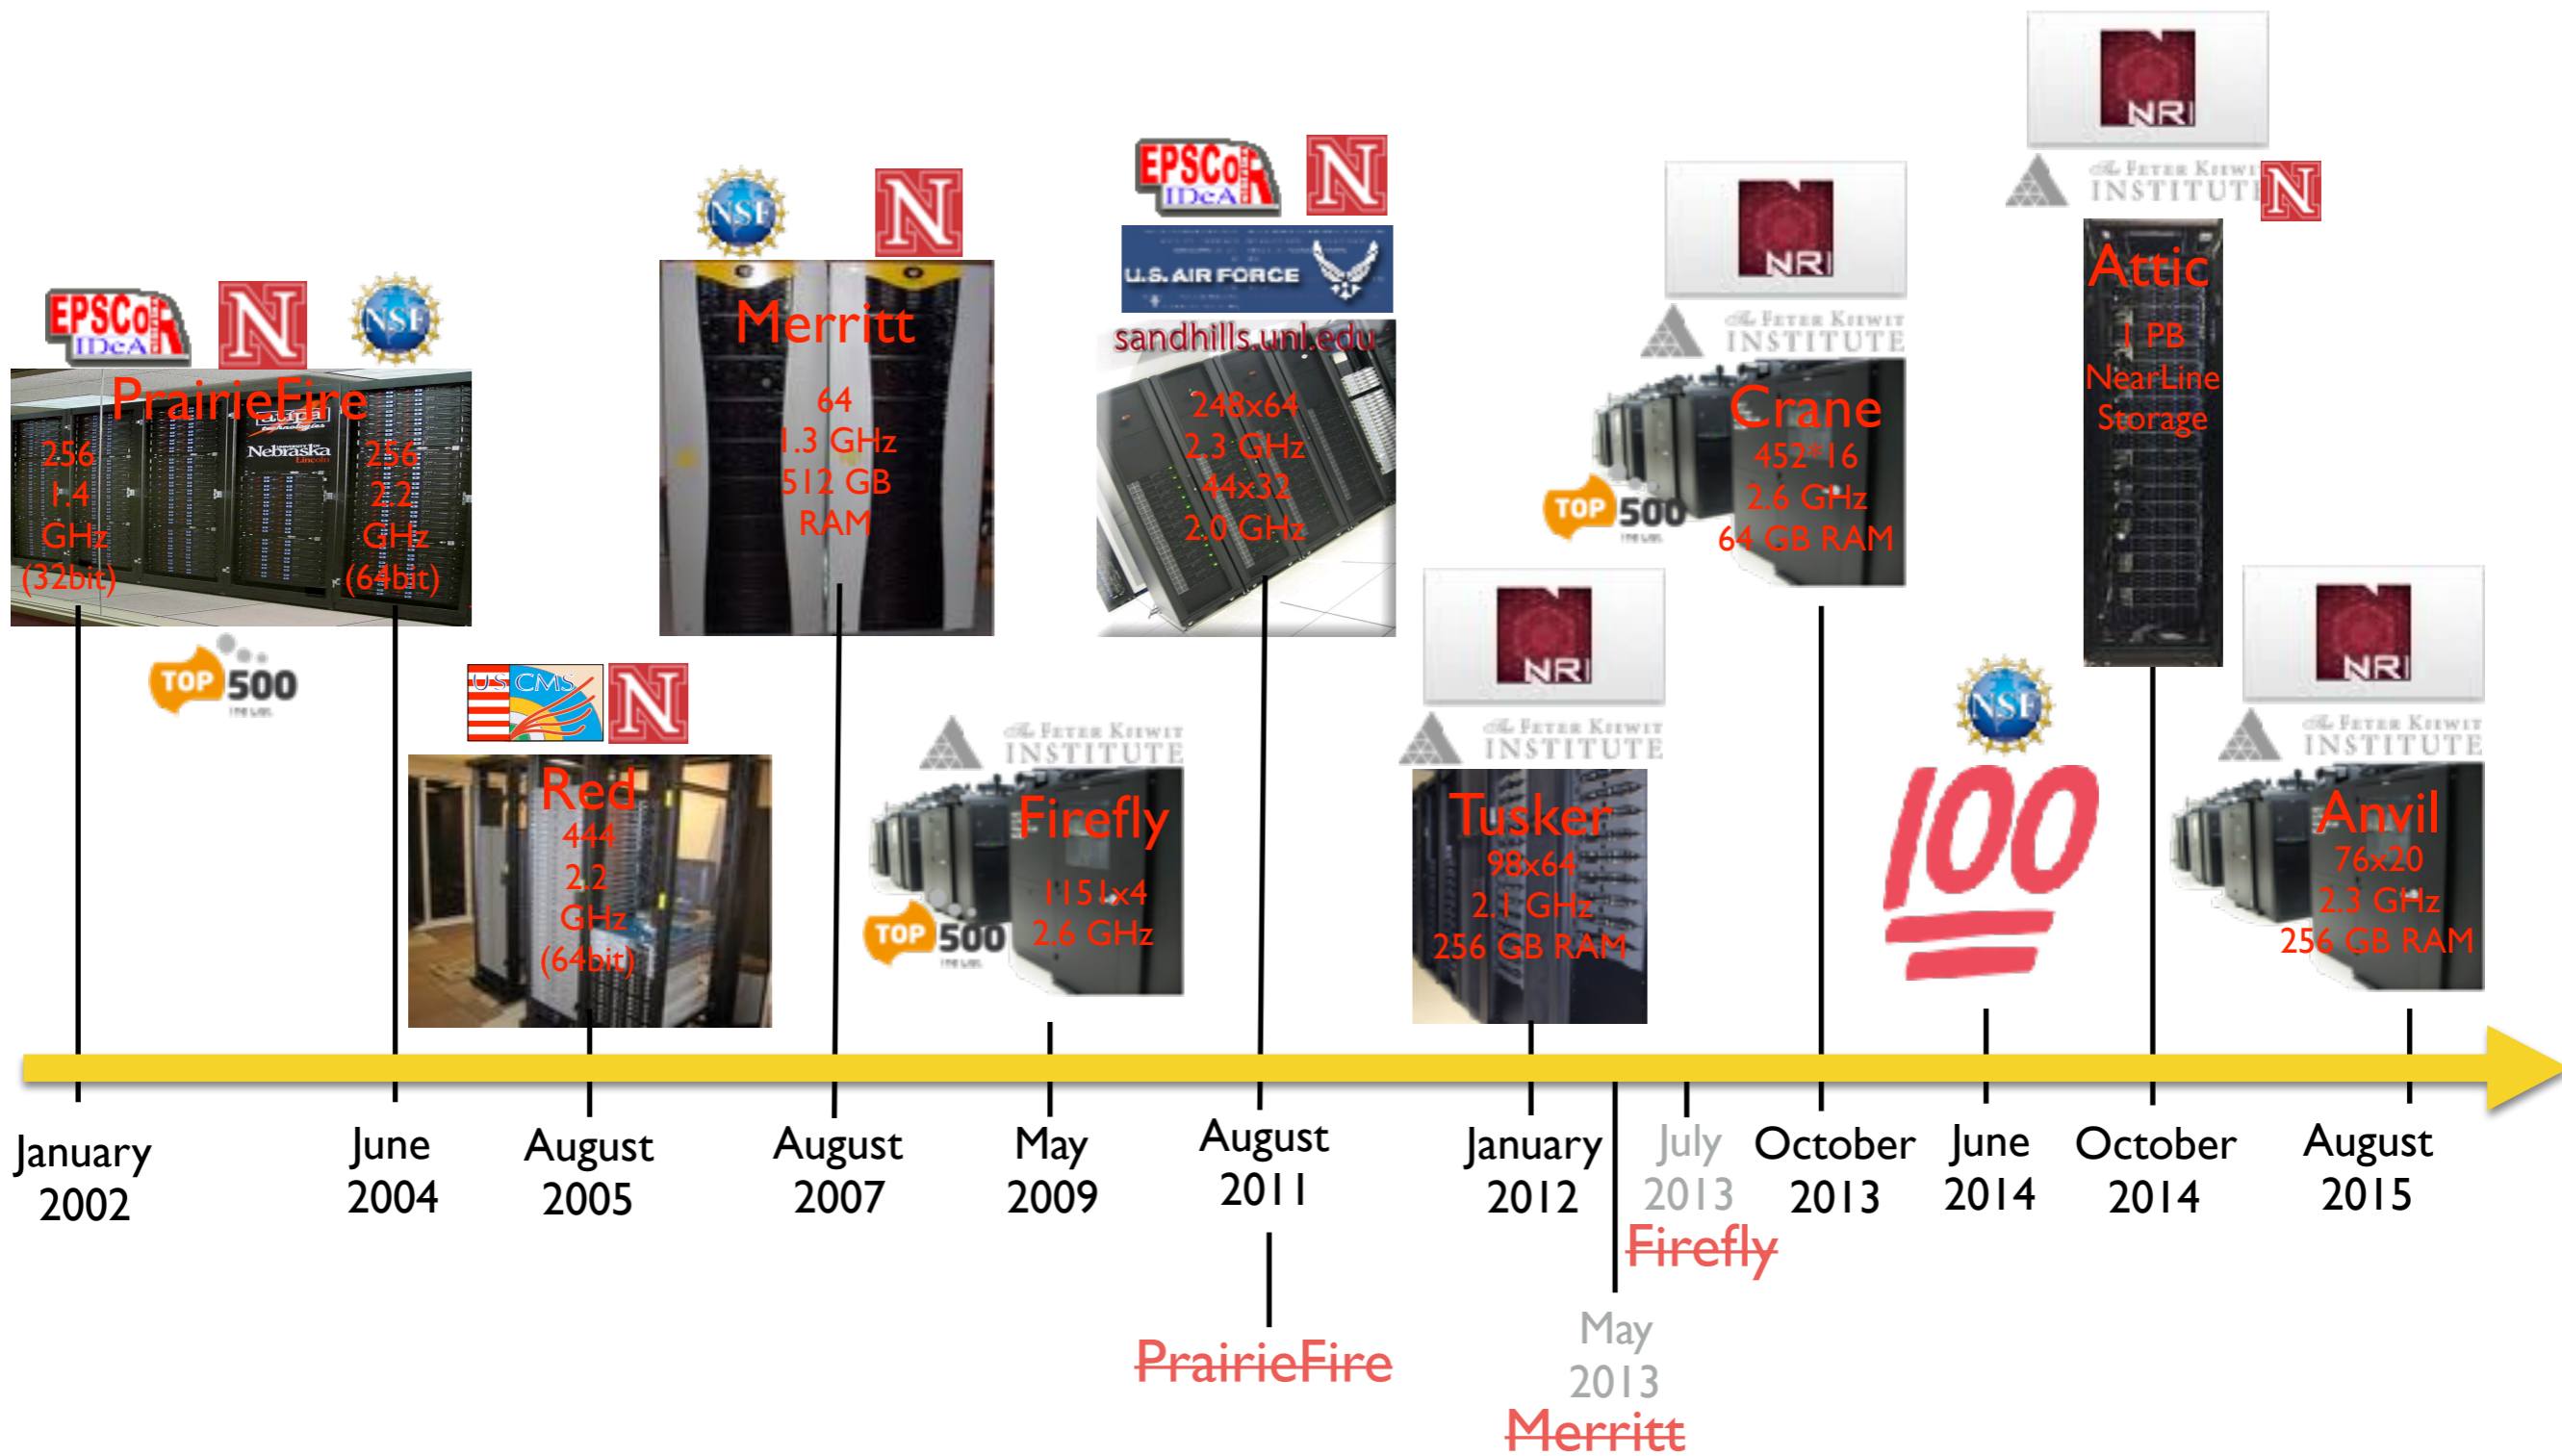
HCC On Ramp

1. Nebraska Faculty Connection (HCC Group)
2. Log in (HCC account and Duo set up)
3. Moving and editing your data (Globus)
4. Launch jobs! (module, SLURM) or (Jupyter Notebooks)

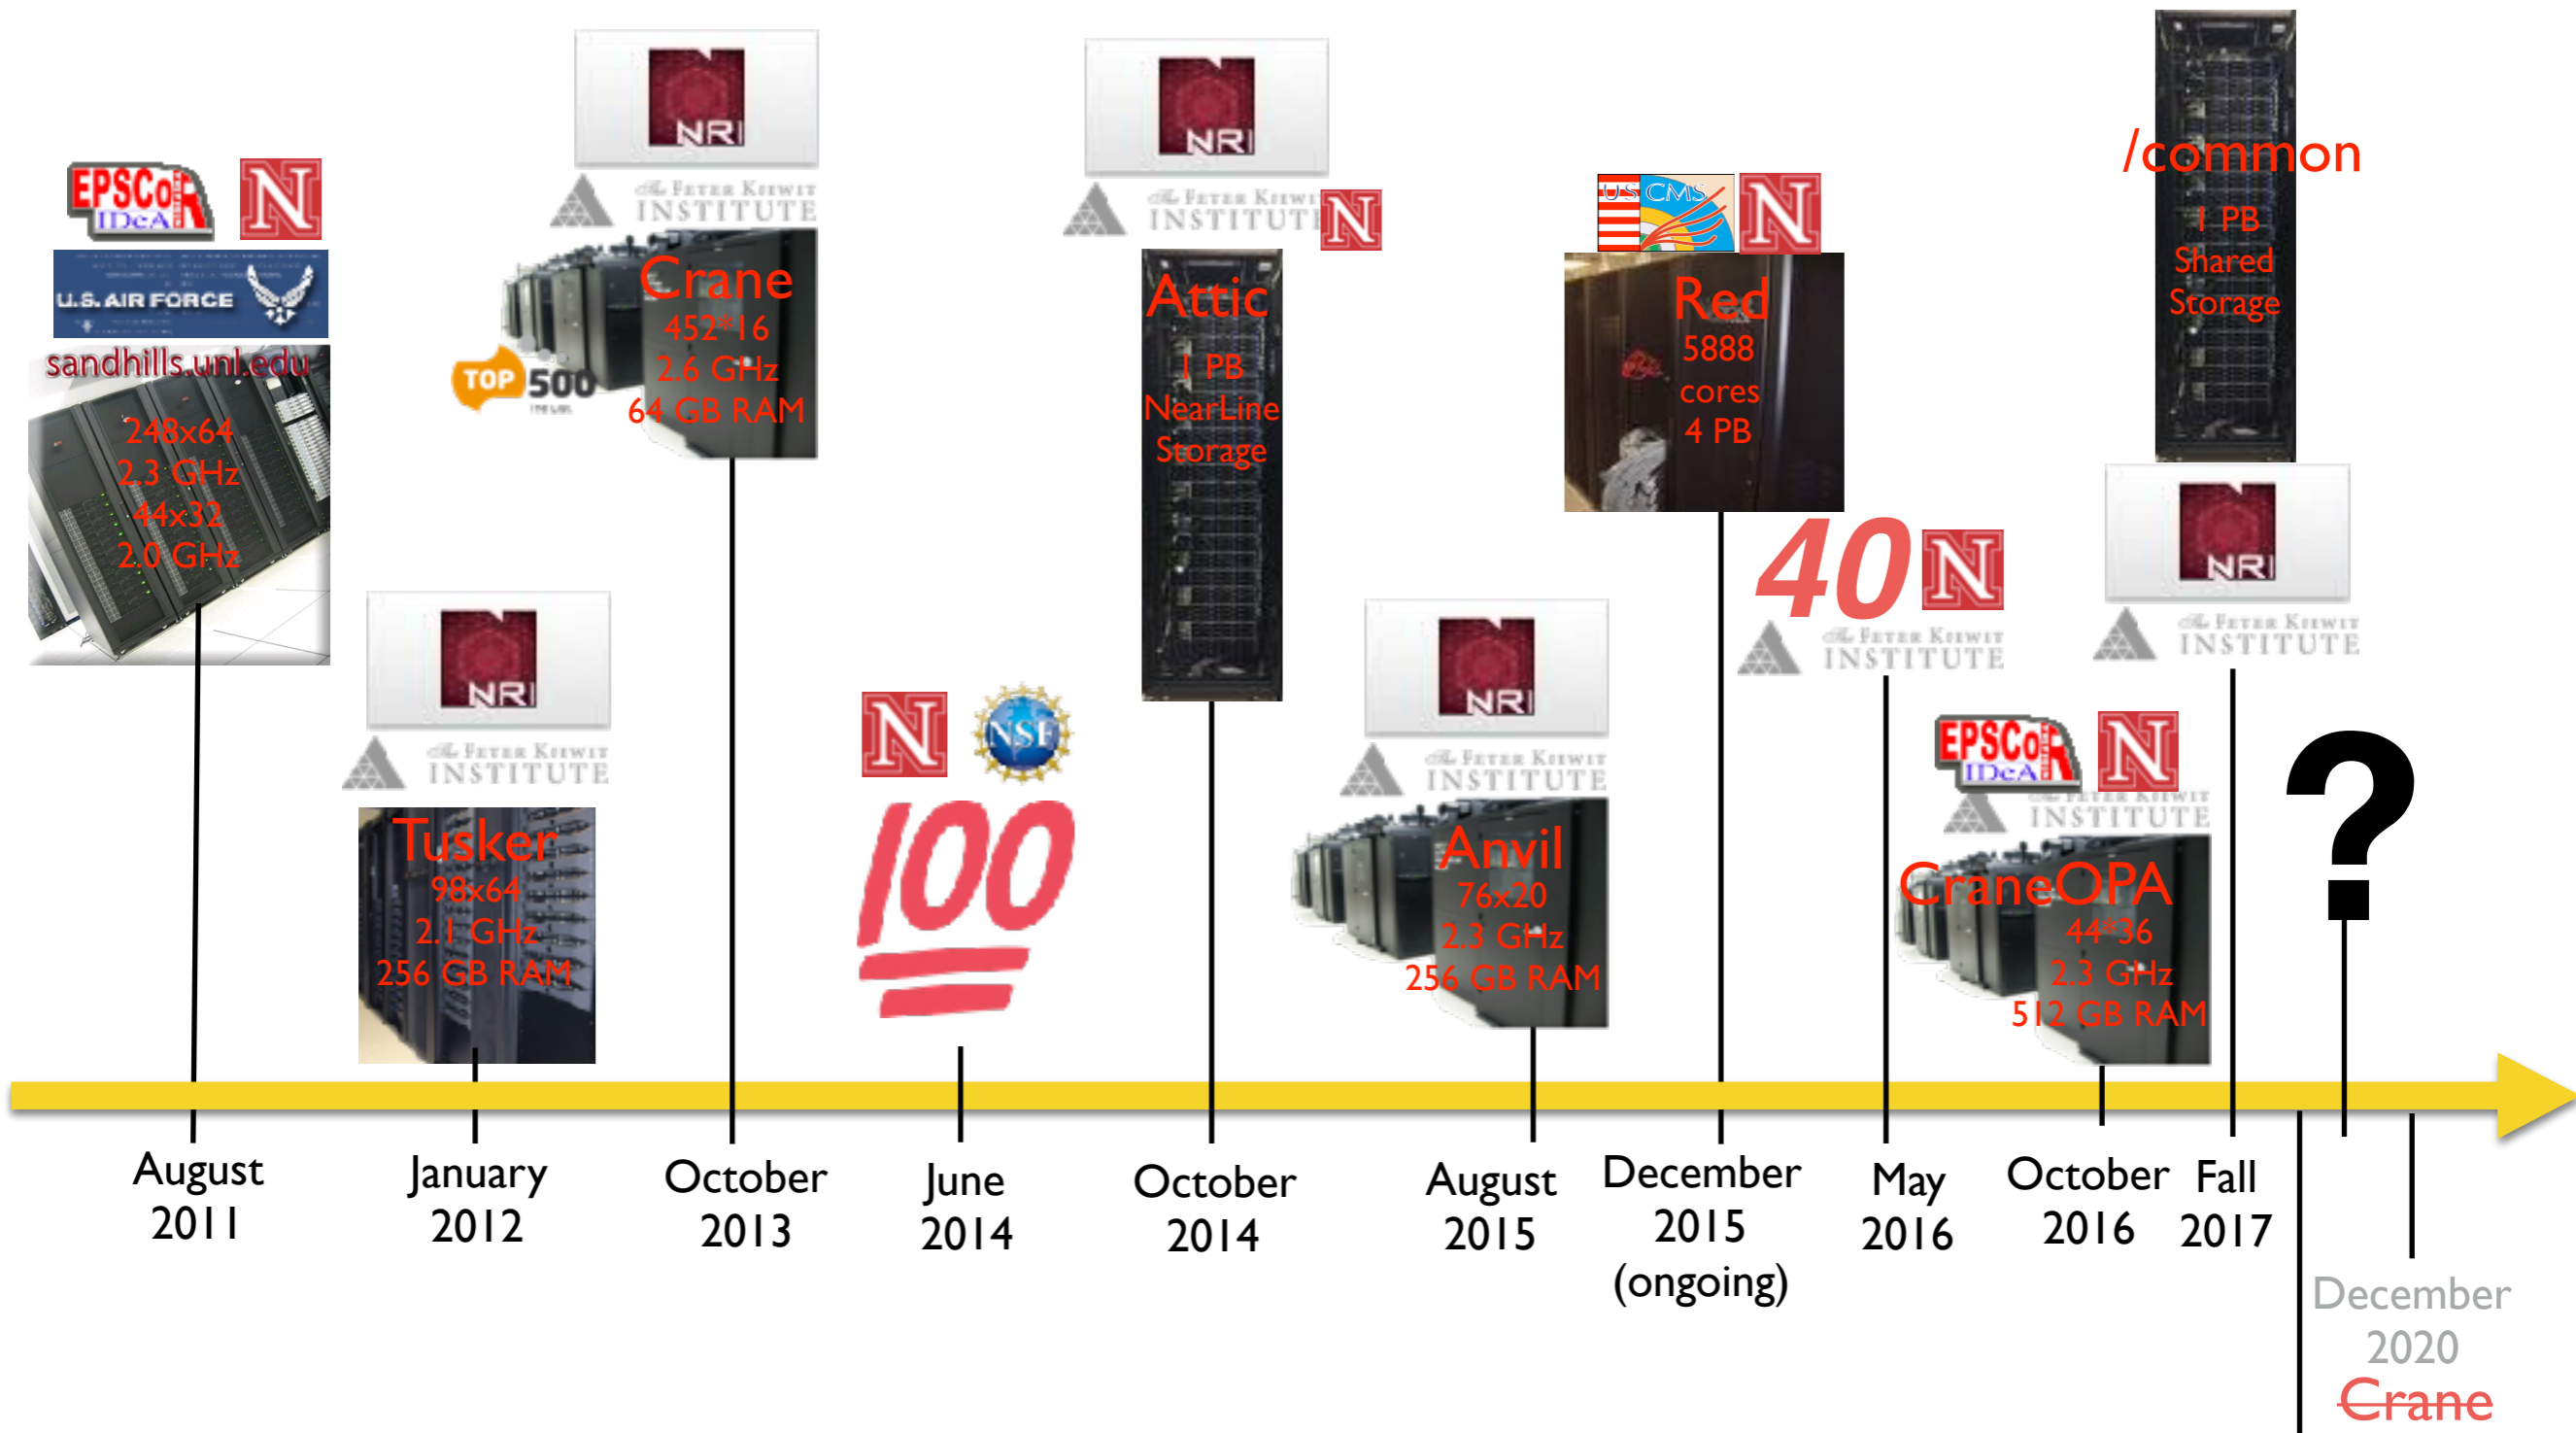
Enough about HCC ...

- What is your area of study?
- What question do you wish you could answer?
- Do you currently use HCC resources?
- What type of computing (storage/networking/visualization) do you do now?
- Are you experiencing any bottlenecks?

Parallelism Rules



Equipment Timeline





Equipment Timeline

January 2019
Tusker
Sandhills

More to come...

- HCC is buying incrementally, annually
 - MRI or other grant funding could enhance this
 - What should we target?
- Establishing an annual process
 - Gather input in fall
 - purchase in spring
 - add incrementally until next spring




It takes
ca. 6 days
to transfer
6,800 TB
at 100 Gbps

**Science is a Team
Sport**

LHC and CMS

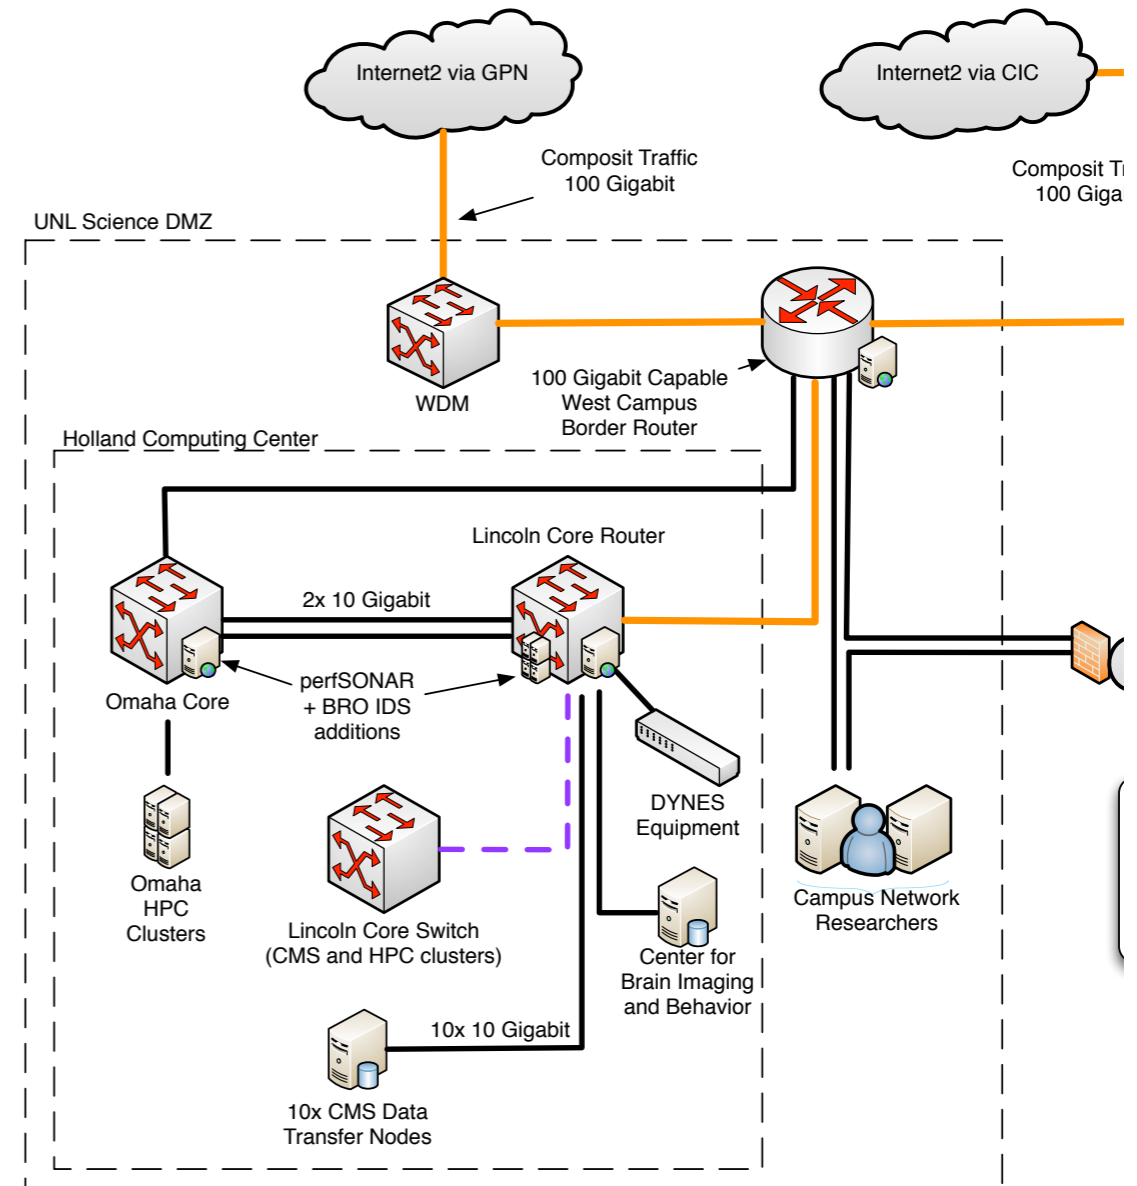
- In Run 1 the LHC created 15 PB of data/year
- In Run 2 expected output is 30 PB/year (1 GB/sec)

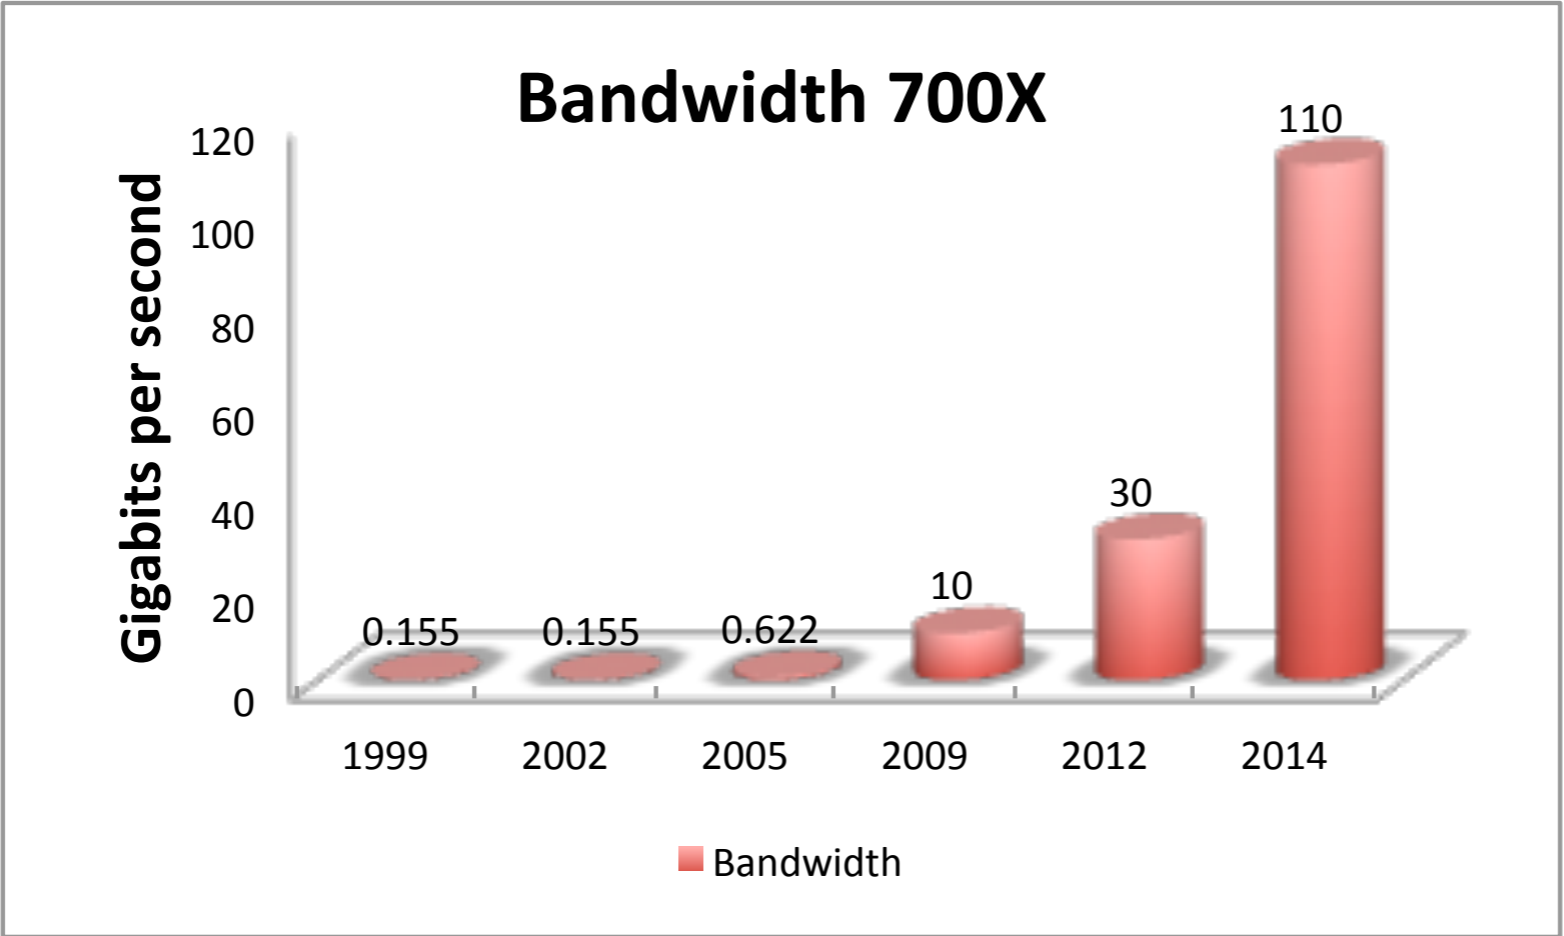
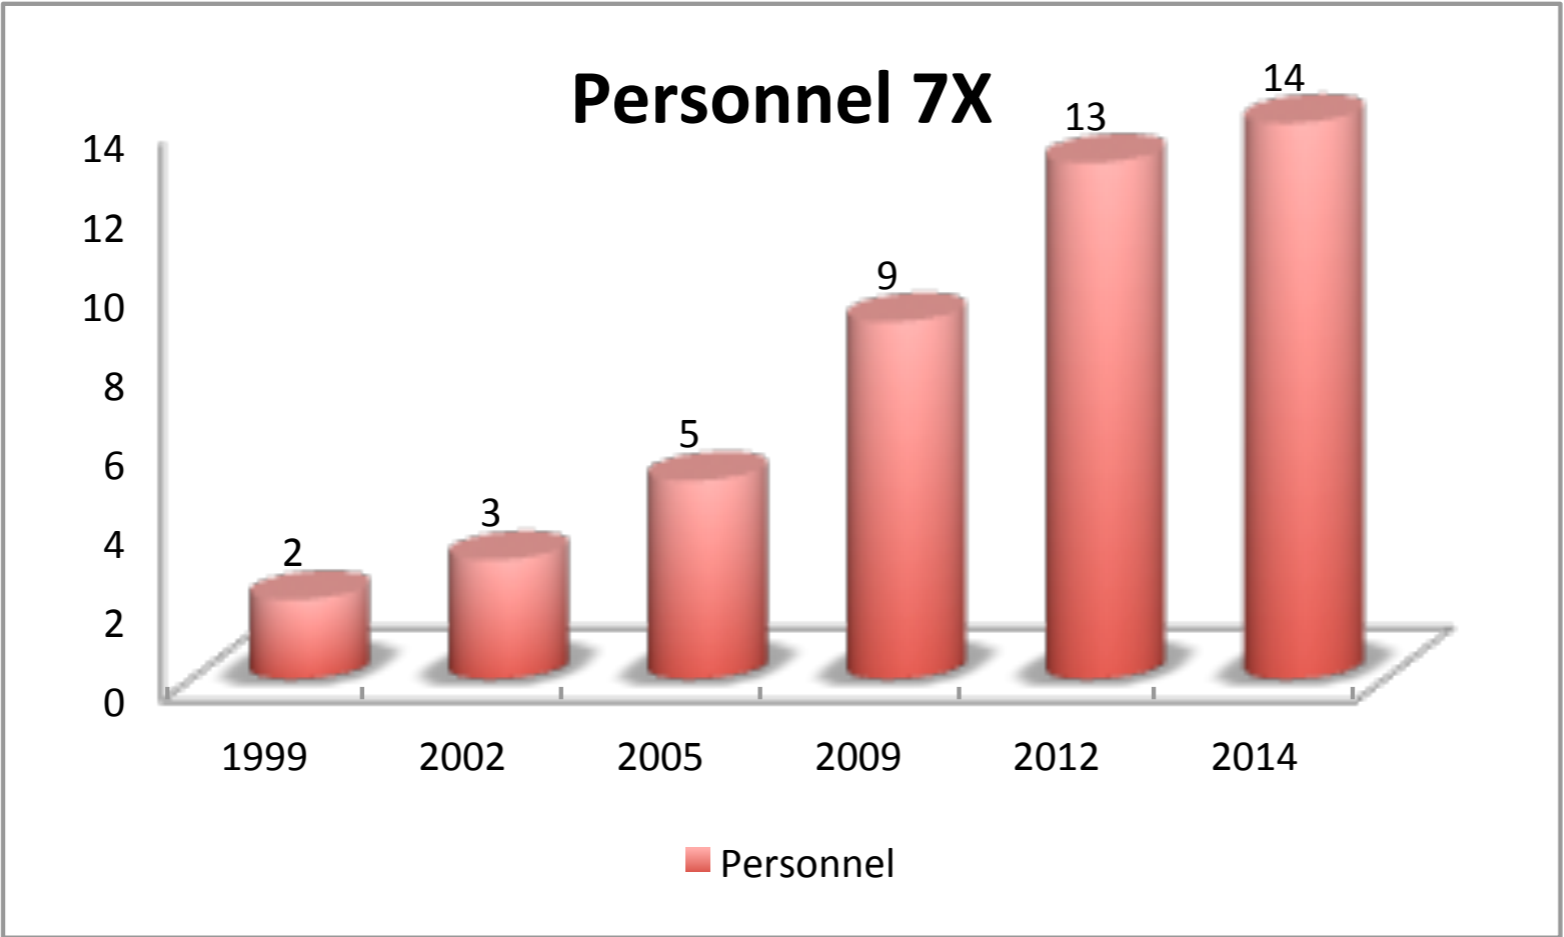





Collaboration Infrastructure

- Network Nebraska
- 100 Gbps between HCC and Internet2 via GPN
- DMZ to/from HCC
- OSG access to HCC resources
- DTNs, Globus Connect







Local speed record to date
42 (+10!) Gbps

- Software Defined Networks
- PerfSONAR
- Performance Testing

Getting Started