NCAR

Looking towards the future –
NWSC-3 Procurement

Irfan Elahi

Head, High-End Services

Computational and Information Systems Laboratory (CISL) National Center for Atmospheric Research (NCAR)

irfan@ucar.edu

Thanks to many folks at NCAR for their slide contributions

Disclaimers

Nothing in this presentation or subsequent request for information should be construed as a Request for Proposal (RFP) and no offers are being solicited at this time.

UCAR reserves the right to discontinue plans to procure the NWSC-3 Computing Platform with no responsibility to interested parties.

NWSC-3 Strategic Objectives

- Provide a higher level of performance for Earth systems science (ESS) workflows
- Invest in technologies with potential to transform the ESS landscape and fosters the acceleration of new technologies suited for ESS applications
- Enable application optimization and refactoring efforts required to make use of these technologies
- Ensure that users remain productive throughout the transition to NWSC-3, and beyond

NWSC-3 Highlights

- Highly efficient, usable, and reliable HPC resource that achieves 3-5x the sustained performance of Cheyenne
- High performance storage and parallel file system that achieves a balanced compute/storage ratio
- Take advantage of NCAR's existing HPC and data storage resources
- Enable the transition to future architectures via new processors, storage technologies, and software
- Provide flexible technical options; future expansions of HPC system, storage, maintenance and services
- Deploy the system in 1H-2021 for a production target of July 1, 2021
- Estimated funding $30M via Best Value procurement
- Deploy the systems at the NCAR-Wyoming Supercomputer Center in Cheyenne, WY

NWSC-3 Budget

- **$30M Procurement Budget**
  - Traditional HPC Compute & all associated software
  - Heterogenous/throughput compute environment
  - Storage and File System
  - 4-year maintenance support and licenses
  - Delivery, Installation and acceptance
- **$2-4M Facility and fit-up budget**
  - Vendor independent facility work
  - System specific facility fit-up/preparation
- **Disk/Tape based Archive**
  - The NWSC-3 budget does not include Disk/Tape Archive systems

NWSC-2 Supercomputer “Cheyenne”

Current HPC Environment

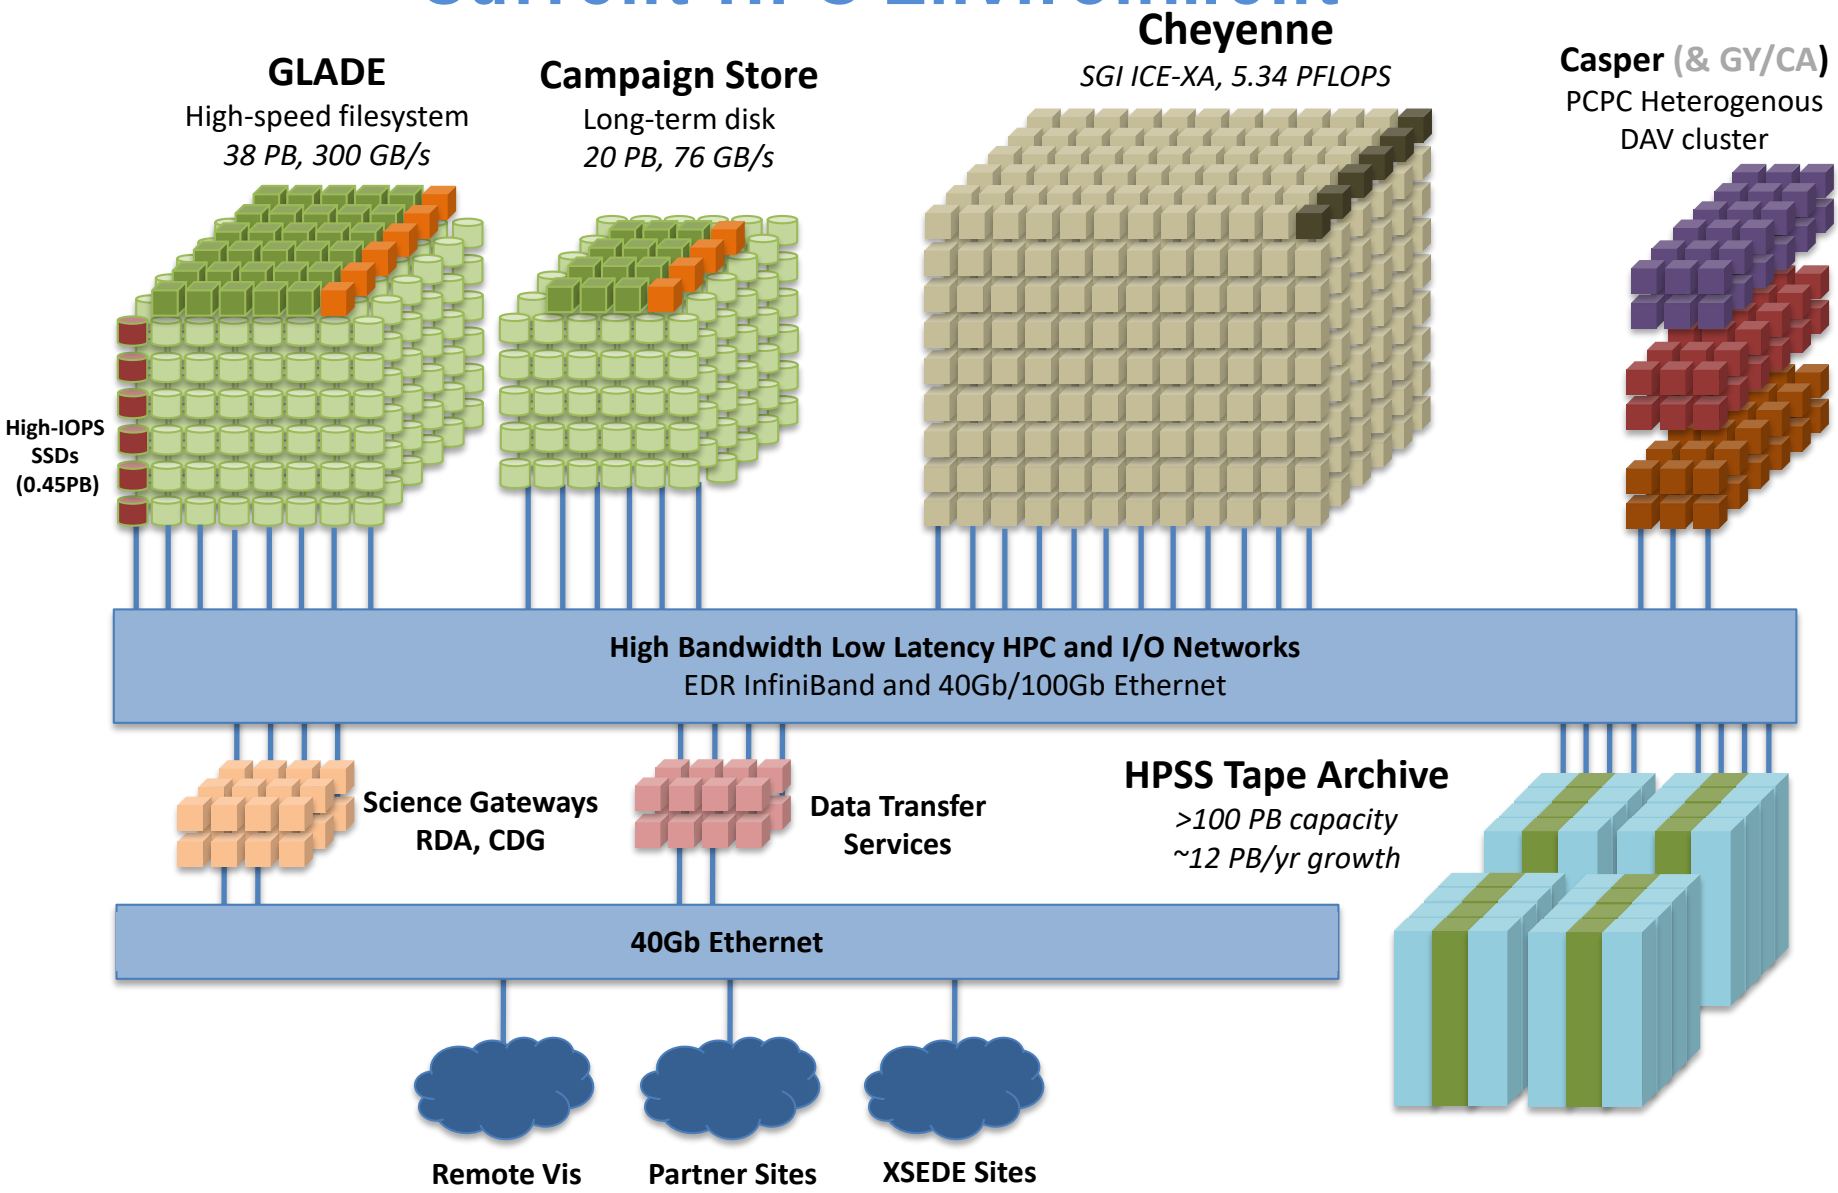

NWSC-3 Project Schedule *(Subject to change)*

- **Late 2018 – Mid 2019:-** Benchmark Design
- **Late 2018 - Mid-2019 :-** Technology Briefings and Architecture Co-Design
- **Feb 2019 - May 2019 :-** Science Requirements Process & Workload Study
- **Late 2019 :-** RFP Release
- **Early to Mid-2020 :-** Vendor Selection and Approval
- **Mid-2020 - Early 2021 :-** Facility Preparation, Vendor System Build and Factory Trials
- **Early 2021 - July 2021 :-** System Delivery, Installation and Acceptance
- **July 2021 :-** Production start, 6-month overlap with Cheyenne
- **December 2021:-** Decommission Cheyenne

NWSC-3 RFP Timeline

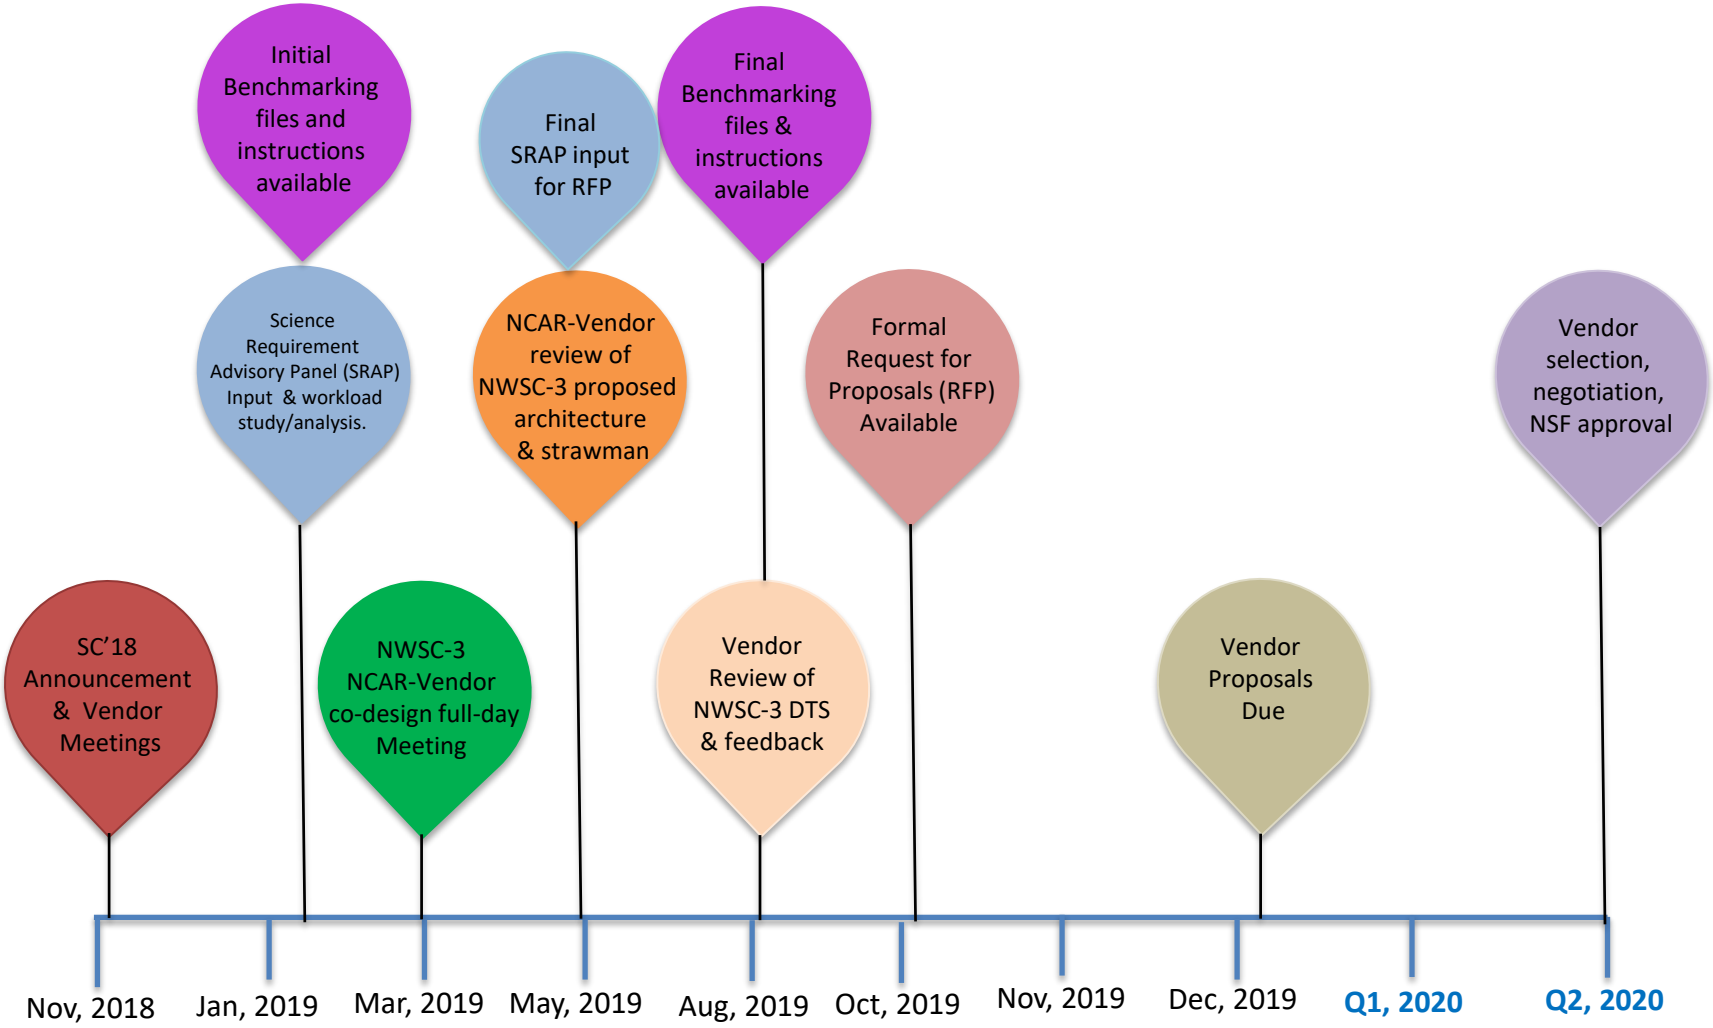

Managing the many choices…

❑ Possible New Compute Architectures

- **CPUs:** Intel XEON, IBM POWER9, ARM (Cavium ThunderX), AMD (Epyc)
- **Coprocessors:** GPUs, FPGAs, NEC Aurora
- **Network:** OPA V2 (Intel), Cray proprietary, InfiniBand (HDR), Ethernet
- **Machine Learning:** Intel Nervana NNP, NVIDIA tensor cores, Google TPUs

❑ Memory & Storage

- **High-speed, stacked memory**
- **High IOPS Non-volatile Storage:**

❑ Cloud Computing Considerations

- **On-Premise Cloud**
- **Commercial Cloud**
- **HPC in the Cloud**

NWSC-3 Attributes

(Focus on a design that enhances the end-to-end rate of science throughput}

- ✓**Heterogeneous hardware**
- ✓**High IOPS Storage**
- ✓**High Bandwidth Memory**
- ✓**Application containers**
- ✓**Cloud bursting capability**
- ✓**Storage tiers**

Growing and Emerging Cloud Workload

- **Use Cases**
  - High Availability (Antarctic Forecast on Cloud when HPC down)
  - Bursting to cloud for urgent runs (e.g., Hurricane forecasting)
  - General small job overflow for users willing to use more expensive allocation (e.g., WRF job swarms)
- **Developing Sample Containers to go with our applications**
  - On- or Off-Premise
  - Capturing complex environment needs
- **Data Discovery, Curation, Analytics and Storage – more to come**

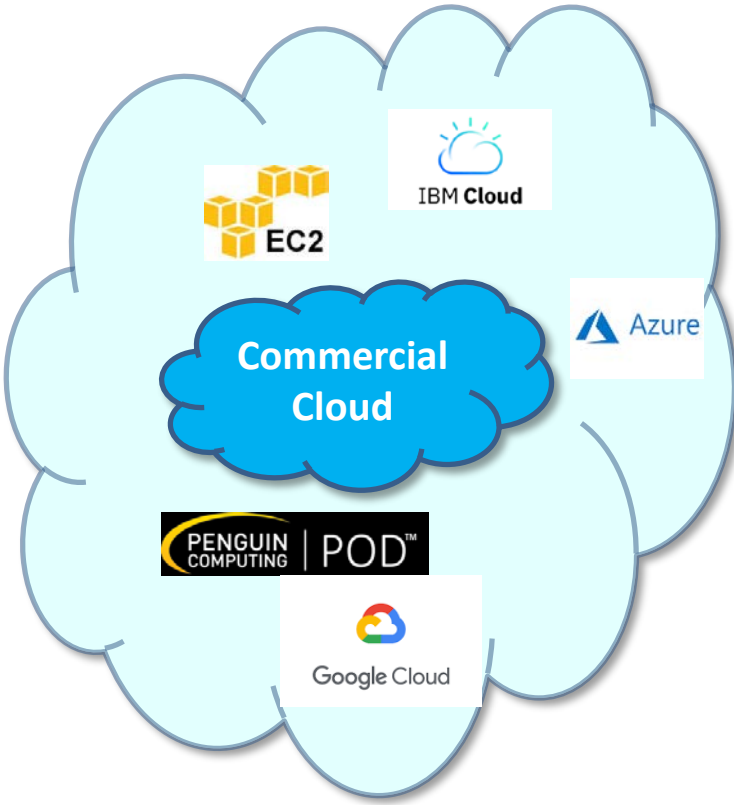

Workflow Centric Architecture

User defined job placement

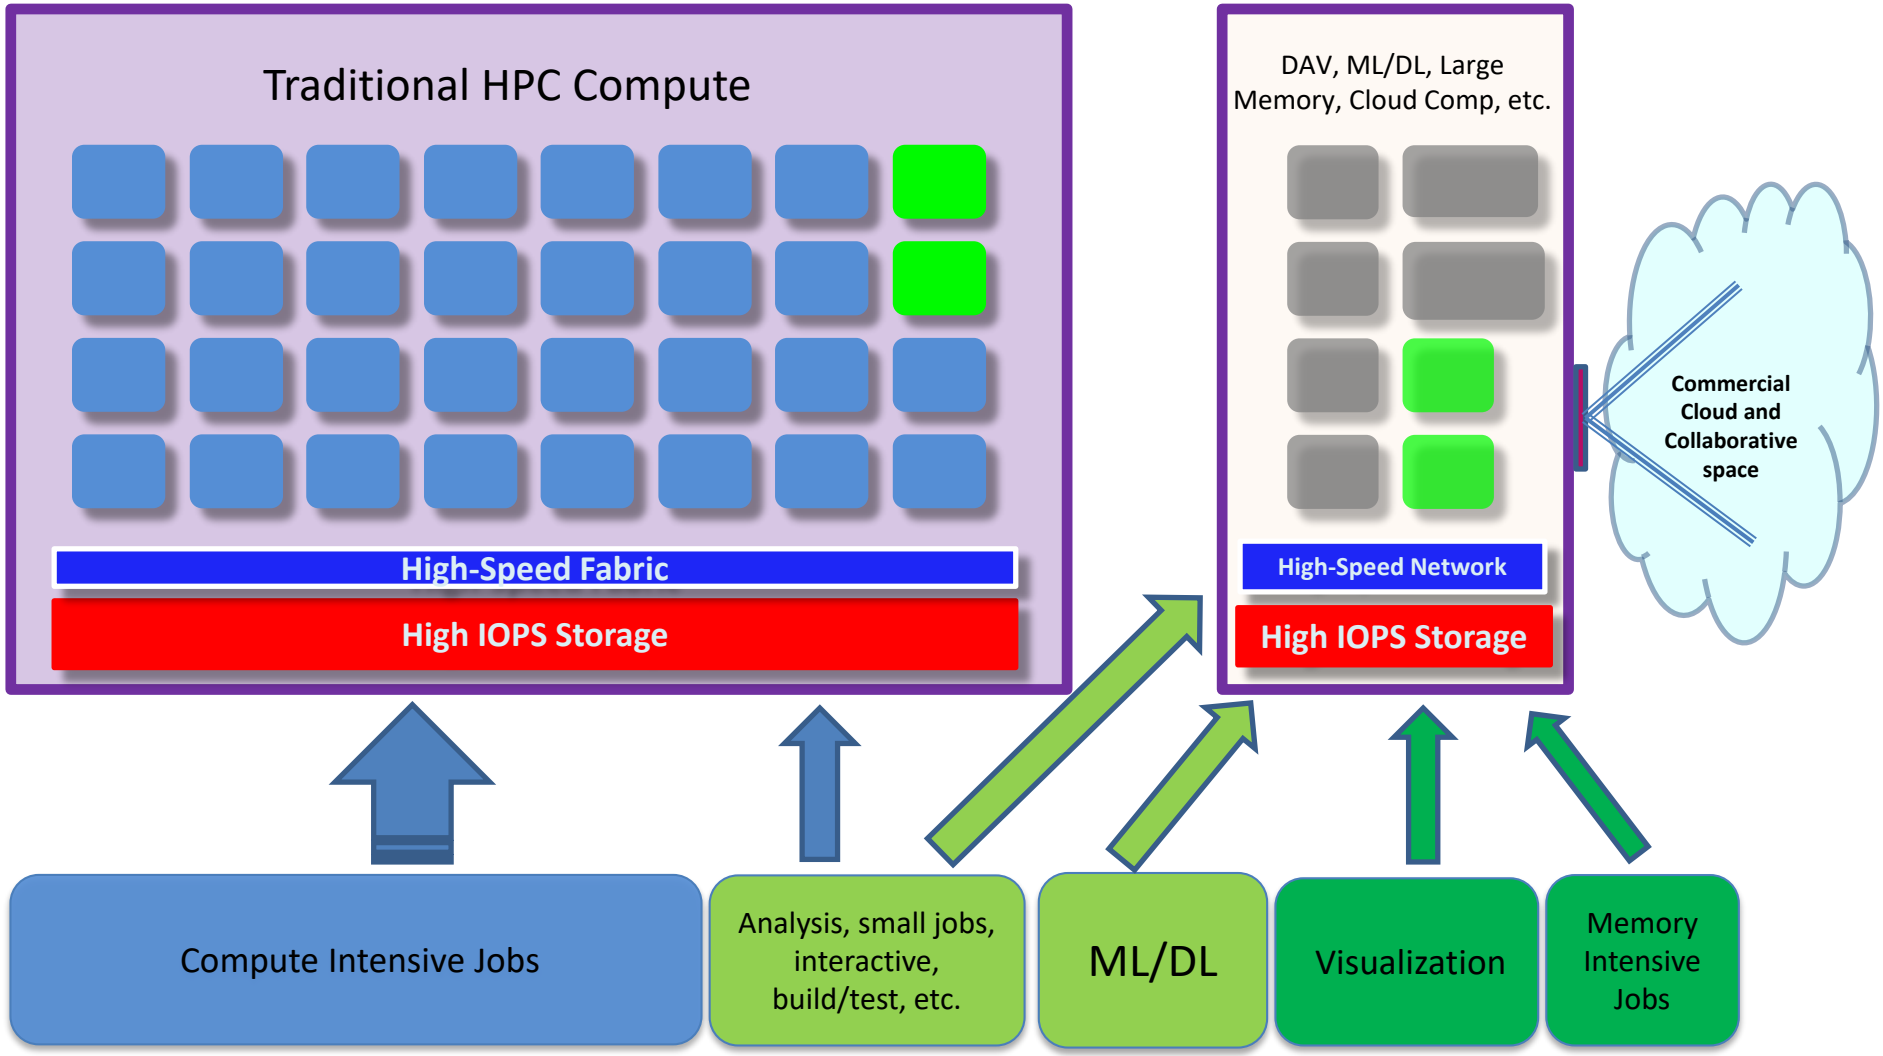

NWSC-3 Benchmarks

❑ Benchmark Characteristics

- Best value procurements through competition
- Lightweight benchmarks, portable codes
- Lower barrier of entry for vendors with easier-to-use benchmark suite

❑ Benchmark Components

- Kernels and Synthetic (hardware characteristics)
- Capability (”scalable science” runs)
- Capacity (throughput)
- Data Science and Deep Learning (*new!*)

Summary
NWSC-3 Procurement Strategy

- **Increase Vendor Pool**
  - Keep pricing honest in procurements through competition: lightweight benchmarks, portable codes
  - Lower barrier of entry for vendors with easier-to-use benchmark suite
  - Best value procurement
- **Create Environment to consider many solutions**
  - Retire as much technical risk as possible (hardware and system software)
  - Prepare Applications to handle wide-range of systems
- **Keep an eye on TCO (Total Cost of Ownership)**
  - Energy costs will influence decision (consider GPUs where possible)

Thank you

Questions?

Backup Slides

Growing Portfolio of Small Jobs

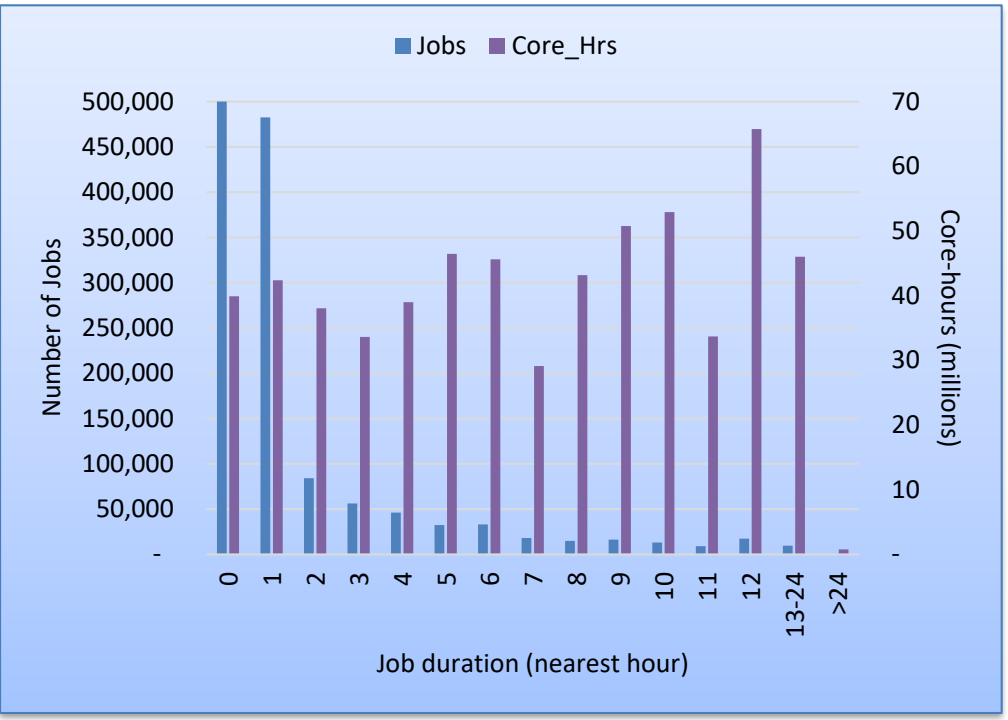

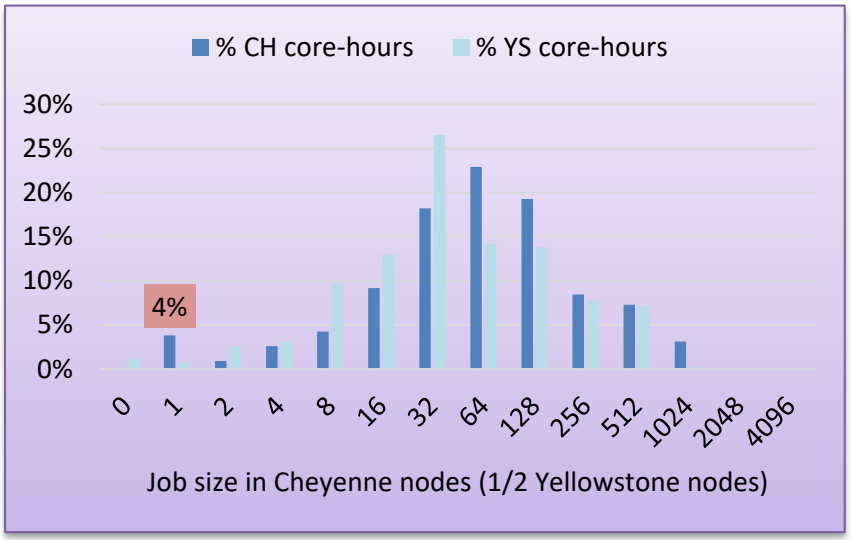

- **Pure "Capacity" Workload**
- **10.8M jobs (93% of all jobs) run <30mins**
  - but 7% of the core-hours
  - 3 projects responsible for 85% of these core-hours
- **4% of jobs single node**
- **Small jobs are creating jitter and noise (system instability) for larger runs**

Budget Breakdown Structure

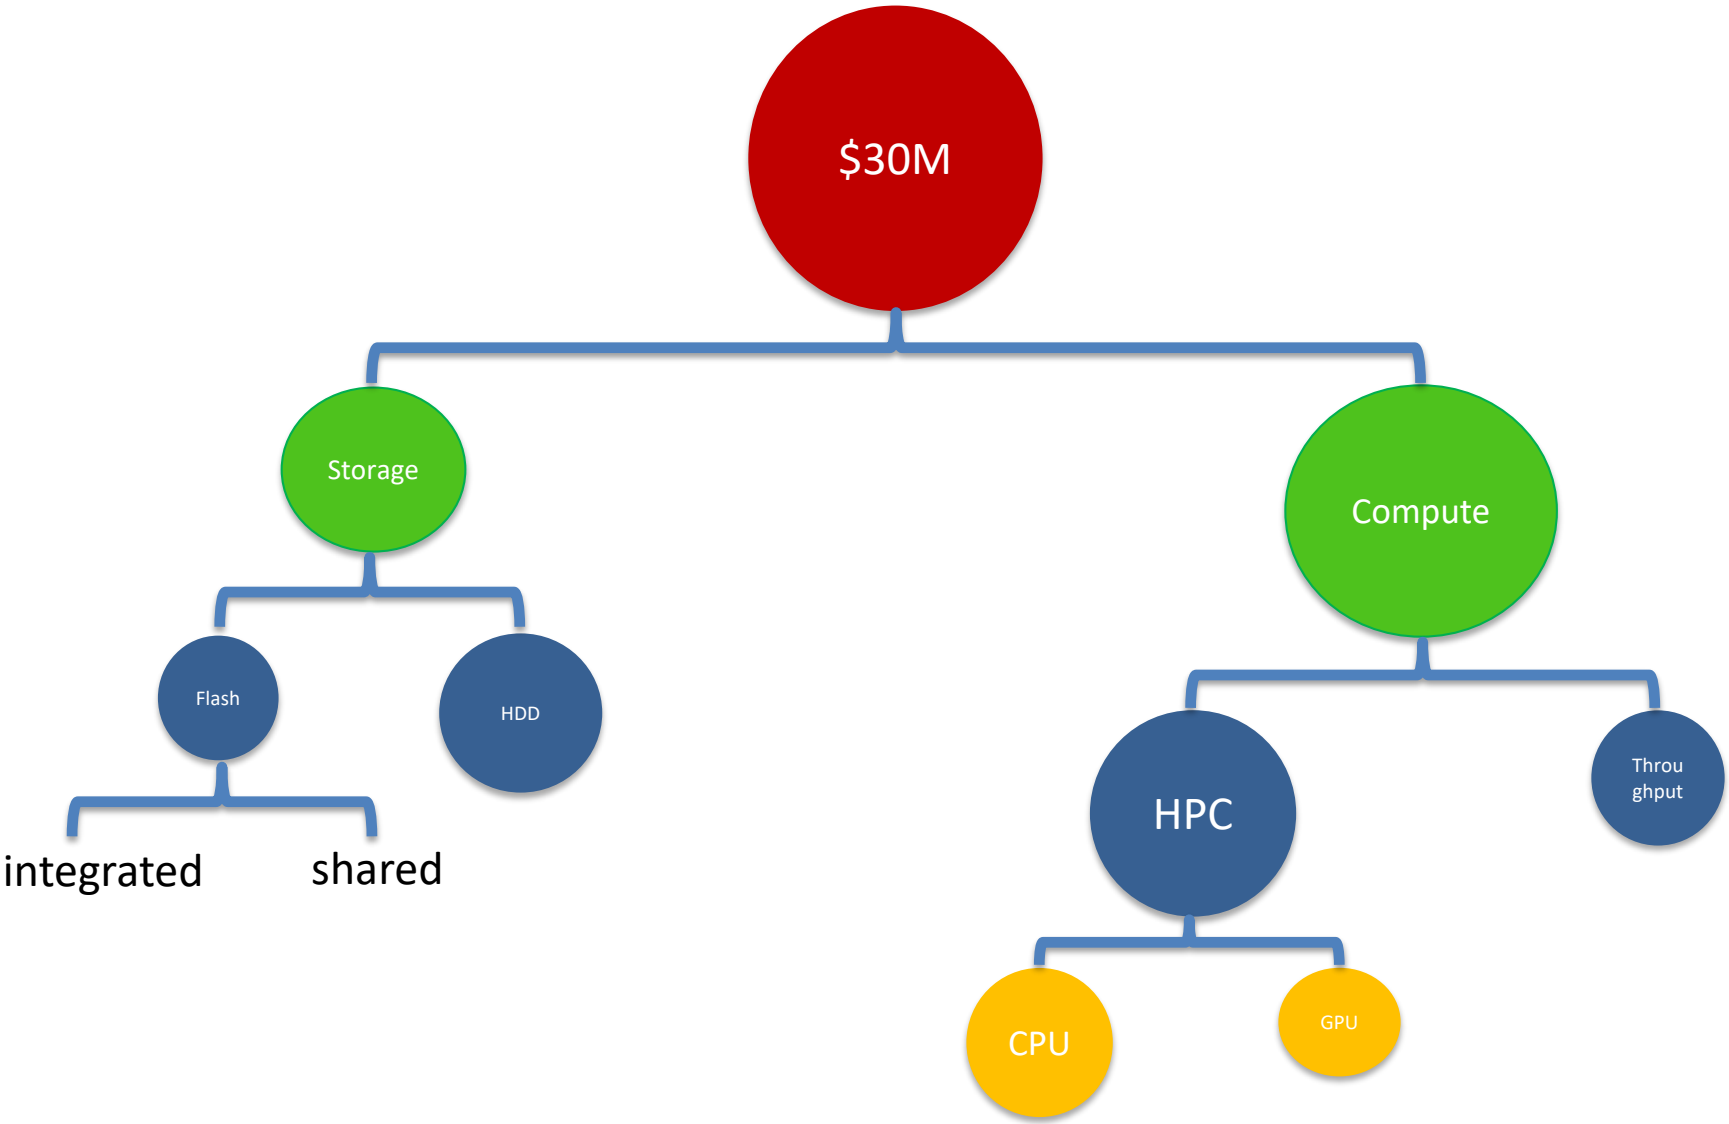

NWSC-3 Procurement Process & Timeline