

# Next HPC-System at PC2

**Robert Schade**

Paderborn University, Germany  
Paderborn Center for Parallel Computing

PC2 User Meeting, Paderborn – 20 October 2023



Paderborn  
Center for  
Parallel  
Computing

- What comes after Noctua 1 and Noctua 2
- Timeline and Process
- Requirements for the next system and how we came up with them

# What comes after Noctua 1 and Noctua 2

As one of the nine NHR centers in Germany we have funding for

- general, scientific and application support,
- PhD-positions in the NHR-graduate school,
- regularly new HPC systems and infrastructure

Ministerium für  
Kultur und Wissenschaft  
des Landes Nordrhein-Westfalen



Thus, the next HPC system is due:



Noctua 1, inaugurated 2018



Noctua 2, inaugurated 2022



Next HPC system at PC2,  
to be inaugurated in early 2025

(Noctua 1, our new data center (building X), and part of Noctua 2 were funded from the Forschungsbau-Projekt "Noctua")

# What comes after Noctua 1 and Noctua 2



Noctua 1, inaugurated 2018

Noctua 2, inaugurated 2022

Next HPC system at PC2,  
to be inaugurated in 2025

Noctua 1: 2018-2025

Noctua 2: 2022- ~2027

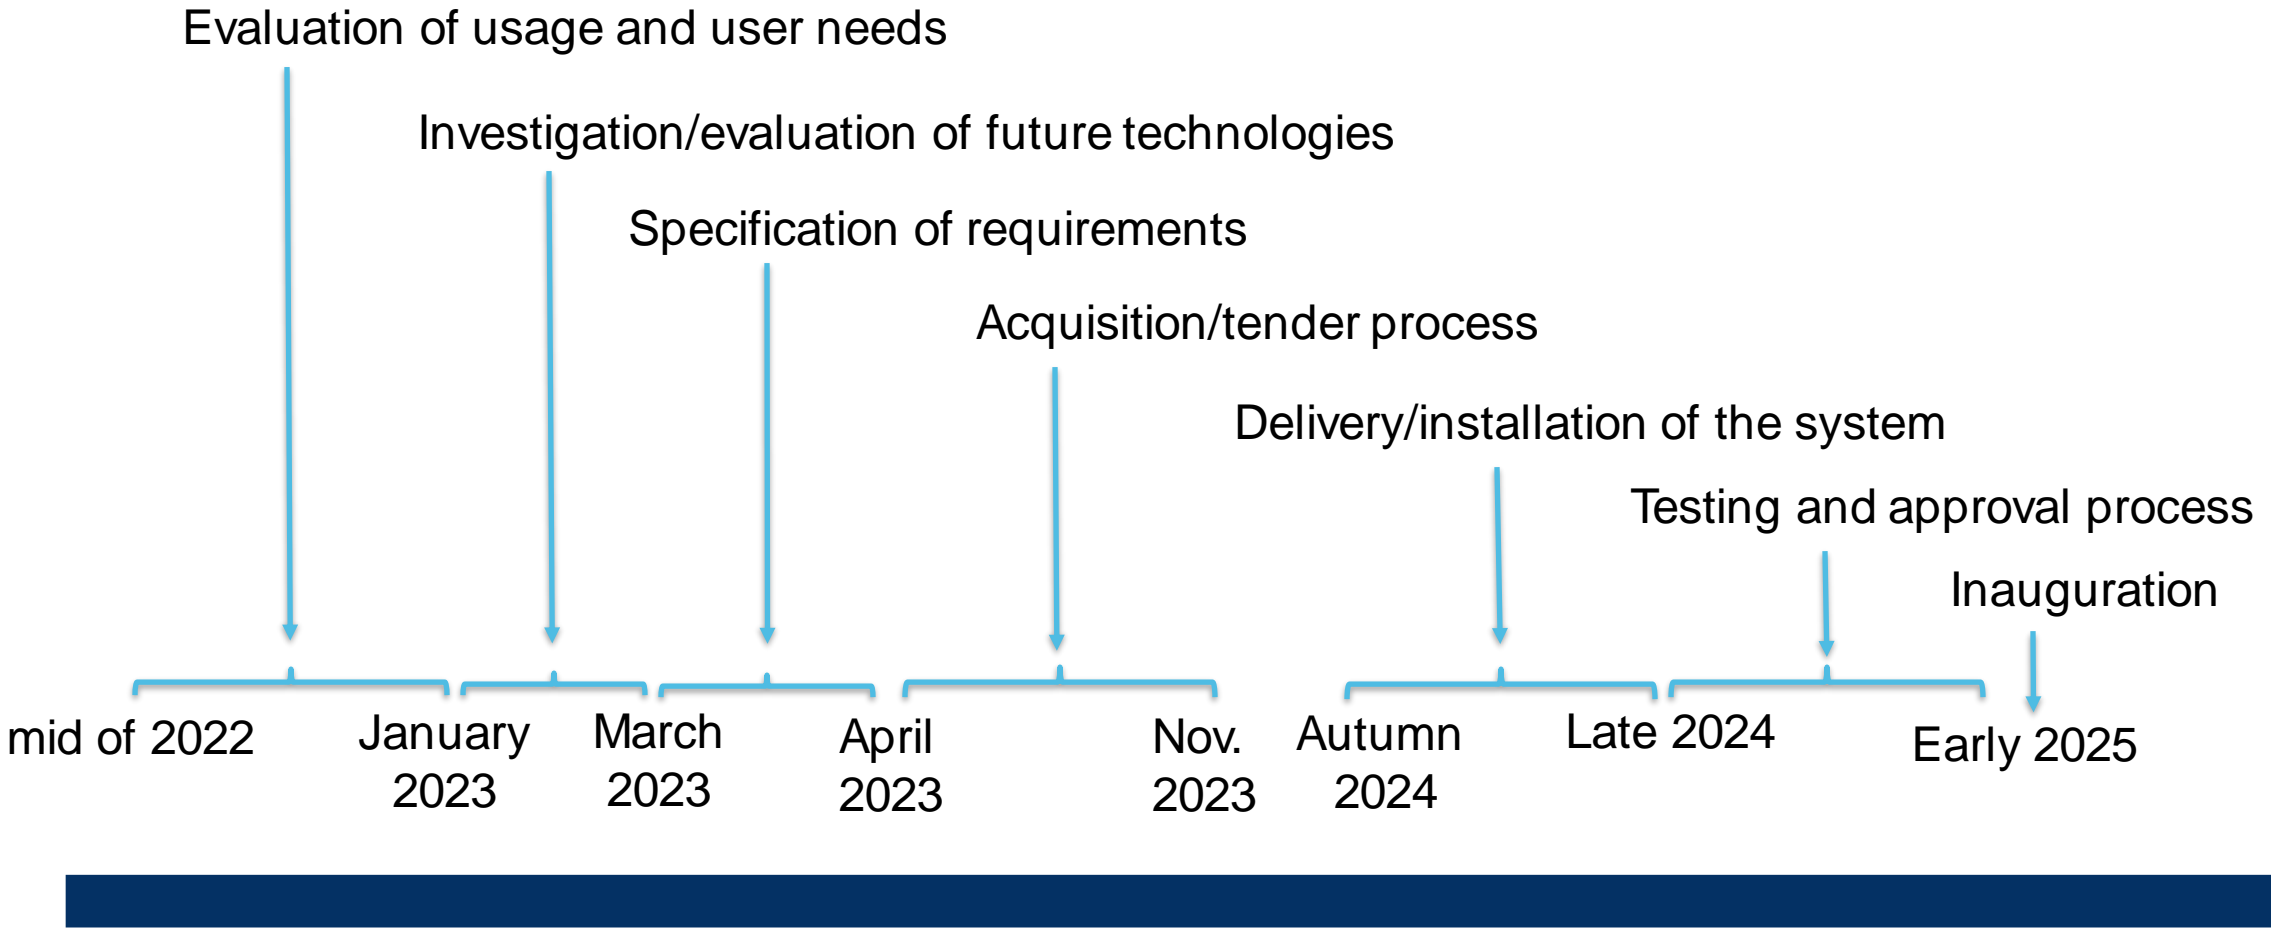
Next HPC system: 2025- ~2030

Ministerium für  
Kultur und Wissenschaft  
des Landes Nordrhein-Westfalen



# Timeline and Process

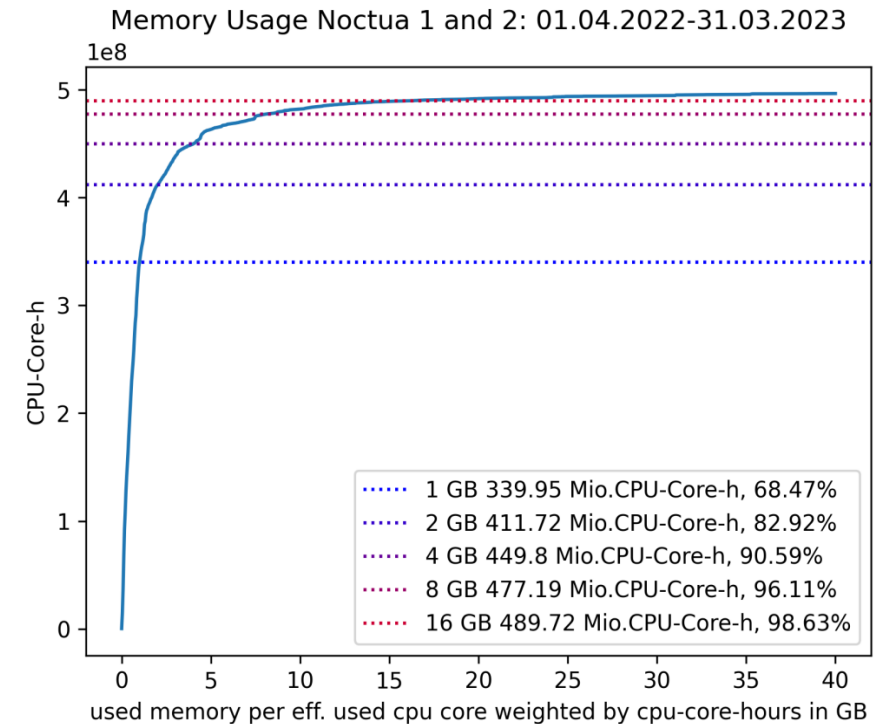
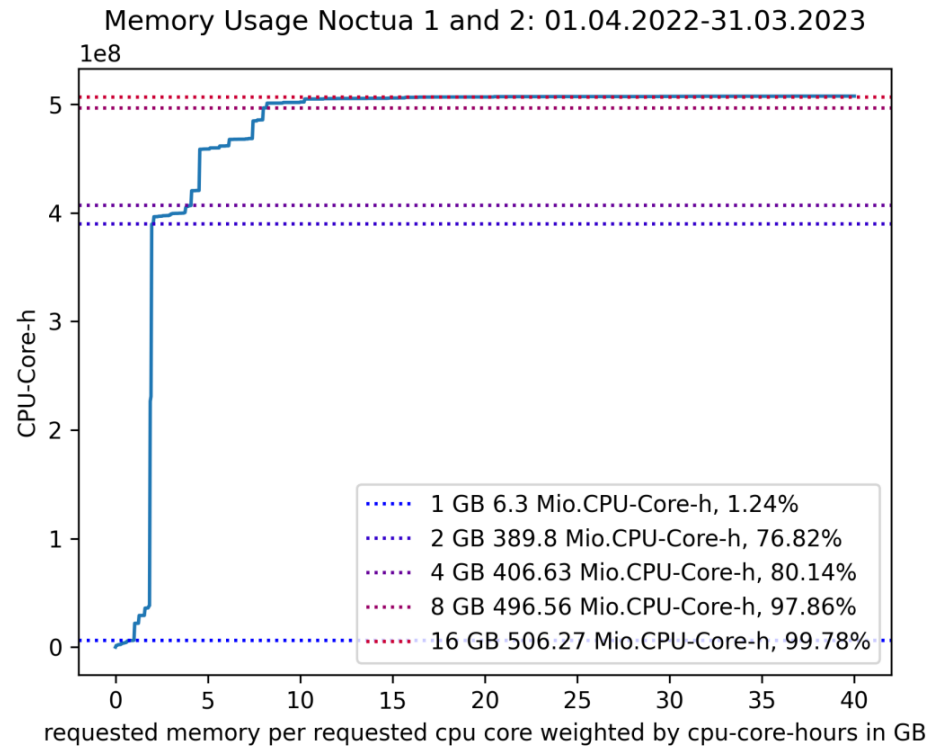
## Steps:



(Dates, especially future ones, are subject to change.)

# Evaluation of usage and user needs

Example: What should be the size of **main memory** in the compute nodes?

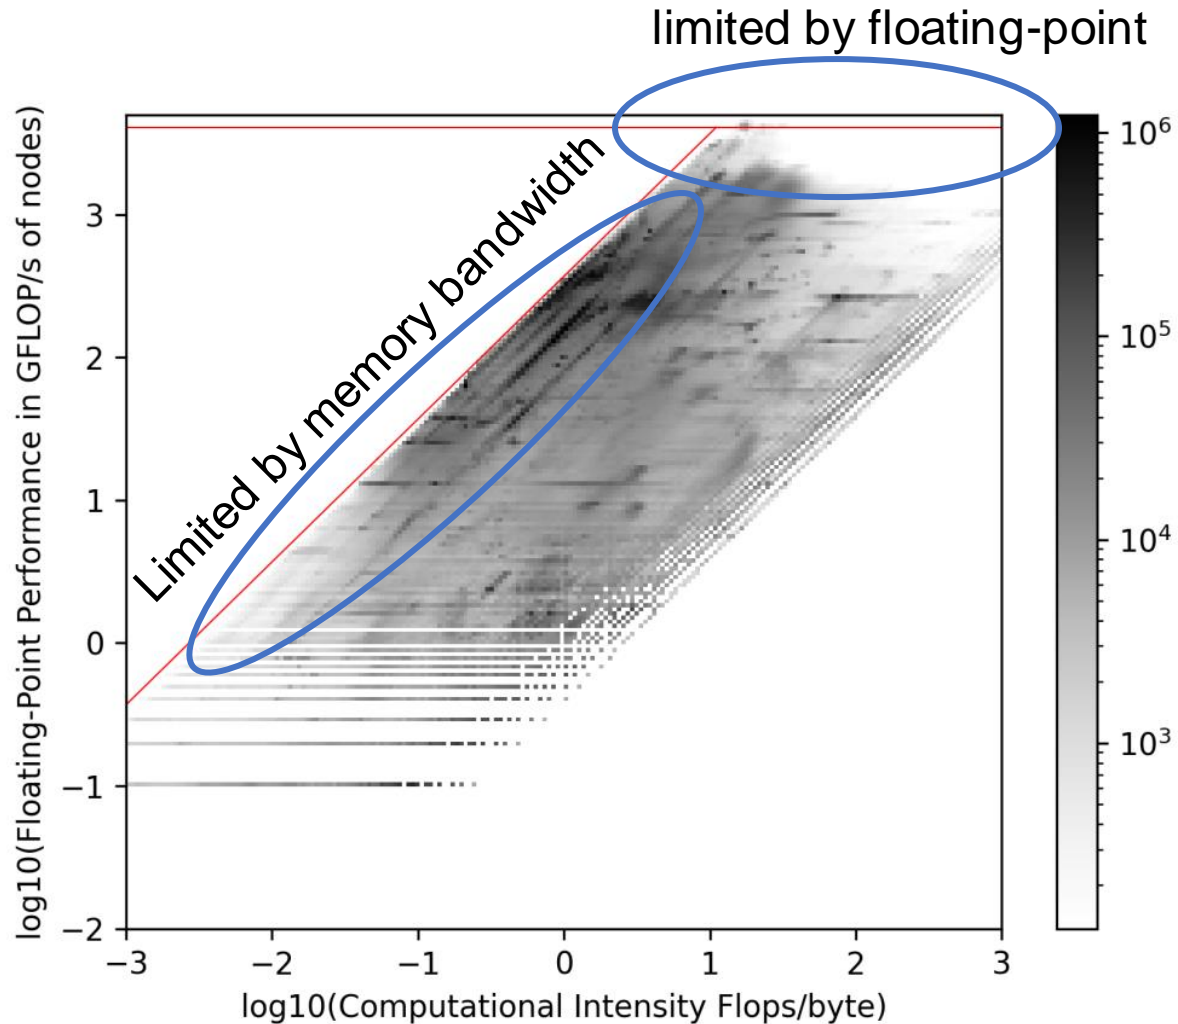


-> 2 GB/core are enough for 77-83% of all compute time used

Resulting requirement: at least **2 GB/core** plus a significant number of **large-memory nodes** (~8 GB/core)

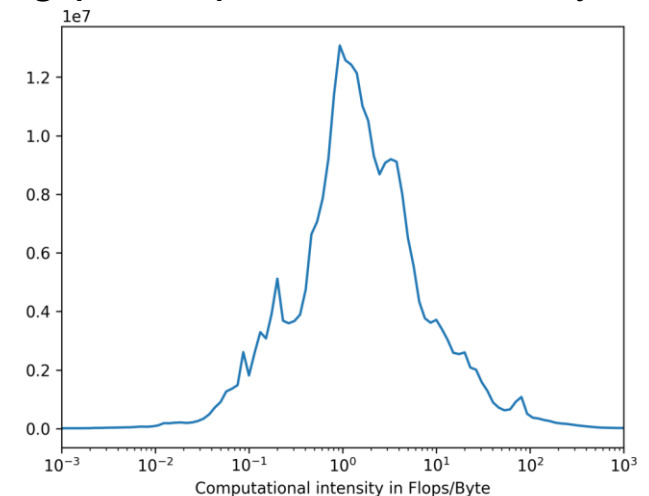
# Evaluation of usage and user needs

Example: How important is **memory bandwidth** versus **floating-point throughput**?



Resulting requirement:  
Memory bandwidth of the  
compute nodes  
is much more important  
than the floating-point  
performance

Histogram of computational intensity  
(Floating-point operations/memory bandwidth)



# Evaluation of usage and user needs

- Which applications/use cases are important?
  - Statistical usage of applications, e.g. from usage of environment modules
  - Requirements in compute-time proposals
- Resulting Benchmarks:
  - All nodes: HPL, HPCG, memory bandwidth, MPI bandwidth/latency
  - CPU-nodes:
    - VASP (CuC\_VdW, SCF for PBE-DFT on 98 atoms with 5 k-points)
    - CP2K (H2O\_512, ab-initio MD for 512 H<sub>2</sub>O molecules in TVZ2P-GTH)
    - QuantumEspresso (GRIR443, DFT for gold surface with 112 atoms and 2 k-points)
  - GPU-nodes:
    - LAMMPS (Reaxff, Hexanitrostilbene crystal with 1.2 Mio. atoms)
    - Pytorch (Bert-large training)
  - Storage benchmarks



# Requirements for the next system

Example: How to avoid the **storage problems** that we had with Noctua 2?

Noctua 2 problems with the storage system:

- When the parallel files system (/scratch) is **used heavily**, the interactivity, i.e., navigating directories, editing files and so on was very slow

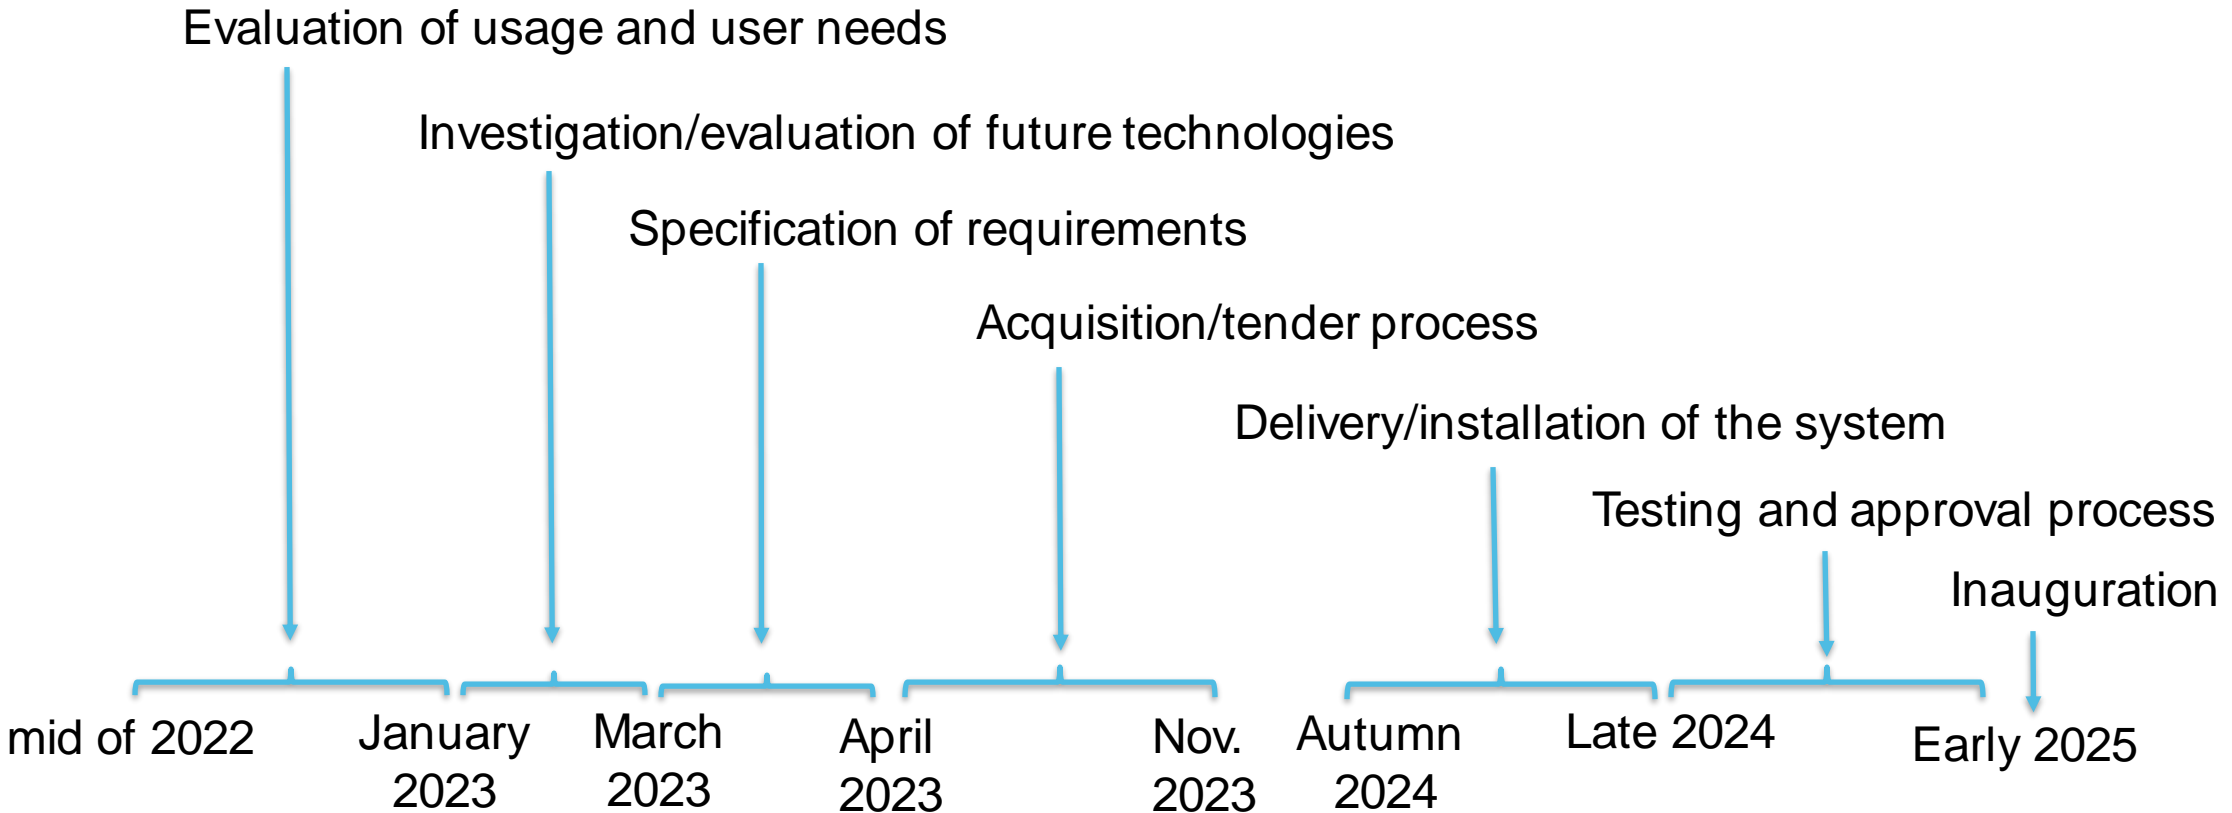


Measures taken in the acquisition to avoid this in the future:

1. More **performance** in general
2. Additional **benchmark case** for storage that simulates the problematic scenario
3. Additional features are required to control **quality of service** and **priority** of IO

# Timeline and Process

## Steps:



(Dates, especially future ones, are subject to change.)

## Guided Tour: Data Center

**Short** tour (~20 Minutes):  
Please gather at the **stairwell**.



Longer tour (~60 Minutes):  
Please gather in the **corridor**.



Thank you all for being here today!