



Leveraging exascale technology to advance NCAR science within CISL

John Dennis, Garth D'Attilo, Brian Dobbins, Irfan Elahi, Falko Judt, Ben Kirk, Sheri Mickelson, Matthias Rempel, and Jeremy Sauer

November 7, 2023

Exascale Tiger Team



- Goal: determine best course of action with respect to the utilization of exascale technology to advance NCAR's science objectives
- Exascale technology
 - GPU-enablement
 - Lossy data compression
 - Parallelization of post-processing workflows
 - Cloud computing
- NCAR-wide committee
 - Garth D'Attilo (ACOM), John Dennis (CISL), Brian Dobbins (CGD), Irfan Elahi (CISL), Falko Judt (MMM), Ben Kirk (CISL), Sheri Mickelson (CISL), Matthias Rempel (HAO), Jeremy Sauer (RAL)

Exascale Tiger Team (con't)

- Create recommendations based on two sources of information
 - UCAR wide survey
 - Laboratory specific focus groups:
 - “What science would you perform if you could magically eliminate limitation?”
- One overall set of recommendation for NCAR directorate
- Seven laboratory specific recommendations
- Starting place for further discussions

Background on computing trends



- Exascale system versus exascale technology
 - Unclear if NCAR will ever purchase Exascale system
 - NCAR scientists will have access to systems based on exascale technology
- Cost of electricity to run HPC centers is significant
 - After NetZero initiative
 - 2030: 63% of NCAR's carbon emissions will be from supercomputer
 - GPU-based computing is approximately 3x more energy efficient than CPU
- Big science → Need to be energy efficient

Overall findings of ETT committee



- NCAR staff aren't broadly practicing exascale software development citing the lack of:
 - Explicit prioritization from leadership
 - Funding
 - Staff skills and familiarity with new technologies
 - Time to shift to new paradigms while still accomplishing existing workload
 - Meaningful cooperation between scientists and scientific software engineer with respect to modern HPC and exascale practices

Overall findings of ETT committee (con't)



- Novel science outcomes/discoveries can be enabled through the adoption of modern HPC and exascale capabilities
- Prioritization, funding and staffing need must be actionably incorporated into LCPO and NCAR strategy
- Larger cross-lab software development activities are a potential opportunity for collective pursuit of culture change

ETT recommendations for NCAR



- Modify NCAR's and LCPO's strategic plan to explicitly include a response to exascale technology
- Foster a co-design culture where scientific objects represent a collaboration between scientific and software engineering leadership
- Identify internal/external funding to advance specific response to ETT recommendations
- Communicate to staff the importance of exascale technology has on the scientific mission of NCAR
- For a complete list of [recommendations](#)



CISL

Exascale technology and CISL science



- Most of CISL science appears to be exascale friendly
- Vast majority of CISL science performed in collaboration within NCAR and broader community
- Significant experience with exascale technology within laboratory
 - GPU-enablement for modeling
 - GPU supported ML/AI research
 - Lossy-data compression
 - Parallel-post processing
 - Analysis of large data volumes

How were “Exascale-friendly science objectives” identified



- Identify science objectives that can benefit from exascale technologies
- May or may not align with overall institution or laboratory science objectives
- Three categories of science objectives
 - Modest: Can be achieved today with a larger allocation
 - Interesting: Not currently possible/easy, but possible in next 3-10 years with significant effort
 - Unlikely: Not possible in the next 20+ years

Interesting exascale-friendly science objectives (CISL)



Science objectives and enabling exascale technologies

| Description | GPU enablement | Lossy data compression | Parallel post-processing | Cloud computing |
|----------------------------------|----------------|------------------------|--------------------------|-----------------|
| High frequency data assimilation | Yes | Yes | Yes | |

Interesting exascale-friendly science objectives (ACOM)

Science objectives and enabling exascale technologies

| Description | GPU enablement | Lossy data compression | Parallel post-processing | Cloud computing |
|--|----------------|------------------------|--------------------------|-----------------|
| Very large chemical networks ~1000 species | Yes | ? | Yes | |
| LES-scale simulations (several days to weeks) | Yes | Yes | Yes | |
| 3-5 km simulations (months to 1 year) | Yes | Yes | Yes | |

Interesting exascale-friendly science objectives (CGD)

Science objectives and enabling exascale technologies

| Description | GPU enablement | Lossy data compression | Parallel post-processing | Cloud computing |
|---|----------------|------------------------|--------------------------|-----------------|
| High-resolution S2S prediction | Yes | Yes | Yes | |
| PPE / Large ensemble | Yes | Yes | Yes | Yes |
| Convection permitting ensembles | Yes | Yes | Yes | |
| Bio-geochemistry in eddy-permitting ocean | Maybe | Yes | Yes | |

Interesting exascale-friendly science objectives (EOL)



Science objectives and enabling exascale technologies

| Description | GPU enablement | Lossy data compression | Parallel post-processing | Cloud computing |
|-------------------------|----------------|------------------------|--------------------------|-----------------|
| Broad usage of OSSEs | Yes | Yes | Yes | |
| Observational campaigns | | Yes | ? | ? |

Interesting exascale-friendly science objectives (HAO)

Science objectives and enabling exascale technologies

| Description | GPU enablement | Lossy data compression | Parallel post-processing | Cloud computing |
|---|----------------|------------------------|--------------------------|-----------------|
| Resolve earth's ionosphere at mesoscale range | Yes | Yes | Yes | |
| High-resolution space weather prediction | Yes | Yes | Yes | |
| Maximize ChroMag observational data | Yes | Yes | Yes | |
| Global solar coronal magnetic field evolution | Yes | Yes | Yes | |
| Magnetotail reconnection | Yes | Yes | Yes | |

Interesting exascale-friendly science objectives (MMM)

Science objectives and enabling exascale technologies

| Description | GPU enablement | Lossy data compression | Parallel post-processing | Cloud computing |
|----------------------------|----------------|------------------------|--------------------------|-----------------|
| 1-3 km global w/ ensembles | Yes | Yes | Yes | |
| LES @ continental scale | Yes | Yes | Yes | |

Interesting exascale-friendly science objectives (RAL)

Science objectives and enabling exascale technologies

| Description | GPU enablement | Lossy data compression | Parallel post-processing | Cloud computing |
|--|----------------|------------------------|--------------------------|-----------------|
| 1 km Regional downscaling at continental scale | Yes | Yes | Yes | |
| LES modeling in urban environments with lagrangian particles | Yes | Yes | Yes | |
| Multiscale wind turbine/farm modeling | Yes | Yes | Yes | |

Specific challenges for CISL



1. Understanding how institution and community wide software limitations impact the use of exascale technology
 - a. Every 5-year SRAP discussions with community is not sufficiently frequent to adapt to rate of technology change
 - b. Most of NCAR do not understand how technology change will impact their science
 - c. Limited interactions with EOL
2. Very limited adoption of lossy-data compression
3. No GPU-enabled analysis tools (VAPOR & GeoCAT)
4. No GPU-enabled data assimilation capability

Specific recommendations for CISL



1. Understanding how institution and community wide software limitations impact the use of exascale technology
 - a. More frequent interaction with broader NCAR user community about technology trends/opportunities
 - b. Explicit and targeted outreach to EOL
 - c. Increase office hours of CISL staff at ML and CG
2. Very limited adoption of lossy-data compression
 - a. Increase investment in lossy-data compression
3. No GPU-enabled analysis tools
 - a. Develop plan to fully GPU-enable VAPOR and GeoCAT
4. No GPU-enabled data assimilation capability
 - a. Develop plan to GPU-enable DART
 - b. GPU-enable DART



Questions?

dennis@ucar.edu