

# WHAT I WISH I KNEW THEN

**Dan Stanzione**

Executive Director, TACC

Associate Vice President for Research, UT-Austin

CI-Compass Workshop

January 2024



# THE LEADERSHIP CLASS COMPUTING FACILITY

- ▶ Our facility is pretty much 100% CI, instead of it being a small component.
- ▶ That said, nearly half the money is acquisition of the datacenter – 20MW of power, liquid cooling for 150KW/server rack.
- ▶ Major components are the physical plant, the compute and storage systems, and people and software.



# “UNEXPECTED OUTCOMES” – SURPRISES ALONG THE WAY

- ▶ First, technology planning 8+ years in advance is sort of a futile exercise.
  - ▶ For instance... I'm not sure “AI” was mentioned in the original proposal (written 2017)
  - ▶ Now LCCF is the “Cornerstone of the NSF NAIRR investment” in 2025.
- ▶ Every vendor roadmap was completely wrong.
  - ▶ Fortunately, we were able to reasonably extrapolate power/cooling requirements.
  - ▶ Mostly \*not\* by believing NVIDIA, Intel or AMD.
  - ▶ The “leaders” in the industry are completely different.
    - ▶ (Mostly the same people, but they’ve all switched companies. Twice.)
- ▶ The trajectory on using the Cloud (or not) has also not really changed – costs continue to shift in favor of building the facility (cloud prices have tripled, effectively, more for AI).

# “UNEXPECTED OUTCOMES” – SURPRISES ALONG THE WAY

- ▶ The most surprising shift from Conceptual to Preliminary design was datacenter strategy.
  - ▶ Moved from an on-campus site built datacenter to a colocation services agreement.
- ▶ Serious increases in construction costs and timelines from campus and (monopoly) municipal utilities.
- ▶ Fortunate convergence of a datacenter building boom, near to our staff, with timing that allowed us to customize the facilities as they are built to meet our needs.



# “UNEXPECTED OUTCOMES” – SURPRISES ALONG THE WAY

- ▶ For what started as a Leadership Computing project became a project where I spent a lot of time thinking about the data lifecycle, and added a lot of pieces to the design around that.
  - ▶ (BTW, spinning hard drives disappeared from the plan along the way).
- ▶ In retrospect, this needed even more attention.
- ▶ Along with the data shift, a lot more interactive computing, persistent services, and other “not leadership computing” modes – we designed that in the hardware plans; software/people costs are higher for these modes though (more than anticipated).
- ▶ ECSS disappearing from the XSEDE ecosystem is likely something we should have compensated for a little more... but that decision was made 4+ years after planning started.

# “BEST PRACTICES”

- ▶ As always, the best practice is never to underestimate the demands of the MREFC process.
  - ▶ Not sure it is theoretically possible to overestimate this.
- ▶ We had a lot of users generating requirements early on.
  - ▶ They were mostly wrong... it is important to filter this.
  - ▶ E.G., almost none of the people at the initial end user workshop talked about AI.
  - ▶ Taking broad trends/directions from this was way more important than trying to deal with the specifics.
    - ▶ If we kept to the details, we would have built a great system for three years ago.



# LESSONS LEARNED

- ▶ Pretty much every time I made a major change that made me feel queasy but was to accommodate a panel or government recommendation I regret not pushing back harder on.
  - ▶ I have found many of the issues some of the stakeholders cared about 5 years ago they have forgotten about or been supplanted by new more urgent issues.
  - ▶ A lot of this is trusting your team's vision of the future, not what past data supports.
    - ▶ For instance, I had no basis in historical data (that I could use for a panel or an auditor) to rationalize a large investment in AI.
    - ▶ Yet, duh.
    - ▶ Turns out we are the subject matter experts.

# WHILE YOU ARE WAITING

- ▶ We have lots of systems available now, or bridge systems available shortly.
  - ▶ Vista will have 600 NVIDIA Grace-Hoppers, available by April (servers hopefully on one of the ships outside. . .).
  - ▶ New Intel and NVIDIA systems this year, new AMD last year.
  - ▶ Archive and online data systems available now.
  - ▶ Frontera still has half a million CPU cores.

STAMPEDE 3

VISTA





# FRONTERA

TACC | NSF | TEXAS

