UT Southwestern Medical Center
Lyda Hill Department of Bioinformatics

BioHPC

# BioHPC Town Hall

## Current business model still roots in the beginnings of BioHPC

- Launched by 2 PIs and 3 department chairs -- 2013
- Generate an infrastructure for ***scientific computing*** for the founding entities under the principles of collegiality and a sustainable economy of scale
- Two pillars:

Liqiang Wang

Simple business plan

- Each entity contributes the hardware it needs ***on average***
- Never reduce power of infrastructure
  - ***20% annual amortization***
- Direct interaction between BioHPC team and end user
  - ***By the people for the people***
  - Entities cover FTE

***Both pillars are still in place and intact***

## But things are changing

| | | |
|---|---|---|
| 3 Departments + 2 PI | → | 25 Departments |
| 656 CPU Cores<br>2 GPU nodes<br>240TB storage | → | 28,000 CPU Cores<br>100+ GPU nodes<br>50PB storage |
| 10 - 20 ‘expert’ users | → | >1000 general users<br>~200 ‘warm’ users |
| 1 PI and his secretary | → | Department-level administration with need for a Bizz Analyst |
| Few software platforms, mostly related to development | → | Several 100 software packages, largely consumer-oriented |

## But things are changing

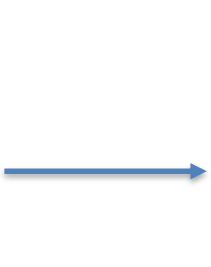

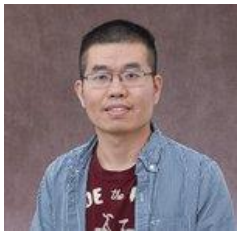

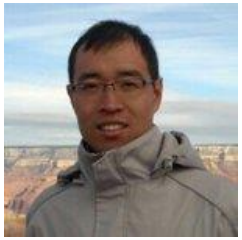

100% investigator covered → 75% Dean's Office + 3 Departments

# 1.2 Million jobs during COVID-19 pandemic

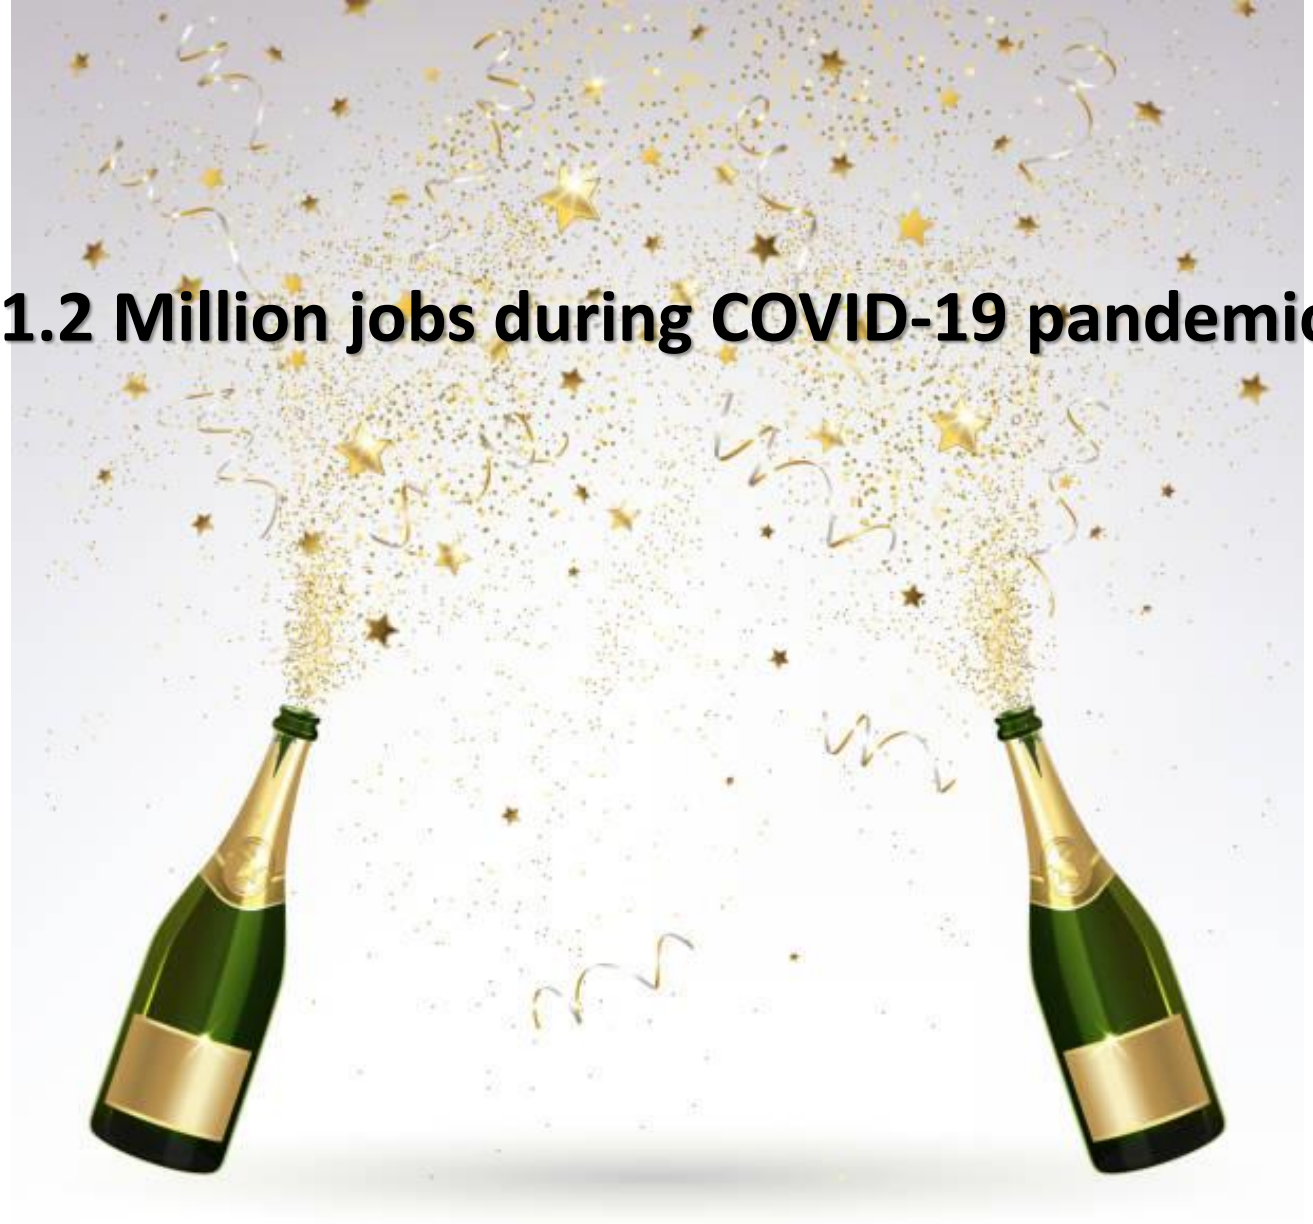

## Things we ***do not want*** to change

- User-centric / science-focused service model
  - BioHPC must remain part of the departmental structure as opposed to part of the administrative structure
  - Easy (and free of cost), direct interaction channels between end user and BioHPC team
- User-owned hardware and amortization
  - BioHPC must retain the flexibility to invest in new technology without the constraints of an institutional budget
- Focus on empowering a heterogeneous user community with scientific computing solutions
  - BioHPC must not become the institutional data storage provider

## Governance

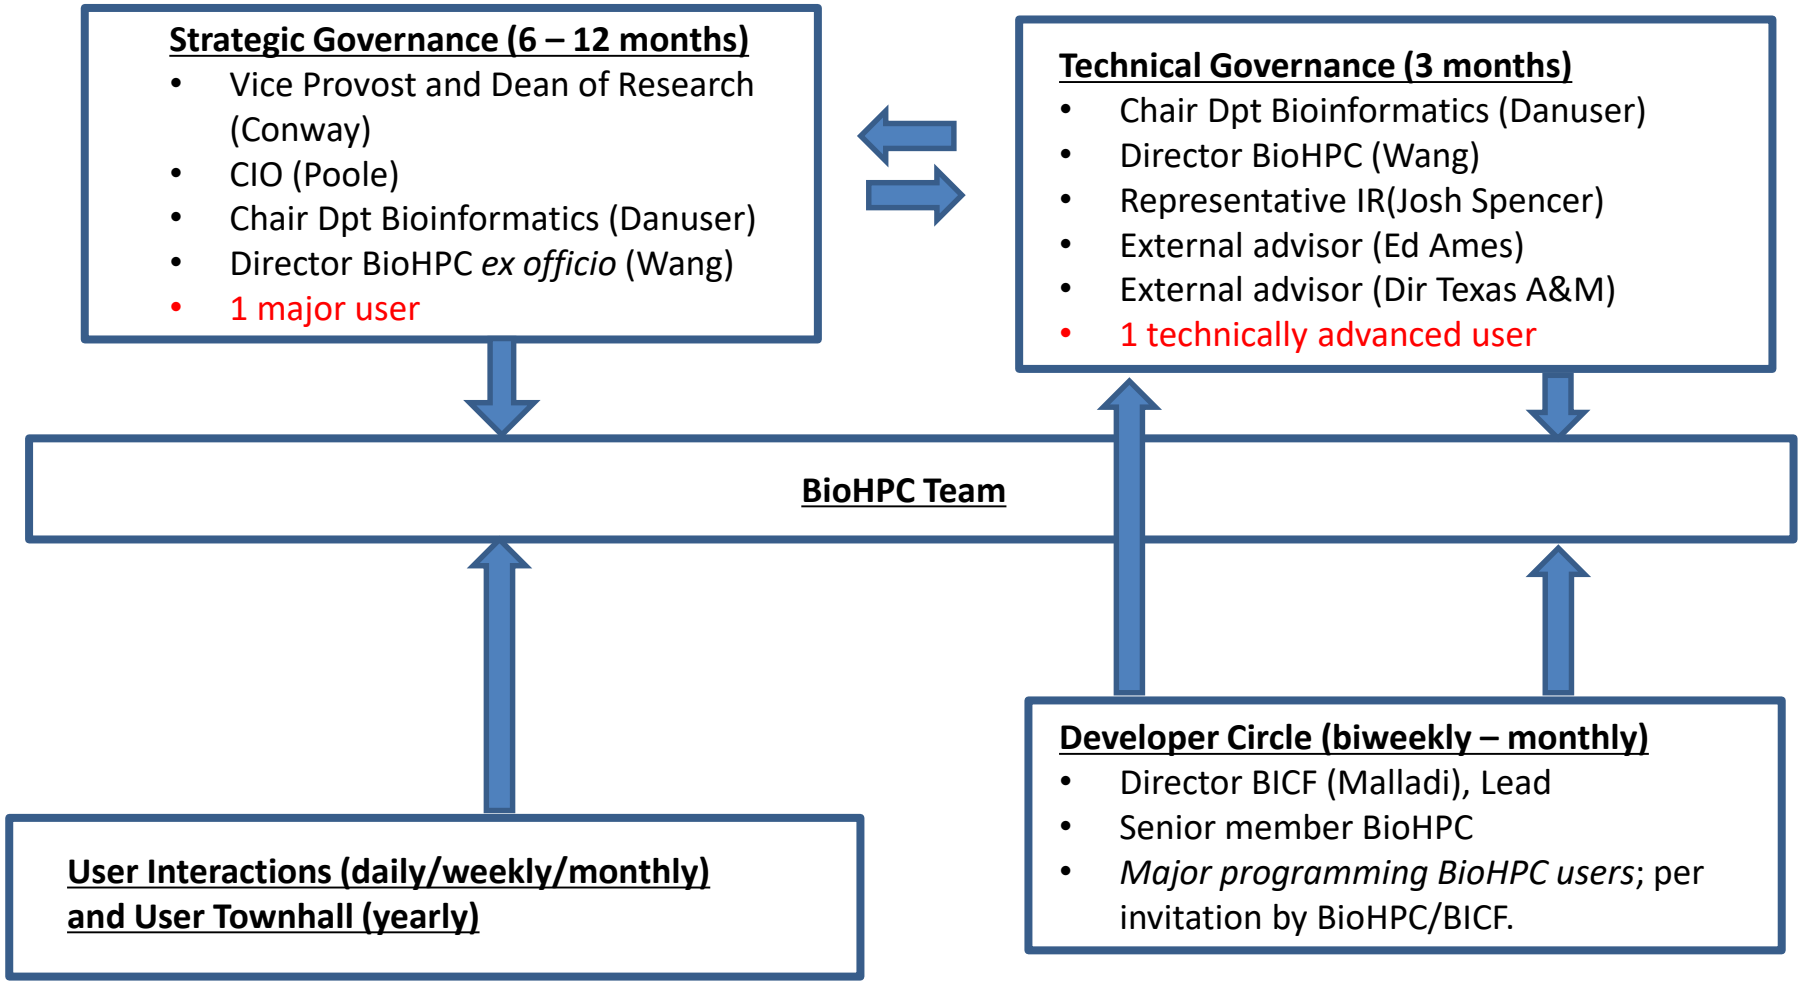
Strategic Governance (6 – 12 months)
• Vice Provost and Dean of Research (Conway)
• CIO (Poole)
• Chair Dpt Bioinformatics (Danuser)
• Director BioHPC ex officio (Wang)
• 1 major user
Technical Governance (3 months)
• Chair Dpt Bioinformatics (Danuser)
• Director BioHPC (Wang)
• Representative IR(Josh Spencer)
• External advisor (Ed Ames)
• External advisor (Dir Texas A&M)
• 1 technically advanced user
BioHPC Team
User Interactions (daily/weekly/monthly) and User Townhall (yearly)
Developer Circle (biweekly – monthly)
• Director BICF (Malladi), Lead
• Senior member BioHPC
• Major programming BioHPC users; per invitation by BioHPC/BICF.

# Data Hall moving plan

## Data Hall moving

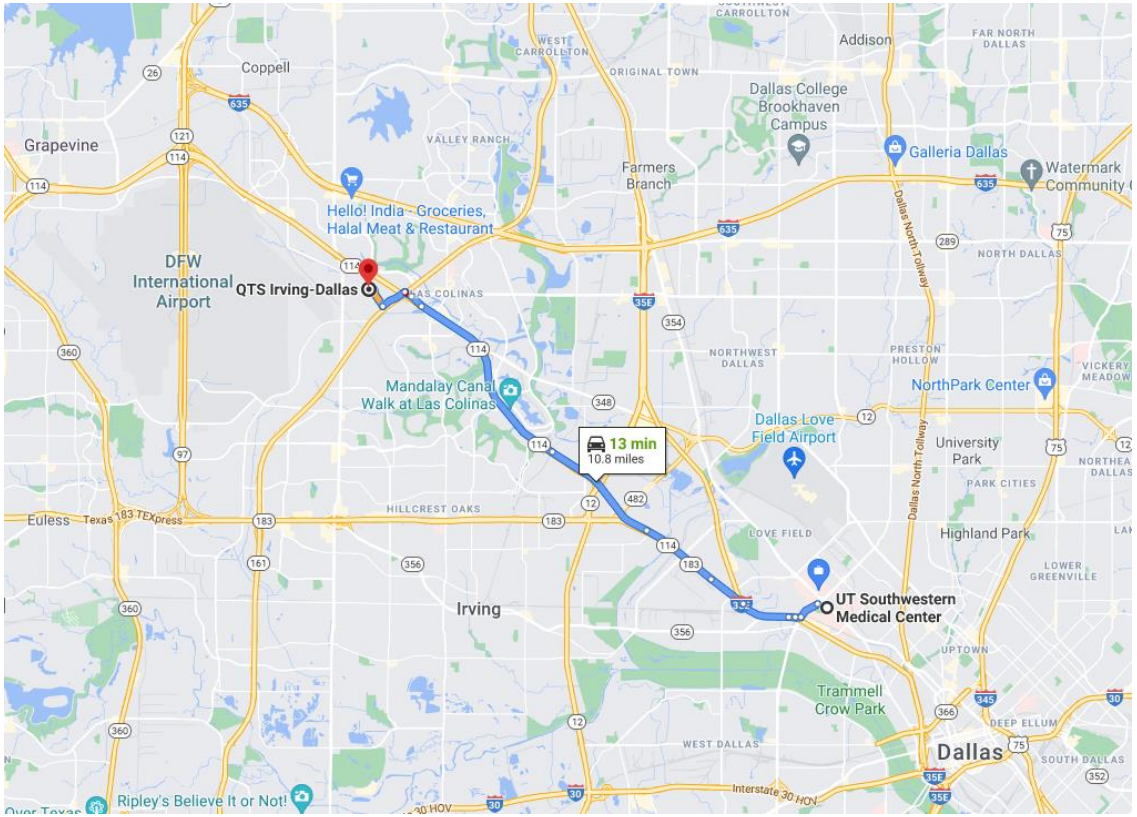


- UTSW Data Hall will move from Bass building to DFW QTS

## Storage moving plan

/home2 Read/Write

/archive Read only

3-4 weeks

/work Read/Write

/project Read only

- Backup all storage 4-5 months at least
- /home2 /work R/W
- /archive /project Read only 3 – 4 weeks

## Computer node moving plan

1, Install core network switches to new location(3.3x faster)
2, Move the Primary Head node, Secondary head node will be Primary
3, Reinstall Primary Head node
4, Move secondary Head node
5, Move compute node by partitions(7 times)

All Shipping with insured mover

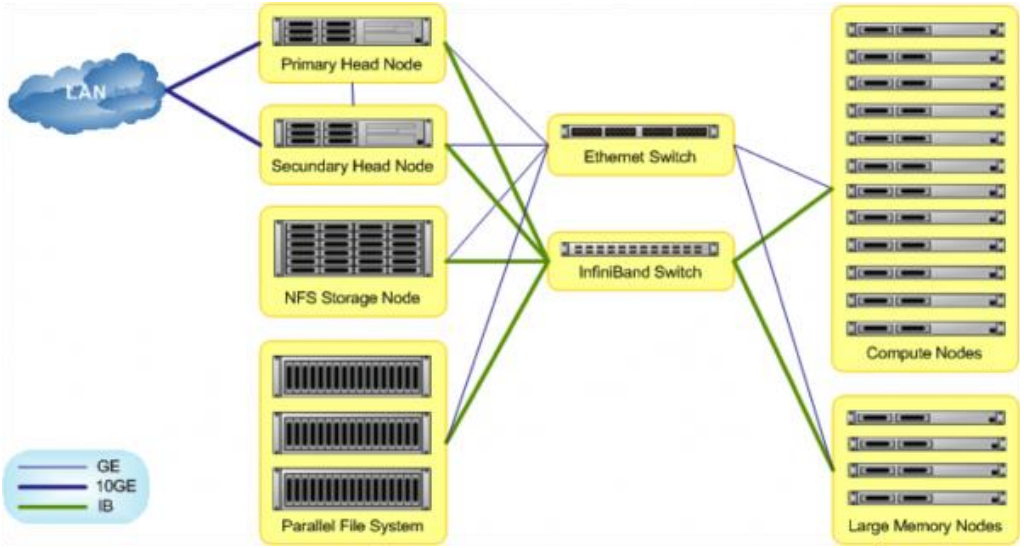

## 3 Downtimes

Move of head-node 1 day

Change /archive /project to read only 1 day
(Read only status will be a month)

Change all storage to read/write 1 day


ESTIMATE ONLY

## Moving plan

Meet Daily during Covid19

POC data moving CUH to BP

POC rolling update



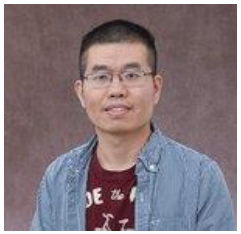

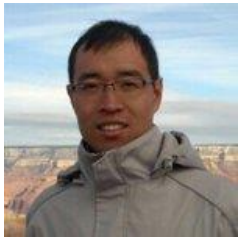

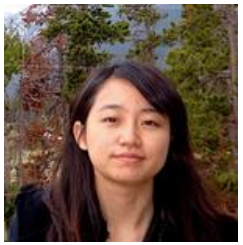



## Communication will be critical

**Beginning November 21 we will give biweekly briefings to all users (format still to be defined)**

# News

## News

BioHPC finished 2.8M jobs vs 2020 Feb 1.6M jobs(1.2M jobs in Covid-19)

Computer nodes in full capacity (Thank you IR.)

16 large memory nodes (Using amortization)

16 GPU nodes with NVidia A100 with 1.5TB memory(Using amortization)

16 GPU nodes with 4 V100 cards per nodes(Bioinformatics contribution)

16 GPU nodes with 4 A100 80GB per nodes(Purchase ongoing)

16PB new storage system online(54% used)

22PB new storage system in installation

1PB all flash new storage system(Purchase ongoing)


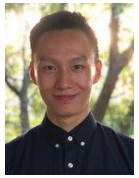

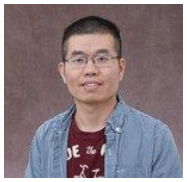

## New Network Topology

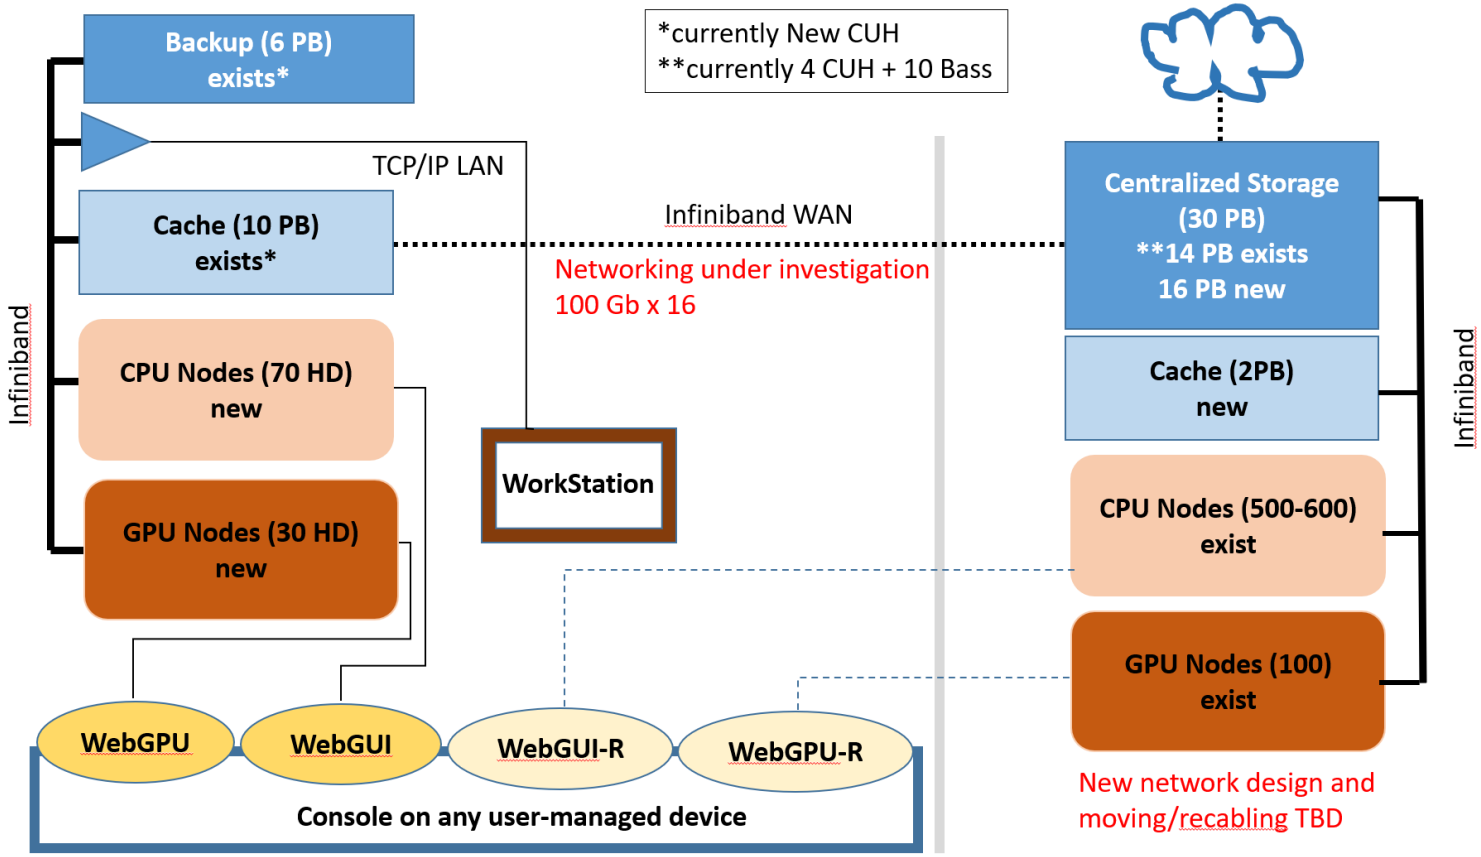


- We will increase network skeleton speed from 6 x 100GB/s to 10 x 200Gb/s
- We will have 16 x 100Gb/s network speed between Campus and QTS

## Multiple-tiers Cloud Storage


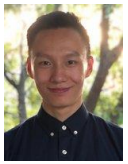


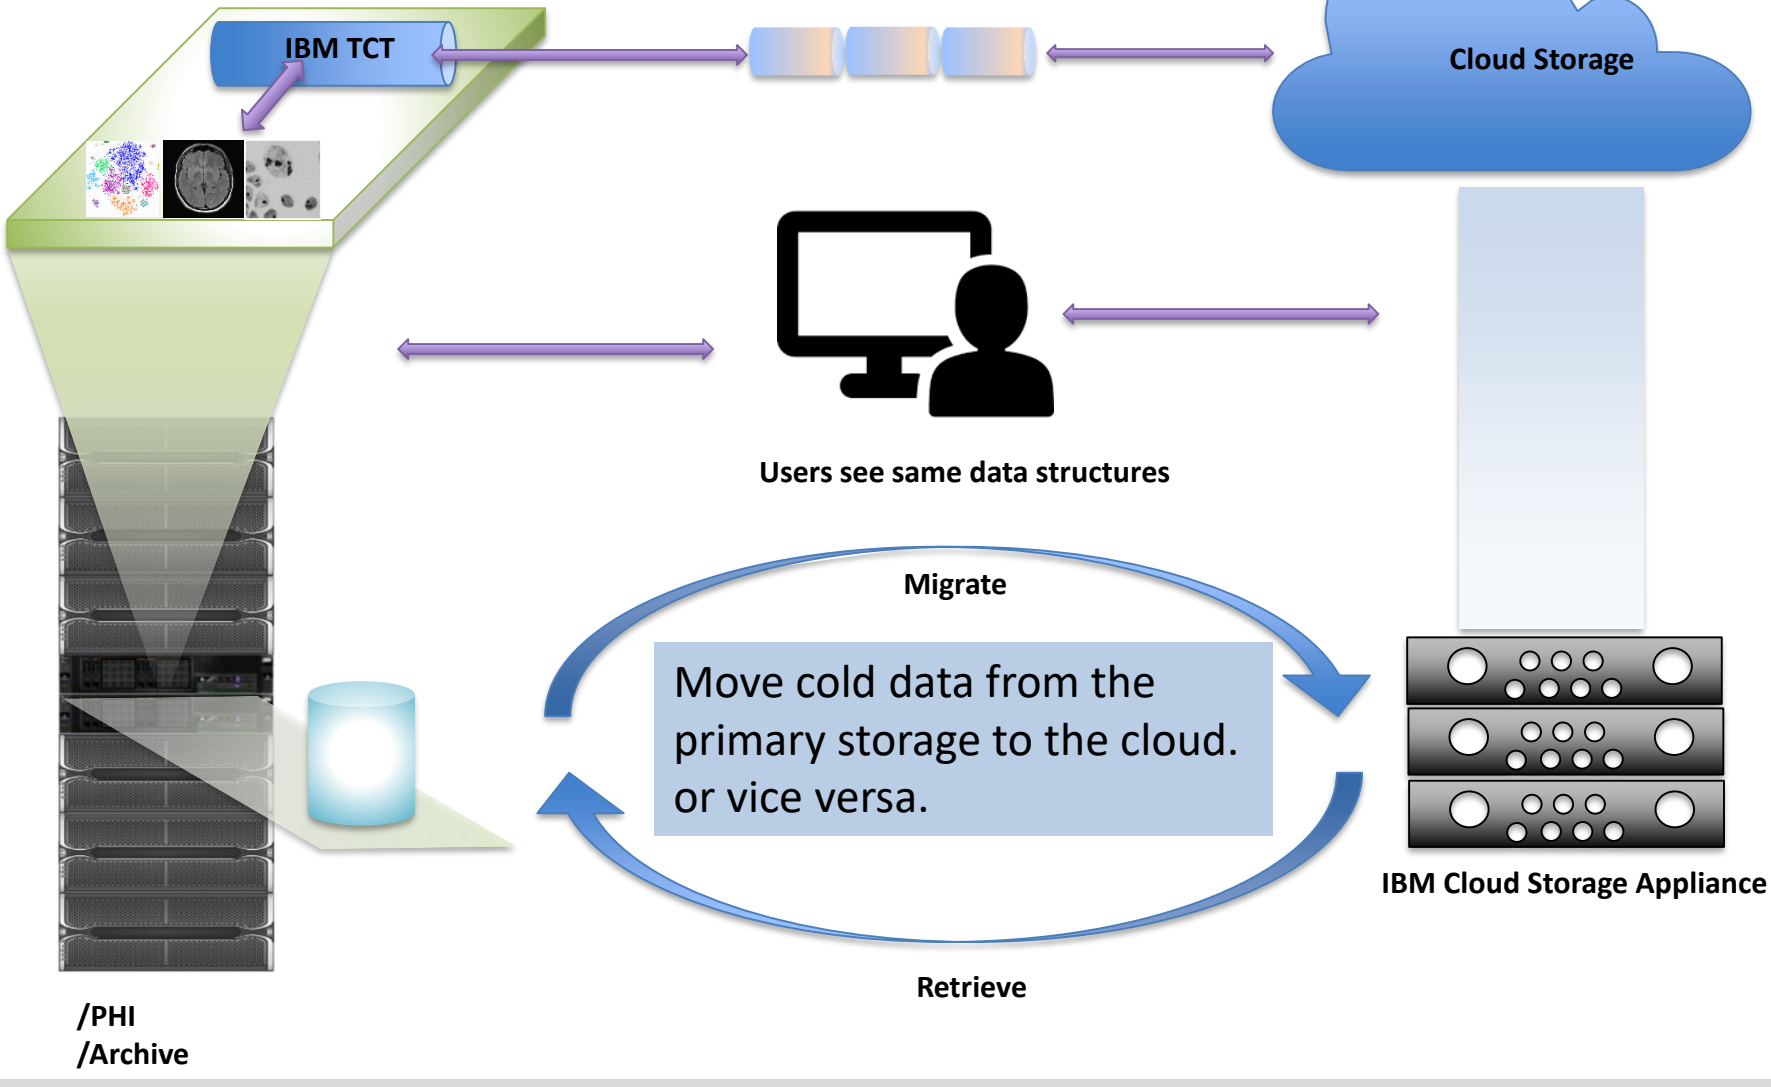

## PHI storage encryption/replication


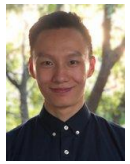


- Files are encrypted before they are stored
- Maser keys are never written to Major filesystem
- Data is safe as long as kept in the **`/PHI`**
- Two sites replication for /PHI and /work
- /PHI encryption can extend to Cloud storage

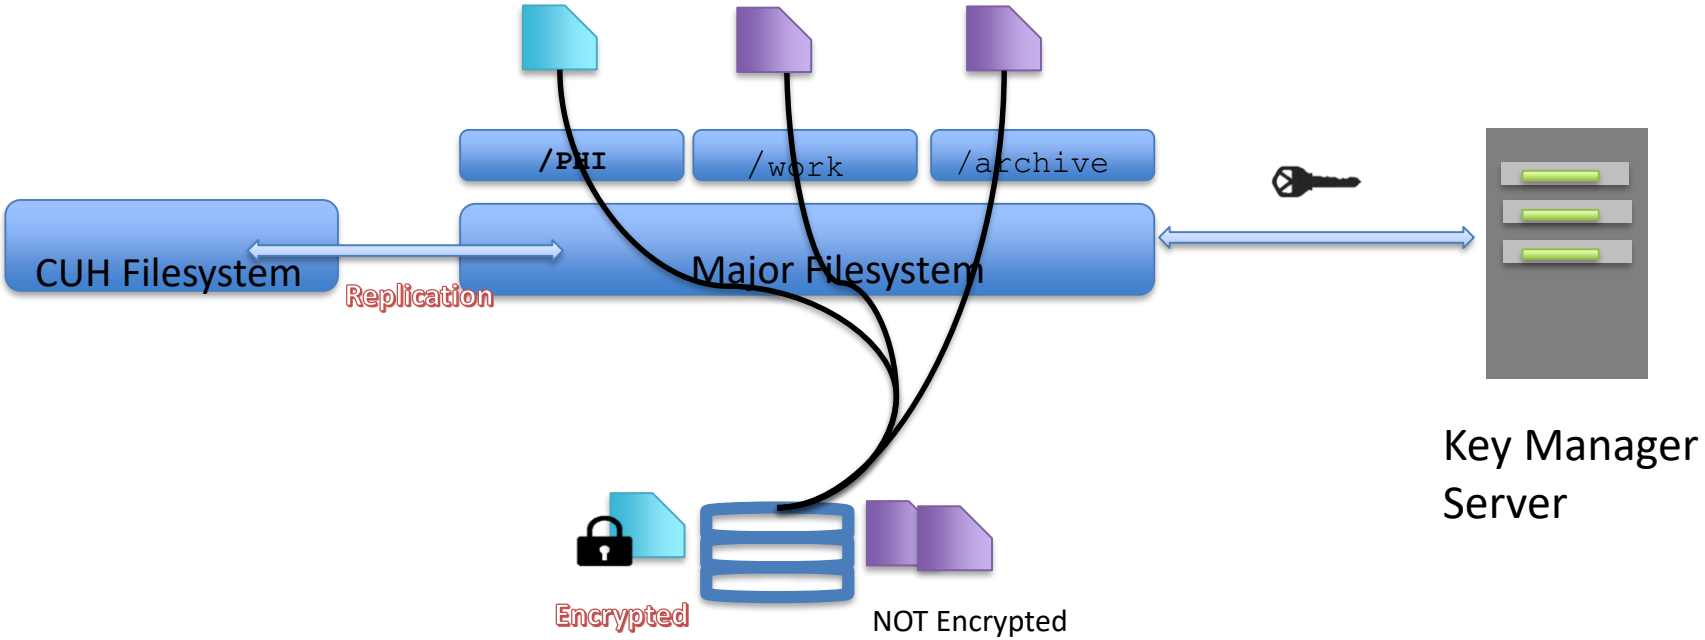

## Container builder and registry

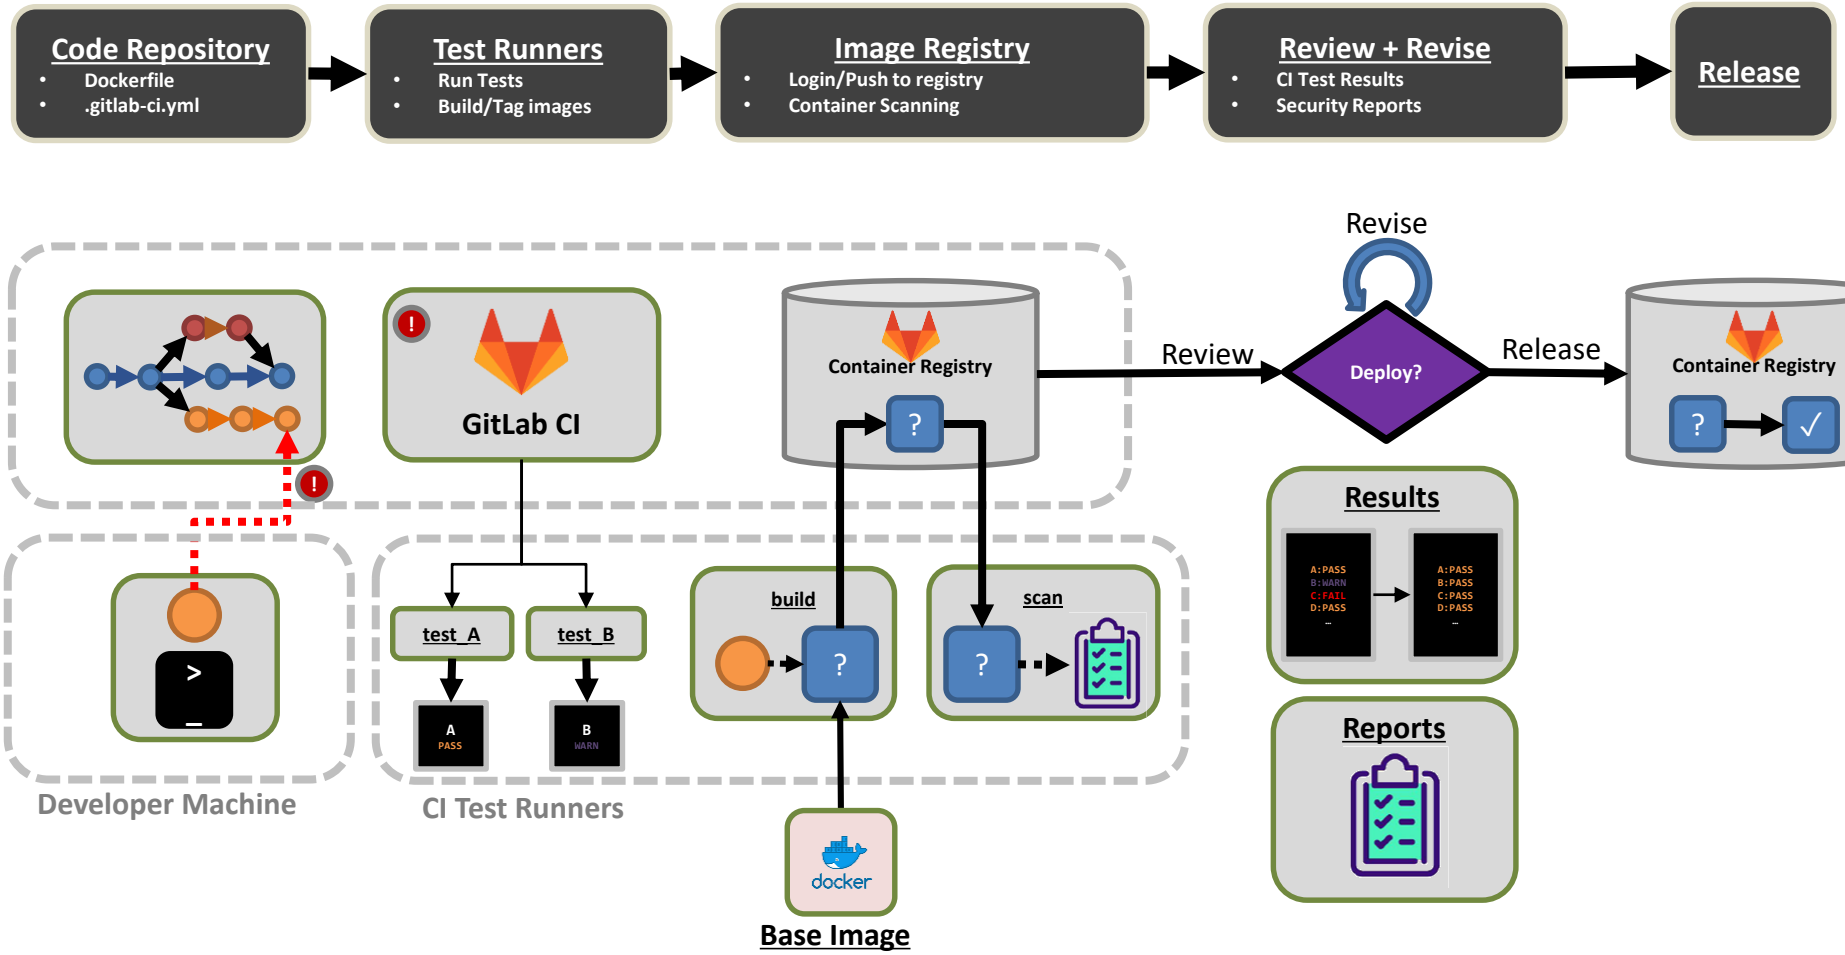

## BioHPC Applications OnDemand




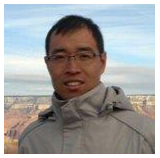

- Providing a wide range of applications through BioHPC portal
- Backend: Integration of SLUMR and a Container-orchestration system or a hybrid cluster

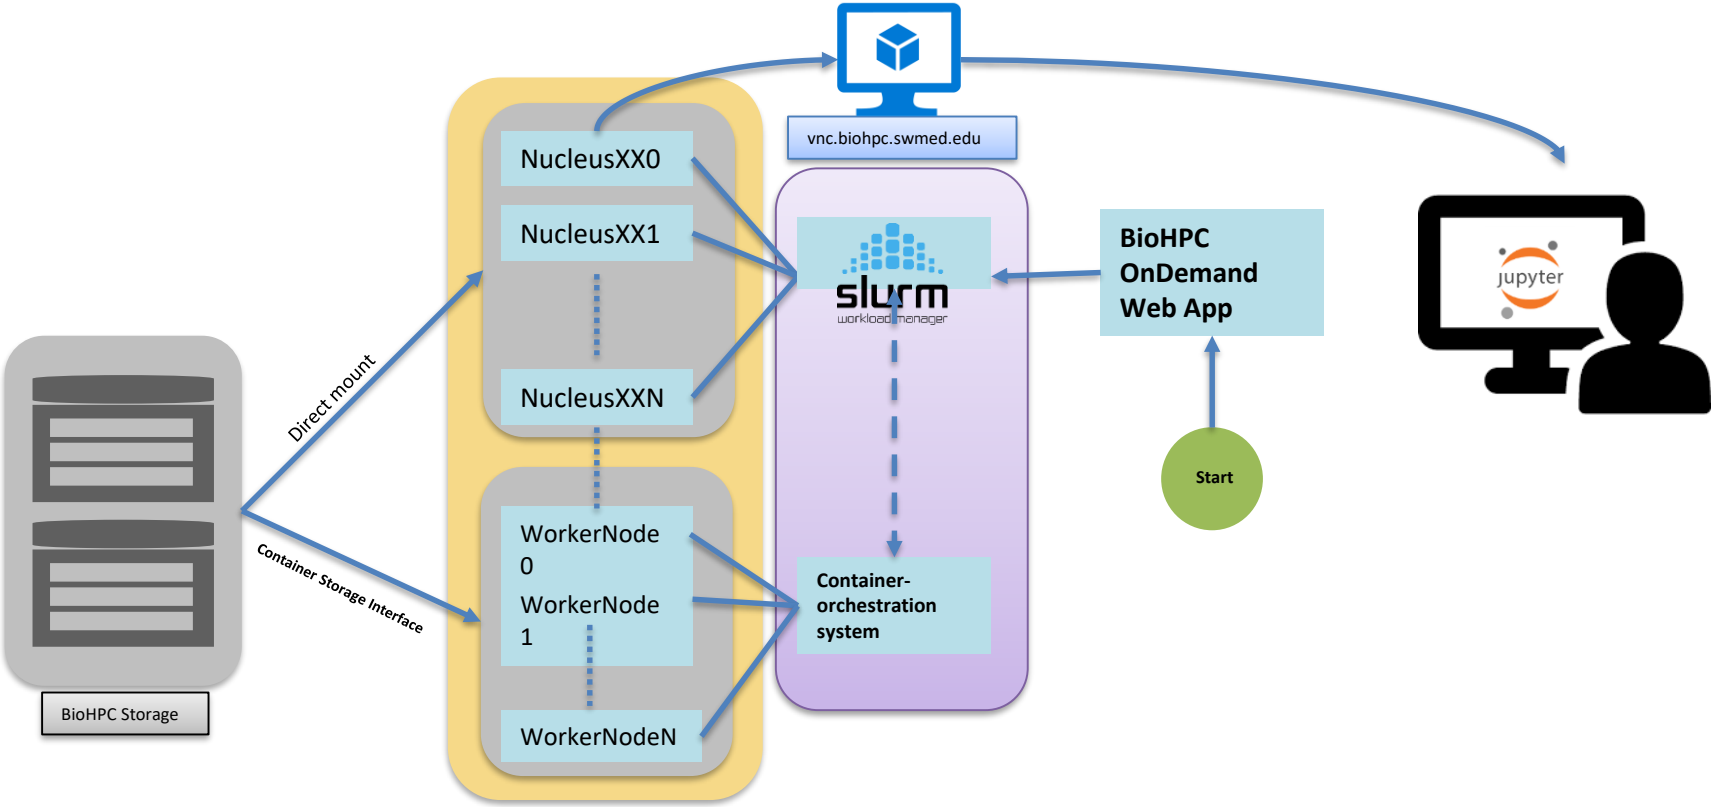

## Cluster Extension to any external location

Extend on-premise cluster to outside infrastructure without changing existing workflow.

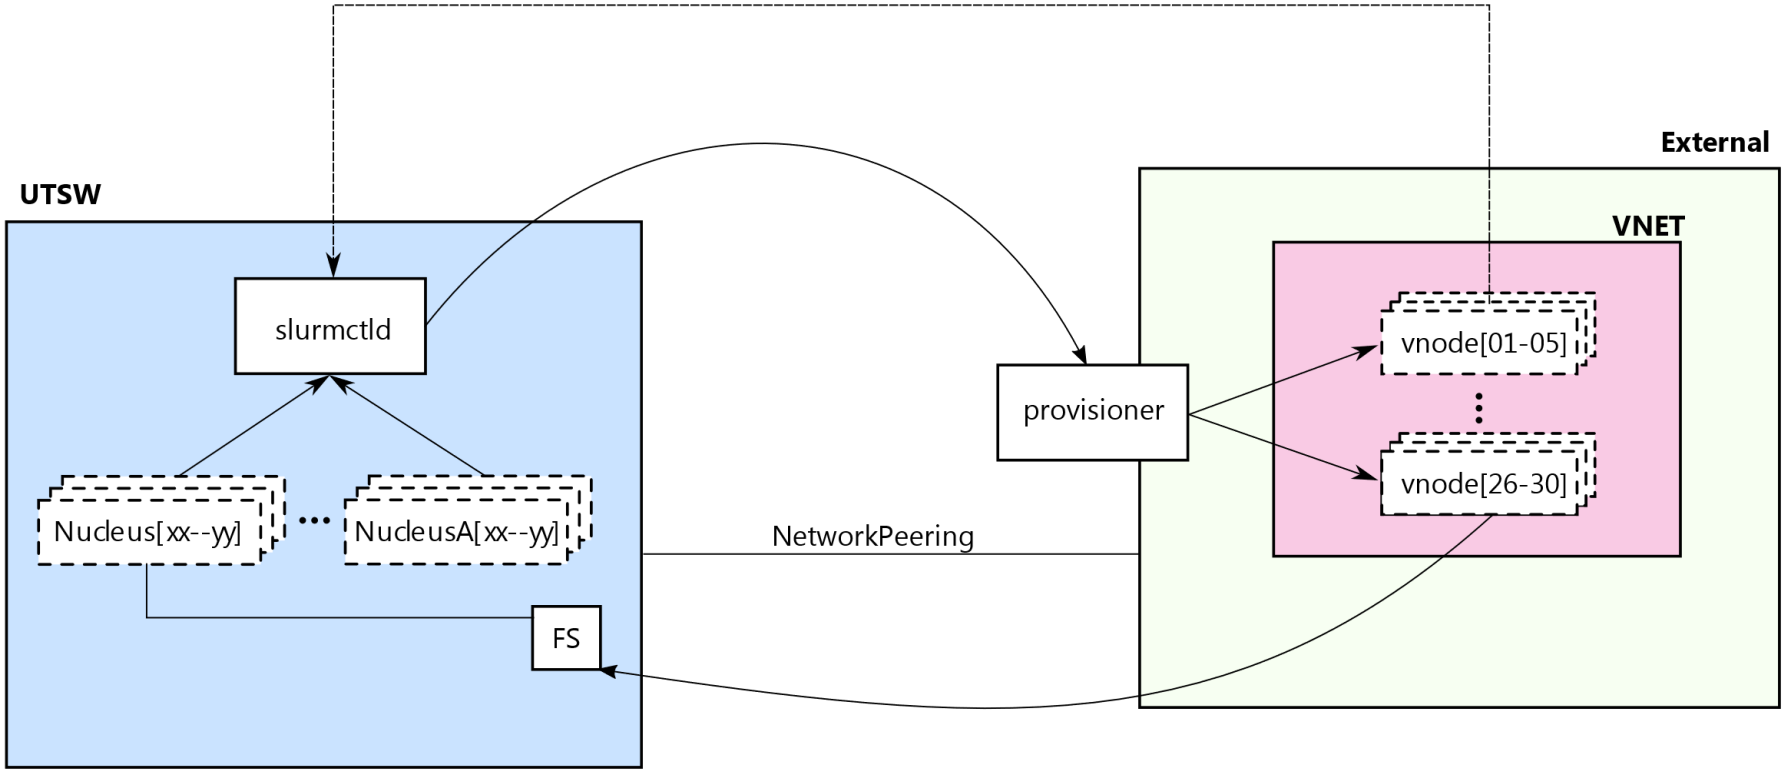

## POC – IBM AFM as HPC Cache

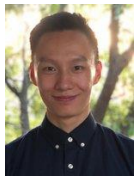

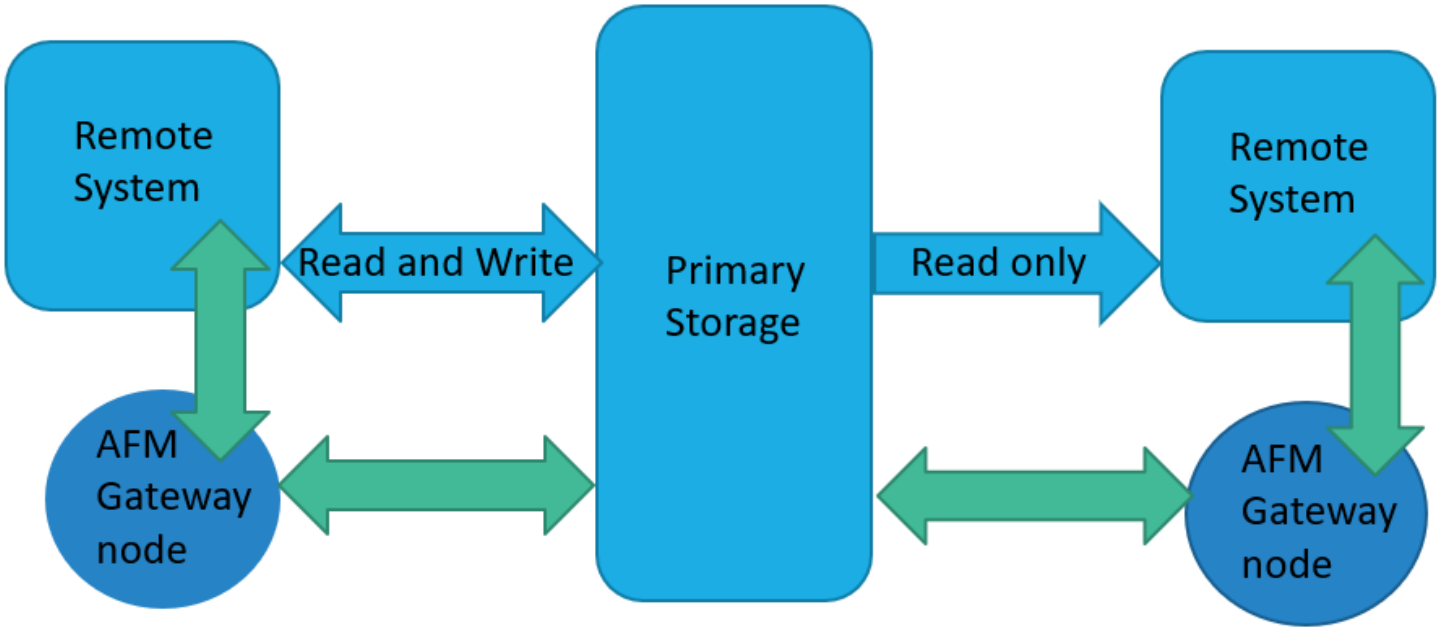


| Direct Mount | AFM HPC Cache |
|---|---|
| 1.4GB/s | 1.4GB/s |

# Astrocyte: A Platform for Scientific Pipelines

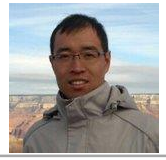



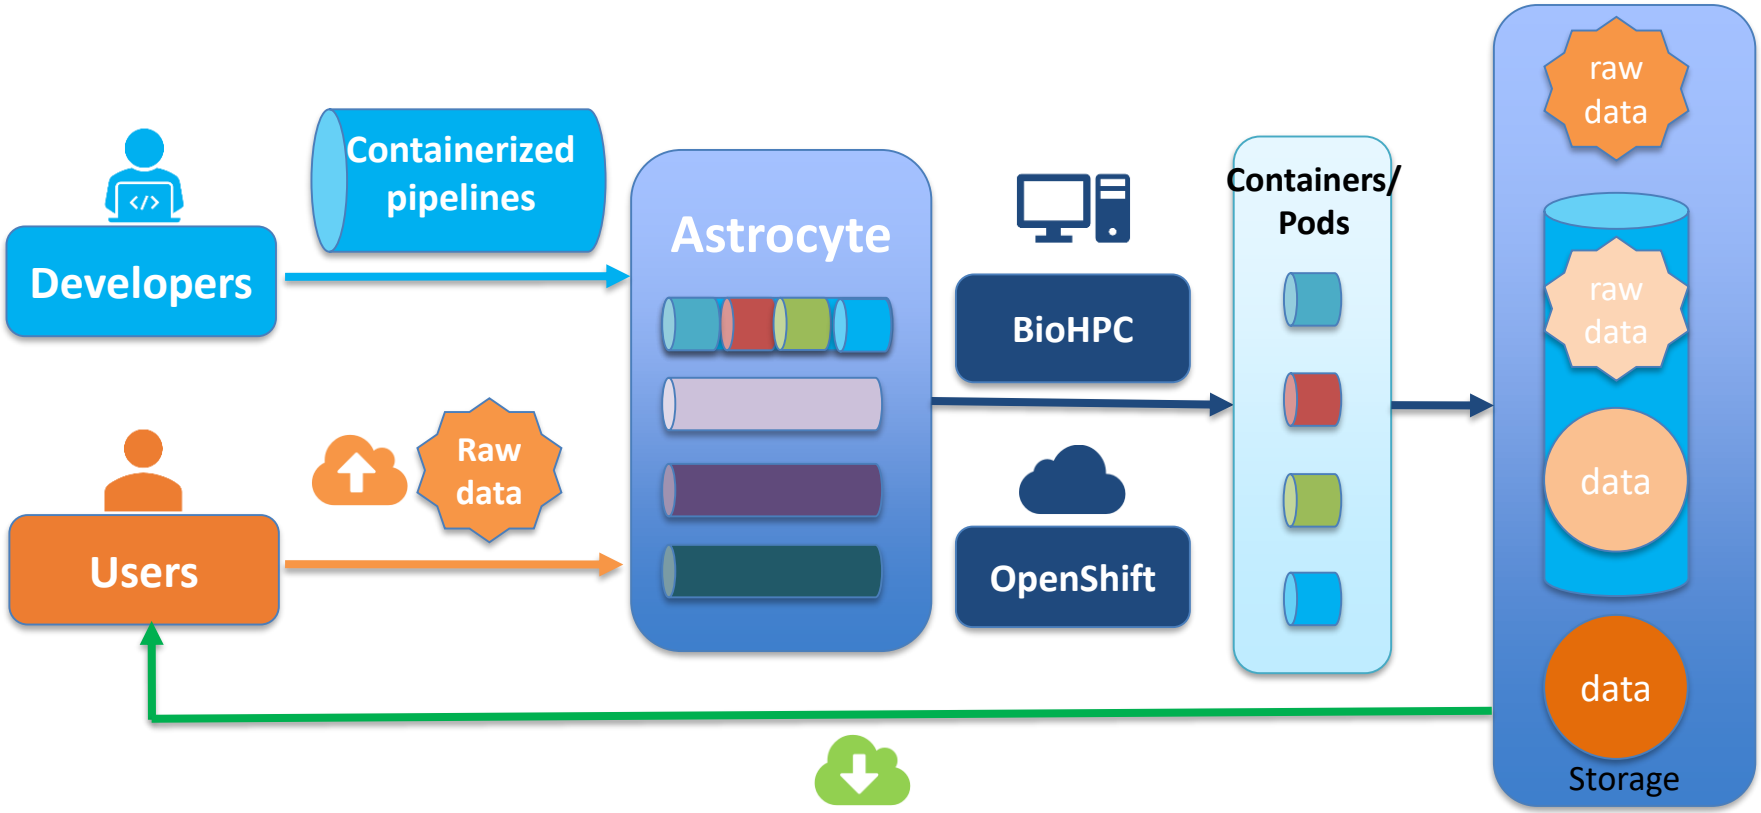


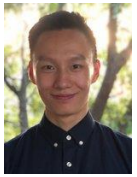
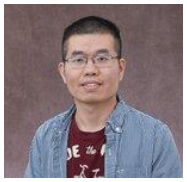

## Windows support

- Our users can use Windows environment on compute nodes by using KVM hypervisor
- KVM support the GPU passthrough to serve high-end GPU resources on Windows
- The feature of allowing users to install their own software on Windows image is coming soon

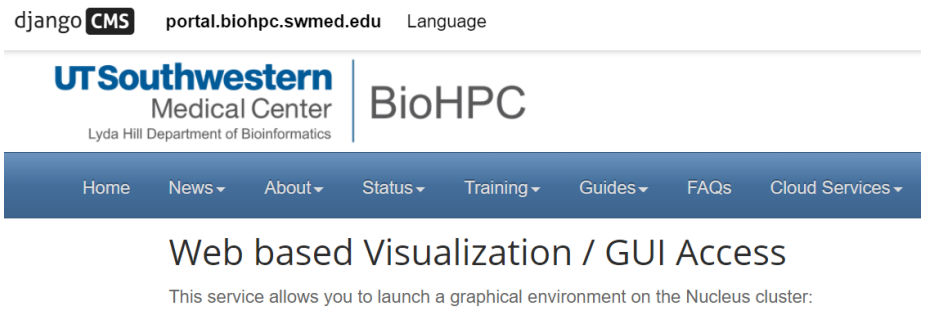


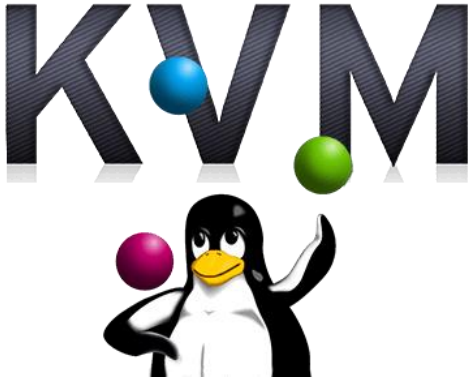

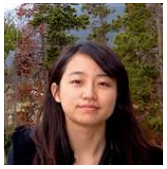

## Federated and swarming learning infrastructure

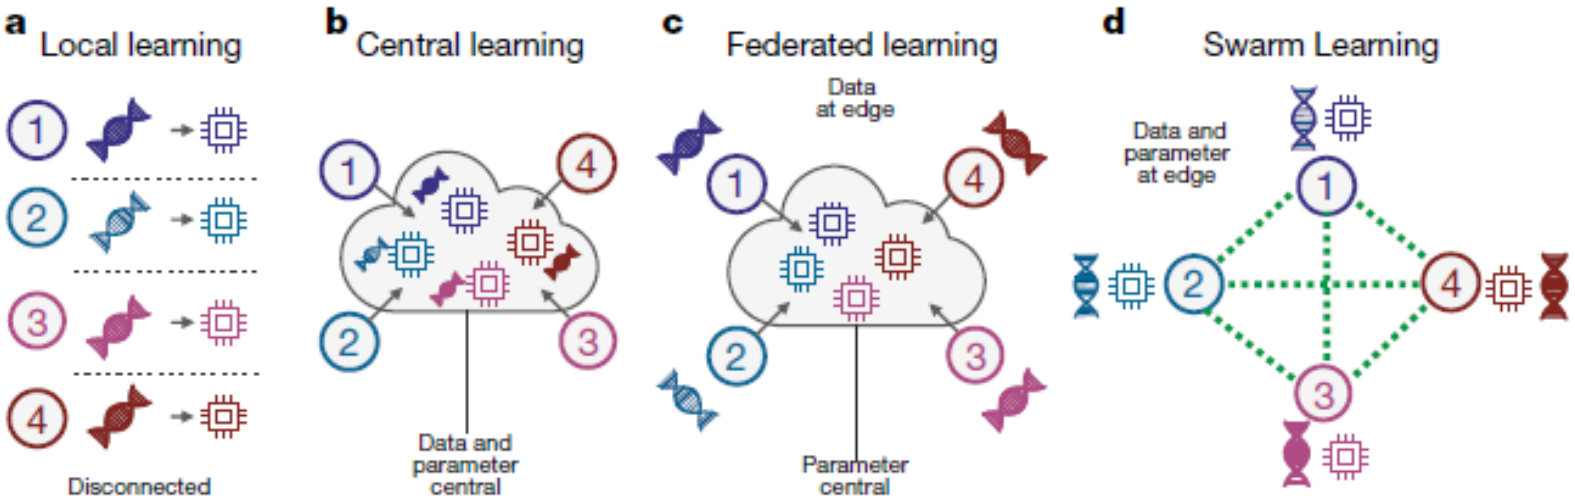


**BioHPC Approach**

- NVIDIA: clara – jupyter notebook ondemand (production)
  - Model / parameter exchange through container repository
  - DMZ server for weights exchange

- HPE: swarm learning – jupyter notebook ondemand

- Mathworks: Matlab – Web application

Image: https://www.nature.com/articles/s41586-021-03583-3

# Q&A