Group: Information Technology Council 2021-2022
Subject: Minutes for the ITC Monthly Meeting
Date: Monday, February 21st, 2022, 9:00-10:00am
Location: Zoom

Recorder: Aaron Biggs
Attendees: Andy Fagg (Computer Science), Aaron Biggs (Provost Office), Nick Key (IT), Ross Mehl (Extended Campus), Richard Sprecker (Drama), Eric Zemke (Libraries), April Dickson (IT), David Horton (IT), Rin Ferraro (Graduate Student Senate), Aaron Baillio (IT Security), Henry Neeman (OSCER)

ITC - Meeting - February 21, 2022

Research Computing - Henry Neeman

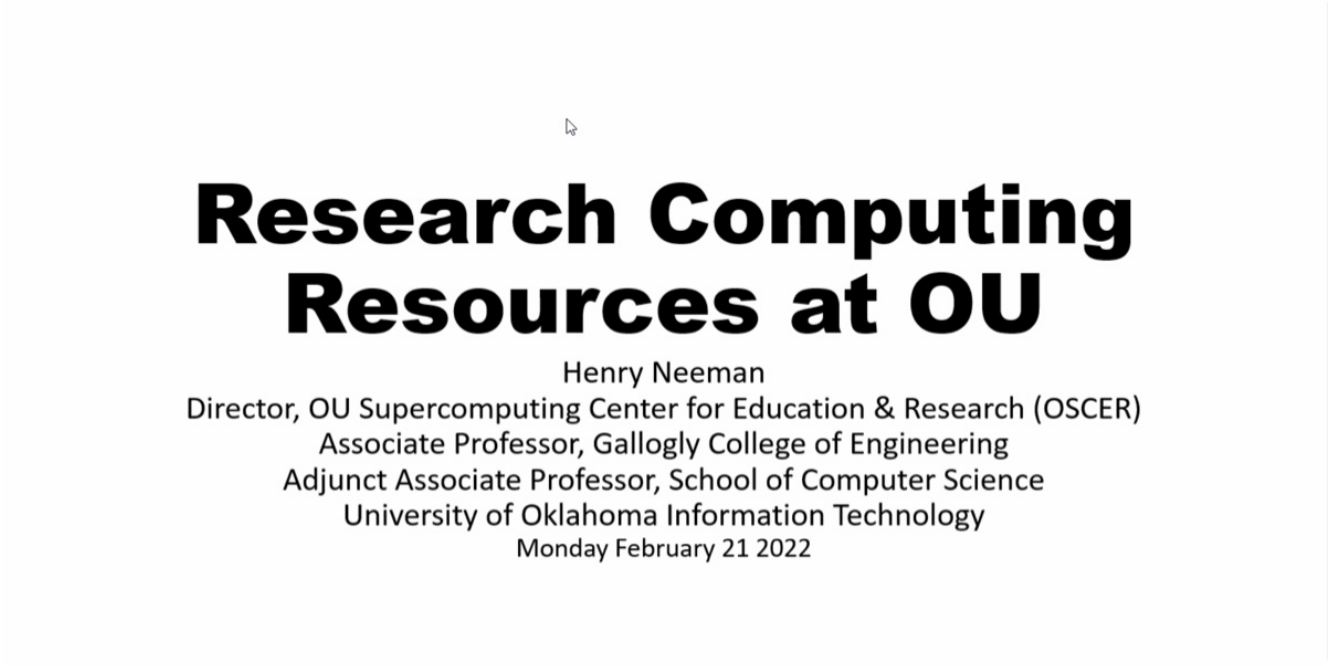

- In this discussion, "Research Computing" is shorthand, because it applies to all of research, scholarship and creative activity.

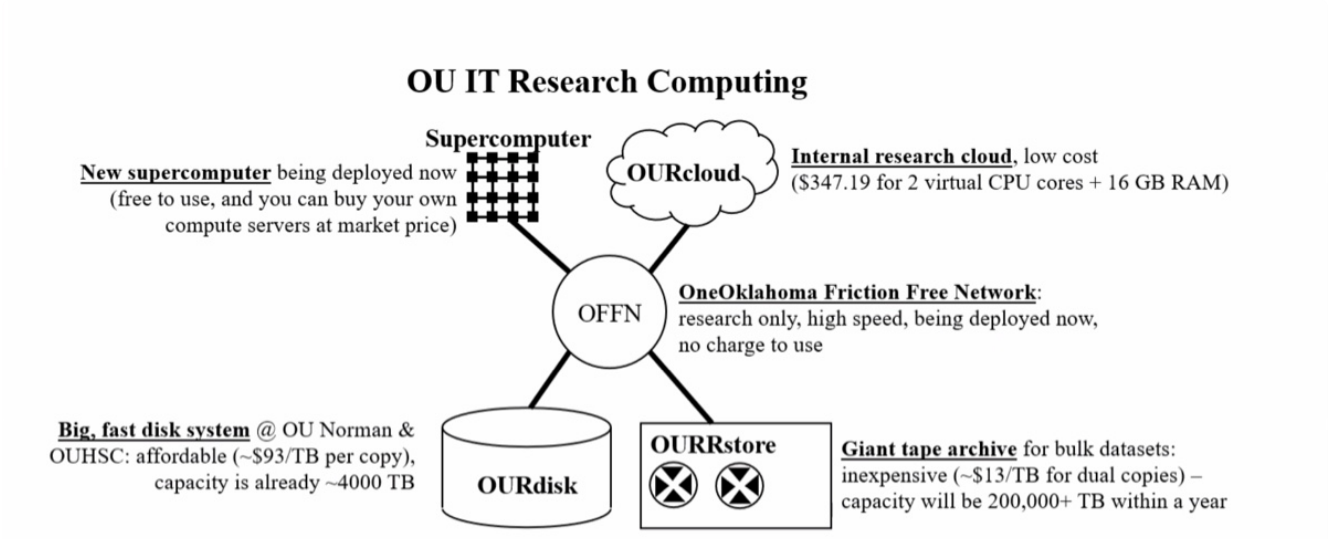

- OU IT Research Computing Summary

    - OSCER = OU Supercomputing Center for Education & Research, an OU IT team
    - Supercomputer Refresh
        - We've purchased network, storage and support servers, which are being configured now.
        - We'll purchase a bunch of compute servers in Jan 2022, to be deployed by fall (late spring?).
        - Shortly will include 33 NVIDIA A100 GPU cards for AI/Machine Learning/Deep Learning.
        - Already 6 OSCER-owned A100 GPU cards, +6 coming; already 11 user-owned, +10 coming
        - "Friendly user" mode expected in summer 2022 our old supercomputer will be active until then.
            - “Friendly user" is the exact opposite of "user friendly."
    - OU Research Cloud (OURcloud): New supercomputer includes hardware/software refresh.
        - ~2 TB RAM, ~250 virtual CPU cores (@ 3:1 oversubscription)
        - Friendly user mode NOW
    - OU Research Disk (OURdisk)
        - Friendly user mode NOW (more features coming).
        - Available NOW locally at OU Norman, soon at OUHSC.
    - OURRstore Tape Archive
        - Friendly user mode NOW (but valuable features still to come).
        - Developing control scripts, reporting scripts, etc now.
    - HIPAA Enclaves for all of these: Planning discussions with OU IT Governance, Risk and Compliance team underway; HIPAA enclaves expected in late calendar 2022.
- HIPAA Etc: New Legally Regulated Enclaves
    - Every OU IT Research Computing resource is going to be split into a legally regulated enclave (HIPAA, CUI, etc) and a non-regulated enclave.
    - Status: We already meet regularly (every other week) with OU IT's Governance, Risk and Compliance team, and OU IT Information Security frequently.
        - We already have a preliminary plan for the HIPAA supercomputer.
    - ETA: Late 2022.
- Current Supercomputer (Schooner) #1
    - Peak speed: 863 TFLOPs* [@ max turbo/boost]
    - *FLOPs: trillion calculations per second
    - 732 compute nodes
    - 1472 CPU chips: Intel Xeon Haswell, Broadwell, Skylake, Cascade Lake, Ice Lake, Rome, Knights Landing, Sandy Bridge
    - 18,356 CPU cores
    - 52 TB RAM
    - 500+ TB global public disk
    - ~3 PB global condominium disk
    - Mellanox Infiniband FDR10 & HDR100

    - ~ 1 microsec latency)
    - Dell N-series Gigabit, S-series 25G/100G Ethernet
    - RHEL 8.3
    - ~50% of the nodes are "condominium" (owned by individual research teams).
    - Status
        - https://ou.edu/oscer/resources/hpc
        - Support Hardware: Network switches, storage, support servers being configured now.
        - Storage: /home and /scratch will be an internal Ceph parallel filesystem, like OURdisk; OURdisk is already available on our old supercomputer, Schooner; flash disk for "burst buffer."
        - Schooner Compute Hardware: Most of the old supercomputer (Schooner) compute nodes will move to the new supercomputer, and new compute nodes are being bought starting this month.
        - New Compute Hardware: Ordered in Jan 2022, expected Apr 2022: 28 compute nodes.
            - Mix of Intel Xeon Ice Lake and AMD EPYC Rome; mostly 64-core, 128 GB RAM, some 256 GB RAM.
    - ETA: Friendly user mode in summer 2022, full production in fall 2022.
        - Meanwhile, the old supercomputer is still in production.
        - Once we get to friendly user mode, we'll gradually shift compute servers from the old one to the new one.
    - HIPAA etc: We'll set aside a small subset of the supercomputer as a legally regulated (HIPAA, CUI etc) enclave (coming in late 2022); the rest will be for non-regulated computing.
- New Supercomputer (Sooner) #2
    - CIO-sponsored general-use compute capacity
        - We purchase new compute servers most years and add them to the supercomputer.
        - We also have all the old "Haswell/Broadwell" (c. 2015) compute servers in production for the lifetime of the new supercomputer, Sooner (through C. 2024).
    - Advantage: Cost to use: ZERO (sponsored by OU's CIO)
    - Disadvantage: Wait your turn in the batch queue.
        - Our goal: Your job waits in the queue for no longer than the runtime you've requested.
            - We can't always achieve that goal, but usually (in CY2021, 15.3% of jobs waited longer than that).
            - Median wait time is less than 25% of requested time.
        - "Fair share:" Your priority in the batch queue is a combination of:
            - (a) how long your job has been waiting in the batch queue;
            - (b) how big the job is;
            - (c) how much work you've gotten done recently (to give everyone else a chance to run too).
- New Supercomputer (Sooner) #3
    - "Condominium" Supercomputing

        - Researcher buys a compute server to go into OU IT's supercomputer at market price, and OU IT houses and maintains it at no upcharge, sponsored by the CIO.
        - Researcher can buy a condominium compute server at any time (with OU IT's help, to make sure it's compatible).
    - Advantage: On-demand access instead of waiting your turn in the batch queue.
    - Disadvantage: Costs money (market price).
        - Market costs are going up.
- New Supercomputer (Sooner) #4
    - Change in Purchasing Approach
        - No leasing => rolling year-by-year additions.
            - We used to do a "forklift upgrade" that took the old supercomputer out shortly after the new one came in.
    - Retain old (2015) supercomputer hardware in production through ~2024 (no loss of capacity in the meantime, more bang for the buck per server).
    - In FY2021, we bought very little compute, but some A100 GPUs (for AI/ML), networks, storage, support components.
    - Includes Infiniband network bridge between old FDR10 40 Gbps and new HDR 100/200 Gbps.
    - Includes Infiniband-to-Ethernet network gateway for storage, especially OURdisk (which isn't part of the supercomputer).
- AI/Machine Learning/Deep Learning
    - Supercomputer GPU Plan: GPU cards for fast AI/ML/DL & number crunching
        - Fiscal Year 2021: 15 NVIDIA A100 GPU cards (4 OSCER, 11 user)
            - OSCER bought 4 A100 GPU cards: 2 × 2-GPU servers, no NVlink.
            - Users bought 11 A100 cards: 2 x 4-GPU servers, 1 × 2-GPU server, 1 × 1-GPU server, most with NVlink (600 GB/sec GPU-to-GPU inside the same server).
    - Fiscal Year 2022: New total will be 33 A100s (12 OSCER, 21 user)
        - OSCER: 8 new A100 GPU cards: 6 A100s with NVlink and 2 A100s without, in 2-GPU servers.
        - Users: 10 new A100 GPU cards (2 x 4-GPU servers, 1 x 2-GPU server), all with NVlink.
    - Fiscal Years 2023 and beyond: Each fiscal year, we'll add more GPU cards, some with NVlink, some without, some bought by OSCER, some by users.
    - NSF Campus Cyberinfrastructure proposal: 7 × 4-GPU servers (28 A100 GPUs)
        - submitted Oct 2021, decision expected March/April 2022.
            - 69 research groups: 18 groups already using GPUs, 51 groups plan
- OU Research Cloud (OURcloud)
    - Purpose: Interactive, web services, databases (e.g., SQL), etc.
        - http://www.oscer.ou.edu/ourcloud
    - Researcher's Price: $347.19 per portion (minimum buy-in)

        - Portion: 16 GB RAM, 2 virtual CPU cores (2/3 physical CPU core 3:1 oversubscribed), good for 7 years
        - Least expensive research cloud offering in OU IT history!
        - AWS costs 3.7 times as much for 3-year reserved, 1.7 times as much for spot pricing!
        - Same cost as buying a 16 GB RAM DIMM stick!
    - Size: Initially 2+ TB RAM, 250+ virtual CPU cores - will grow as needed.
    - OS Options: Linux (many versions), Windows Datacenter (most recent).
    - ETA: Friendly user mode NOW, full production fall 2022.
        - Ordering webpage: https://itsupport.ou.edu/TDClient/30/Unified/Requests/ServiceDet?ID=291
    - HIPAA etc: Legally regulated enclave (HIPAA, CUI, etc) coming in late 2022.
- OU Research Disk (OURdisk)
    - Purpose: Persistant, dedicated disk space on supercomputer, OURcloud, other servers across OU
    - Researcher's Price: $860.03 per portion (minimum buy-in) = ~$93 per usable TB, good for 5 years
        - Least expensive research disk offering in OU IT history!
        - Ordering webpage: https://itsupport.ou.edu/TDClient/35/Norman/Requests/ServiceDet?ID=258&SIDs=1600
    - Speed: 14+ GB/sec aggregate: Fastest research spinning disk offering in OU IT history!
        - An individual sequential write is ~0.7 GB/sec.
    - Size: Initially 3.8 PB @ OU Norman, 3.8 PB @ OUHSC - each will grow as needed.
        - Already requested: ~2.5 PB (~66%)
    - ETA: Friendly user mode NOW, full production spring 2022.
    - Where Available
        - Supercomputer
        - OURcloud
        - Meteorology servers in the National Weather Center
        - Can be mounted on other OU IT systems and non-IT systems on any OU campus.
            - Inside Norman 4PP and OUHSC data centers: very fast (14+ GB/sec).
            - Outside Norman 4PP and OUHSC data centers: much slower.
    - HIPAA etc: Each campus will have an open enclave and a legally regulated enclave (HIPAA, CUI, etc, coming in late 2022).
- New OURRstore Tape Archive #1
    - OU Regional & Research Store: Giant robotic tape archive
        - Business Model: NSF MRI buys HW/SW, researchers buy tapes, CIO covers space/power/cooling/network/labor/maint (same biz model as PetaStore).
            - Currently ~$50 per LTO-7 "Type M" tape cartridge (~7.65 TB usable) => ~$13 per usable TB dual copies
        - Tape Cartridge Slots: Initially ~11,000, will expand to to ~22,000 (11K @ Norman, 11K @ OUHSC)

        - Tape Drives: Initially, ~1.8 GB/sec in aggregate (6 × LTO-8 @ 300 MB/sec/drive) - almost double!
            - LTO-8 drives can read and write LTO-7 "Type M" and LTO-8.
            - Later, we'll add 4 x LTO-9 or LTO-10 tape drives (can read and write LTO-8/9/10?): 3+ GB/sec total.
        - Disk: ~450 TB usable disk front end "landing pad."
        - Resiliency: Secondary copies are exported from OURRstore at OUHSC, shelved.
    - ETA: Friendly user mode NOW, full production spring 2022.
    - HIPAA etc: We'll set aside a small subset of OURRstore as a legally regulated (HIPAA) enclave (coming in late 2022); the rest will be for non-regulated data.
        - 2640 tape cartridge slots, 2 LTO-8 tape drives, 100+ TB disk
        - HIPAA enclave is funded by OU IT, not the NSF grant - not available outside OU.
- New OURRstore Tape Archive #2
    - NEW FEATURES (compared to the soon-to-be-decommissioned PetaStore)
        - Auto-Archiving: User places files in a specific directory, daemon" process archives those files automatically.
    - File Sharing (via Globus license): With a few clicks, a file owner can designate a file to be downloadable by (a) a specific user, (b) a specific group or (c) the whole world. (Files are private by default.)
    - Caching: Files reside on the disk front end until they're the least recently used and need to be cleared out to make room for incoming files popular files are on both disk and tape, unpopular files are on tape only.
    - Disk Purchase: Buy disk that files can live on permanently, for fast downloading.
    - Researcher's Price: LTO-7 "Type M" is 3.6 times bigger, half the cost per TB of LTO-6 on the PetaStore (and we'll add LTO-8, LTO-9 etc as they become viable in $/TB).
        - ~$50 per LTO-7 tape cartridge (~7.65 TB usable) => ~$13 per usable TB for dual copies, plus IDC
    - Working with library to use OKShare and Globuls together. https://libraries.ou.edu/content/upload-your-work-shareok
- OneOklahoma Friction Free Network
    - OneOklahoma Friction Free Network (OFFN)
        - Researcher's Price: ZERO (sponsored by OU's CIO)
        - High speed network for open research only
            - Friction Free: bypasses firewalls because the data is open
                - This allows much higher speed.
        - Funded by 6 National Science Foundation Campus Cyberinfrastructure grants.
        - OU's open research connection to:
            - Other Oklahoma research institutions
                - 20 campuses of 18 institutions: PhD-granting, masters-granting, bachelors-granting, community colleges

        - Research institutions across the US
    - ETA: Friendly user NOW, full production spring 2022.
- Andy: For protected class data, is there a physical separation?
    - Henry: No, we able to segment virtually. Protected super computer is currently separated from open super computer, but that will probably chance to a virtual separation in the future.
- Andy: How will the handoff to new super computer be conducted?
    - Henry: Ramp down old one and turn on new one. Effectively, as the new super computer is moved to friendly mode, we will shift them virtually from the old to new super computer (one rack at a time). As for the user experience, they won't notice much of a change on the old super computer until more resource move to new super computer.
    - Henry: At the request of export control, everyone will have to request a new account to access the new super computer.
        - User directories will be automatically copied over.

Research Computing Survey - Henry Neeman & Nick Key

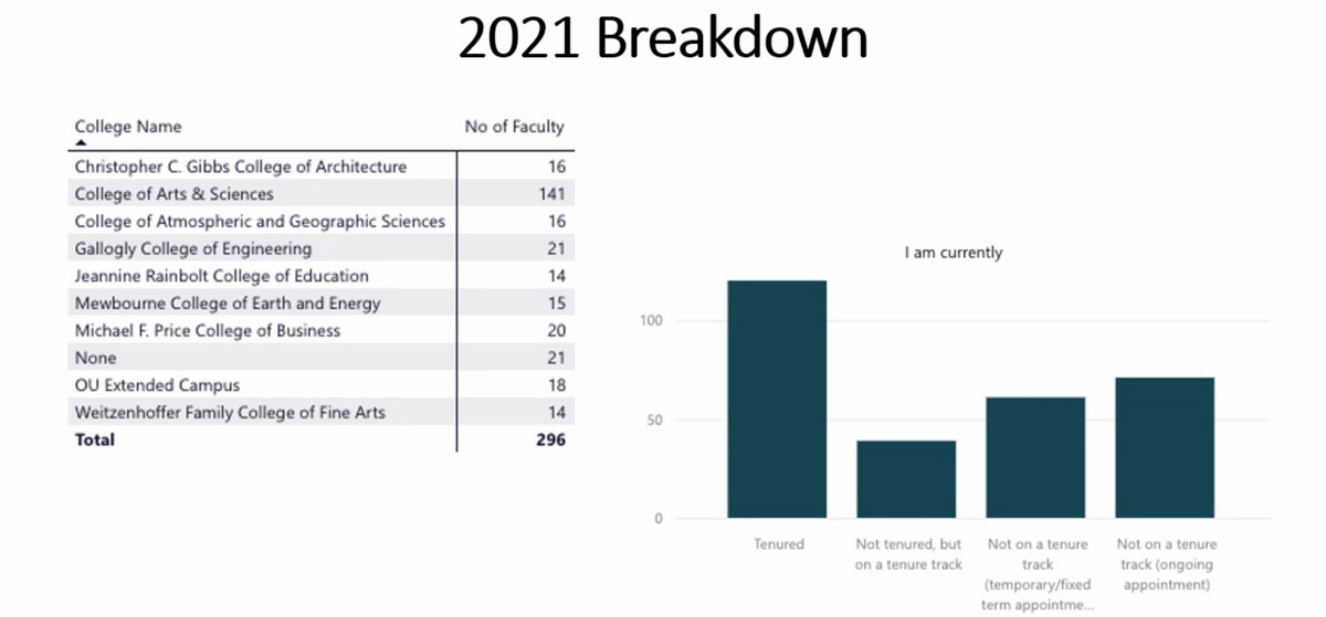

Computational Needs Not Being Met?

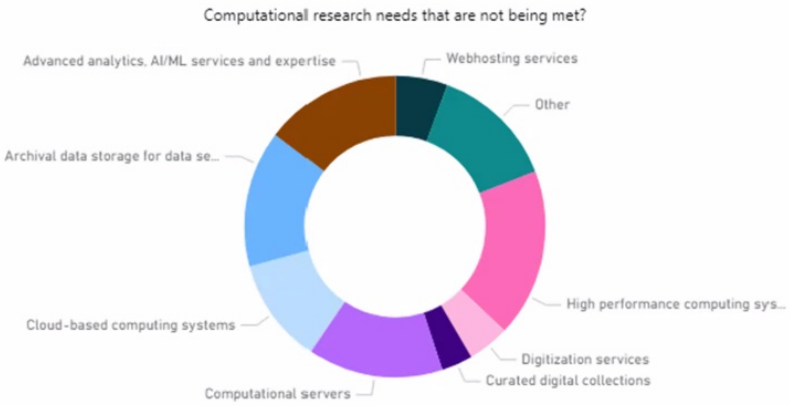

- Many of the new services that are coming online will address many of theses needs.

Computational Needs Not Being Met?

- Cloud
  - Want to buy cloud services.
  - Need workshops on using cloud services.
- Storage
  - Need storage options.
  - Central storage is hard because of firewalls.
- Purchase Issues
  - Need purchase model for storage and compute to match research funding for equipment purchases [but can buy via OSCER].
  - Need policies that allow buying servers.

Need Significant Computational Services Beyond PC?

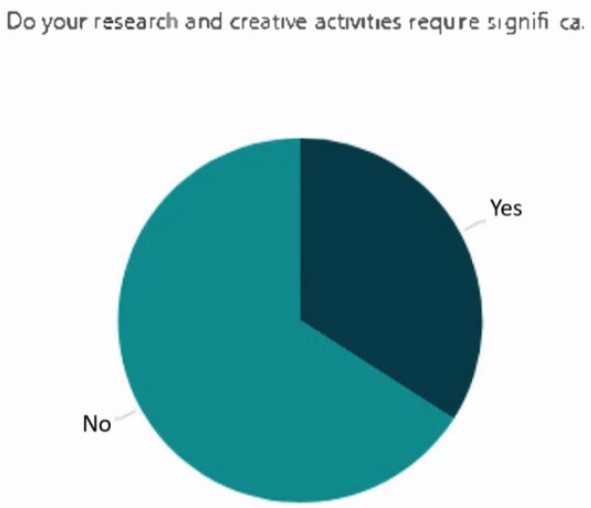

Need Significant Computational Services Beyond PC?

- Self-support instead of using central resources.
- Virtual server hosted by VPRP Office [which is going away, but OSCER can help].

How Are Your Needs Being Supported?

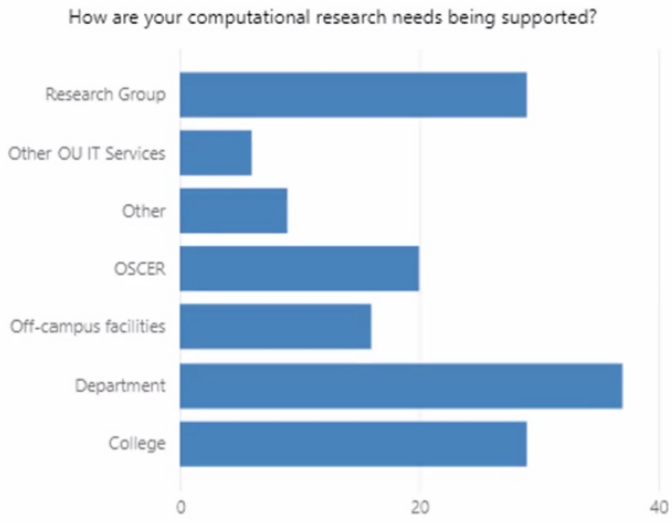

How Are Your Needs Being Supported?

- Need more help from OU for technical problems.
- Need flexible tech purchase policy.
- Need training, help and documentation on a broad variety of tech topics.
- Document sharing
- OU-wide digital asset management system
- Filesharing [coming soon]
- Support for HIPAA data [coming soon]

Benefit from Access to Computational Consulting?

How well does OU/IT meet your research computing needs?

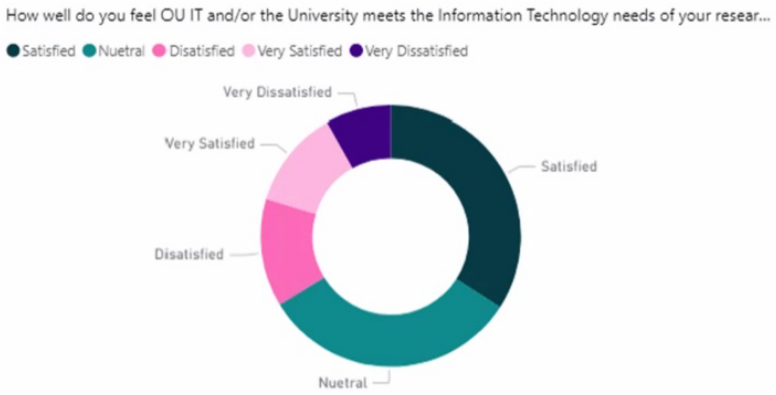

Services you need?

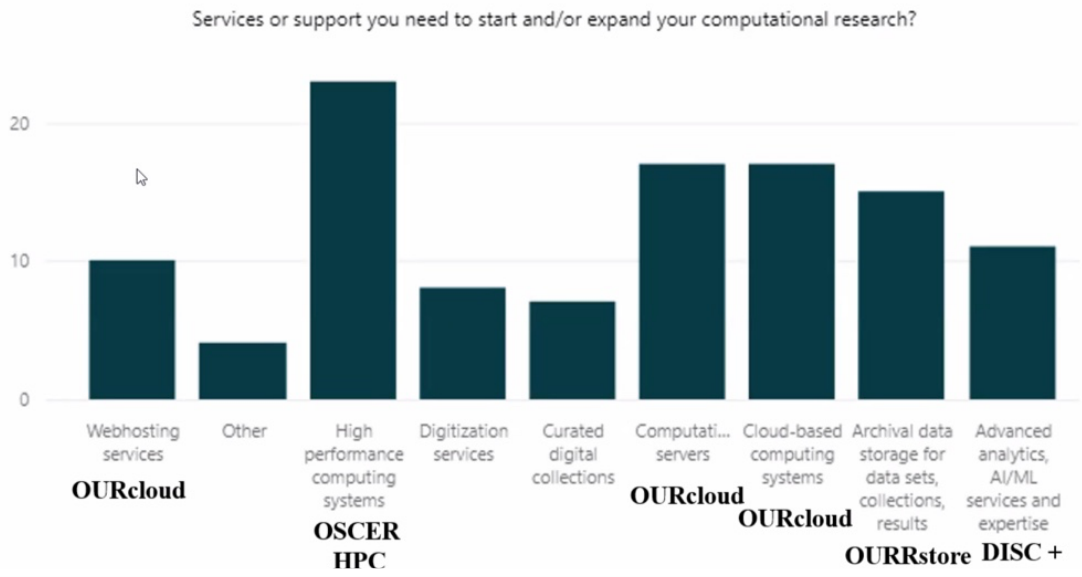

Which resources do you use or plan to use?

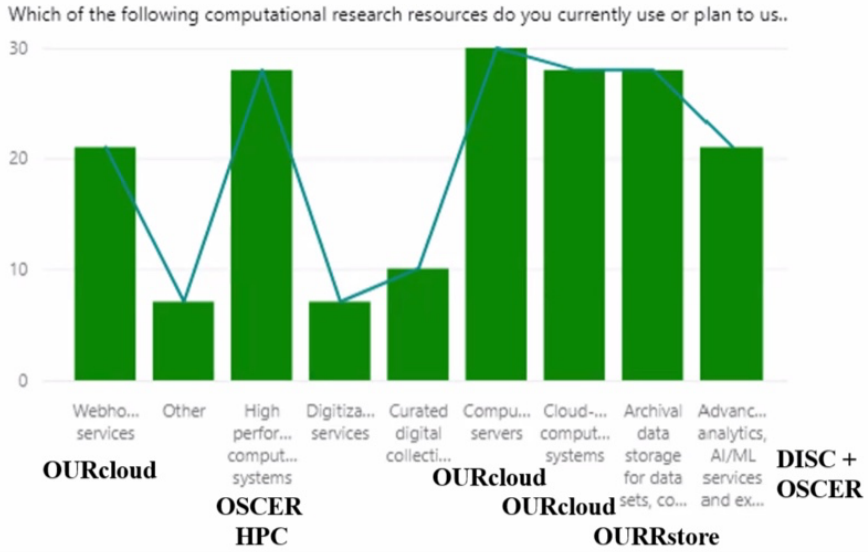

Priority by Importance?

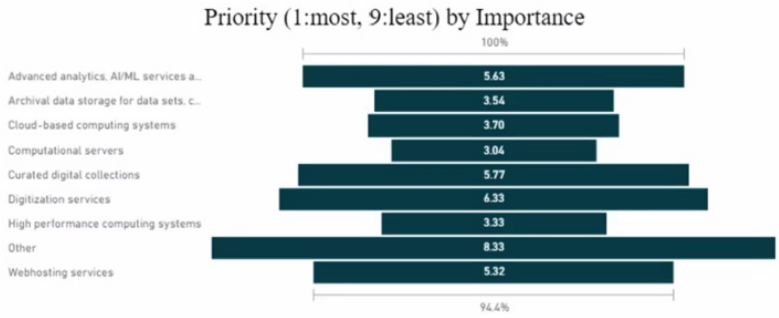

Benefit from access to hireable data scientists?

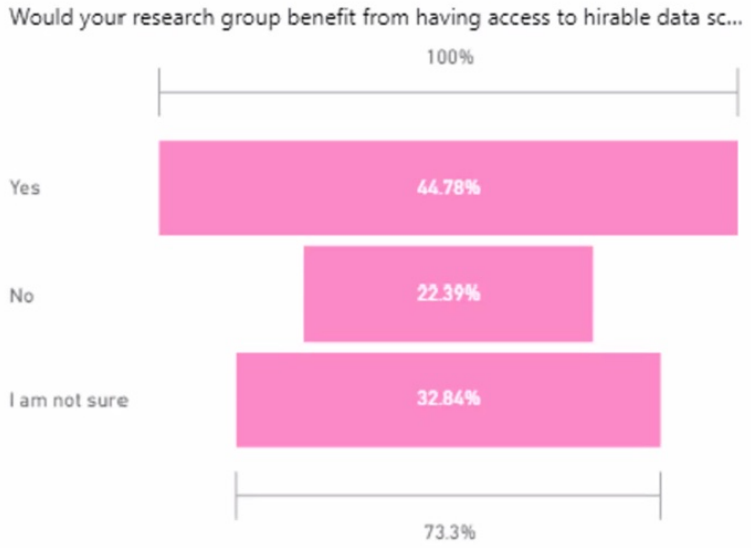

General computing/ training resources needed?

- Lightboards
- Web development support
- Linux system administration support
- GPU hardware support
- Bibliographic applications training
- Statistics packages
- NVivo [and other qualitative analysis capabilities]
- Programming in several different languages
- One-on-one consulting
- Digital imaging

IT Research Support?

- Help with Zoom
- Help with commercial cloud
- Security policies that align with research methods
- Better support for qualitative research
- Tech purchase process
- Digital humanities
- More IT staff to help with research
- More software offerings

David: Security policy align with research methods. This is Challenge. We need research methods to give some especially around protected research. There are changing exceptions from sponsors. This dominates BIG 12 CIO conversations. Expect more regulations even with FERPA data.

- Andy: We don't want to add a lot of friction with research data.
- David: We need to find a balance. We have to find solutions. We are on a journey with all research institutions.
- Henry: Much of future funding will soon be tied to security control capability.

- April: The funding agencies are not looking project-level policies, but institutional policies. Our balanced approach should be documented and auditable.
- April: IT Governance will play a critical role in shaping these policies
- Andy: These policies need to be very well defined. No black boxes.
- David: We need to define these policies collaboratively. Things are changing rapidly. We will continue to work on this.
- Andy: Should we bring in researchers to help with this?
  - David: We need to get them in front of IT Governance research committee.
  - Nick: Hopeful we can get first meeting in March. We can meet before with a tier III committee. We can start asking for some names.

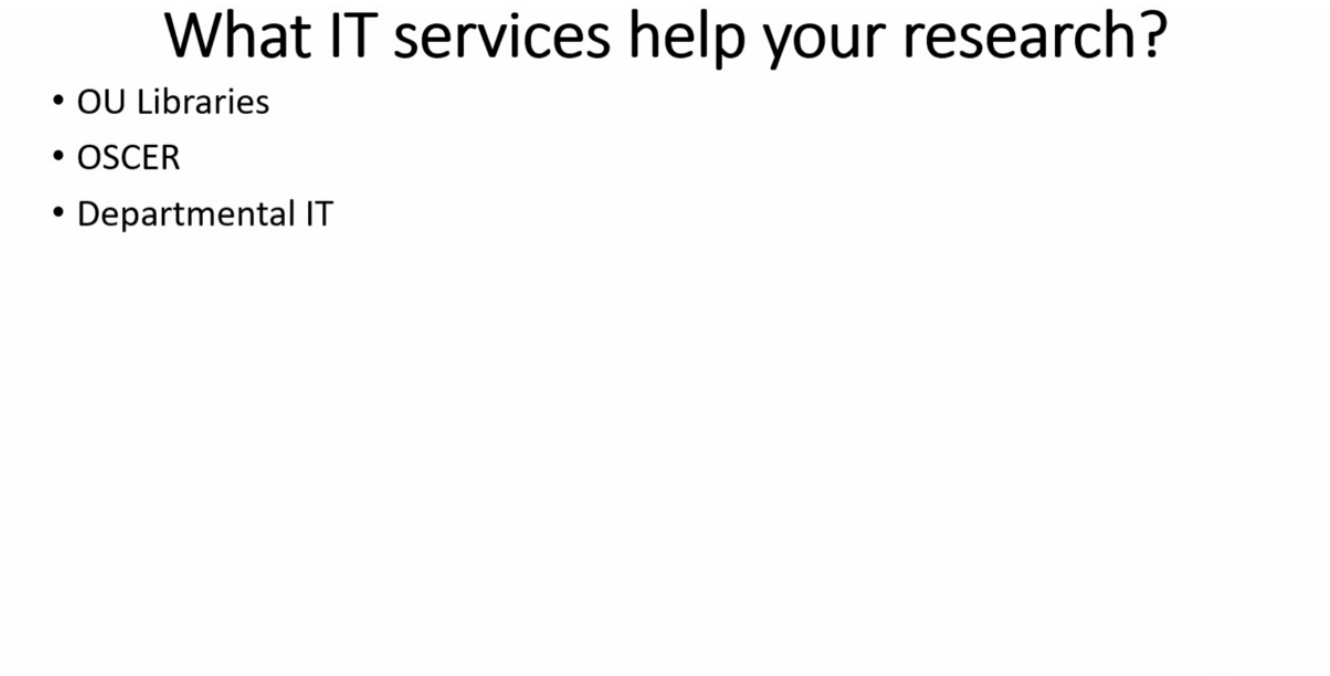

Policy Roadmap Update - Aaron Baillio

- Andy: The draft policy is entirely admin level stuff. The reality is that we have people doing real work via SSH and autonomous system. Those need to be articulated in policy.
  - Aaron: We don't' want to disrupt those use cases (autonomous logging and MFA example). We need to get an inventory of those. There should be an exception clause in the policy. Will verify.
  - Andy: How to request IT security assessments.
  - Aaron: Should just be a ticket.
  - Nick: All security assessment requests can be made through TDX (https://itsupport.ou.edu/TDClient/35/Norman/Requests/ServiceDet?ID=64&SIDs=626).