Lawrence Berkeley National Laboratory

LBL Publications

Title

National Institute of Standards and Technology Requirements (Analysis Report)

Permalink

https://escholarship.org/uc/item/704411nv

Authors

Zurawski, Jason

Schopf, Jennifer

Publication Date

2023-04-21

DOI

10.2172/1971111

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NonCommercial-NoDerivatives License, available at https://creativecommons.org/licenses/by-nc-nd/4.0/

Peer reviewed

eScholarship.org Powered by the California Digital Library
University of California

National Institute of Standards and Technology Requirements Analysis Report

April 21, 2023

Disclaimer

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor the Regents of the University of California, nor the Regents of the University of Texas System, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or the Regents of the University of California or the Regents of the University of Texas System. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or the Regents of the University of California, or the Regents of the University of Texas System.

National Institute of Standards and Technology Requirements Analysis Report

April 21, 2023

The Engagement and Performance Operations Center (EPOC) is supported by the National Science Foundation under Grant No. 1826994.

ESnet is funded by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research. Carol Hawk is the ESnet Program Manager.

ESnet is operated by Lawrence Berkeley National Laboratory, which is operated by the University of California for the U.S. Department of Energy under contract DE-AC02-05CH11231.

This is a University of California, Publication Management System report number LBNL-2001525[1].

[1] https://escholarship.org/uc/item/704411nv

Participants & Contributors

Ron Boisvert, NIST
Tanya Brown-Giammanco, NIST
Artur Chernovsky, NIST
Kamal Choudhary, NIST
Jason Coder, NIST
Ari Feldman, NIST
Jim Fowler, NIST
Ann Leith, NIST
Carol Johnson, NIST
Tom Larason, NIST
Stephen Maxwell, NIST
Randall McDermott, NIST
Duncan McGillivray, NIST
Ken Miller, ESnet
Alan Munter, NIST
Nate Olson, NIST
Andrew Reid, NIST
George Robb, ESnet
Carolyn Rowland, NIST
Jennifer M. Schopf, TACC
Doug Southworth, TACC
Francesca Tavazza, NIST
Jason Zurawski, ESnet

Report Editors

Jason Zurawski, ESnet: zurawski@es.net
Jennifer M. Schopf, TACC, jmschopf@tacc.utexas.edu

Contents

1 Executive Summary

Deep Dive Review Purpose and Process

EPOC uses the Deep Dive process to discuss and analyze current and planned science, research, or education activities and the anticipated data output of a particular use case, site, or project to help inform the strategic planning of a campus or regional networking environment. This includes understanding future needs related to network operations, network capacity upgrades, and other technological service investments. A Deep Dive comprehensively surveys major research stakeholders' plans and processes in order to investigate data management requirements over the next 5–10 years. Questions crafted to explore this space include the following:

- How, and where, will new data be analyzed and used?
- How will the process of doing science change over the next 5–10 years?
- How will changes to the underlying hardware and software technologies influence scientific discovery?

Deep Dives help ensure that key stakeholders have a common understanding of the issues and the actions that a campus or regional network may need to undertake to offer solutions. The EPOC team leads the effort and relies on collaboration with the hosting site or network, and other affiliated entities that participate in the process. EPOC organizes, convenes, executes, and shares the outcomes of the review with all stakeholders.

This Review

In October of 2022, staff members from the Engagement and Performance Operations Center (EPOC) met with researchers and staff from the National Institute of Standards and Technology (NIST) for the purpose of a Deep Dive into scientific and research drivers. The goal of this activity was to help characterize the requirements for a number of campus use cases, and to enable cyberinfrastructure support staff to better understand the needs of the researchers within the community.

Material for this event included the written documentation from each of the profiled research areas, documentation about the current state of technology support, and a write-up of the discussion that took place via e-mail and video conferencing. The case studies highlighted the ongoing challenges and opportunities that NIST has in supporting a cross-section of established and emerging research use cases. Each case study mentioned unique challenges which were summarized into common needs.

The review produced several important findings from the case studies and subsequent virtual conversations:

- The NIST Research Computing Advisory Committee (RCAC) serves as a research focused IT governance organization to define some of the research IT challenges that NIST Laboratories face. The RCAC includes a member from each of the research Laboratories and one from OISM.

- NIST's organizational approach to research IT support is at best federated and at worst disjoint or duplicative. Effort must be made in the coming years to better

address the needs of research groups by a combination of engagement and technology support. Simplifying communication, clearly articulating services, and growing the IT support staff are all critical to the success of NIST.

- NIST storage solutions are a mixture of different technologies and approaches. Currently there is a wide range of storage sizes and mount types, as well as usage expectations. NIST must offer a cohesive strategy to research storage, including longer-term support and continued maintenance.

- NIST does not have a coherent, organization-wide HPC strategy that provides a low-barrier-to-entry resource for researchers. In reality, many research projects across NIST could benefit from the use of HPC resources if there was a lower barrier-to-entry for using computing resources at NIST. If access to HPC resources needs to be expanded, NIST Senior management could consider exploring methods to engage in partnerships with an outside organization that provides these services (e.g. NSF, DOE, NCSA, TACC, or NOAA/Department of Commerce).

- Data mobility in or out of NIST is felt to be the responsibility of the user, and some users can waste time trying to solve technical problems on their own without reaching out for assistance. For example, some IT resources are devoted to maintaining high performance use cases (e.g., Globus), and these solutions are not specific to an area of research. Globus is available across the NIST ecosystem, but many users are not aware of this as an option.

- There is a general lack of technical staff that can serve in a "research IT coordination" role that spans NIST. This role should be a staff member who can address gaps identified in this report and should be normalized to have similar responsibilities across the organization.

- Network challenges negatively impact research projects, and NIST must invest in upgrades to the capacities and services offered within and external to NIST facilities if they want their research teams to remain competitive.

The review produced several recommendations that NIST can consider implementing

- EPOC recommends that NIST invest heavily in developing an architecture for research IT support (sometimes called "data architecture" in R&E parlance) for the research ecosystem, specifically that is, separate from enterprise IT support.

- The Research Computing Advisory Committee (RCAC) presents an opportunity for the NIST labs to get more involved in planning and investment for research IT that impacts their projects and programs. The RCAC should implement a communications campaign across NIST to raise awareness of their efforts.

- It is a widely held perception by researchers that security requirements are opaque and immutable. Each laboratory has an IT Security Officer (ITSO) who should be

able to help the researchers find solutions that balance security and mission requirements. The ITSO role may not be implemented consistently across laboratories.

- OISM and the RCAC should build upon the current capabilities and adjust approaches used to communicate capabilities to the NIST research community, taking into account the hybrid nature of work post-COVID.

- Users of NIST resources often are unaware of the full suite of IT services available to them. It would be beneficial to revisit the methods used to document the resources NIST provides (or doesn't provide).

- Protection of NIST resources (systems, information, etc.) is important, and as such there is a defined process for reviewing the adoption of new technologies. Instruments and control systems being added to networks, as well as software that integrates analysis capabilities, must be reviewed for risks before being fully implemented. It is recommended that OISM better convey the process of this review, the expectations of time, and how the research community can assist to streamline the activities.

- Planning time should be spent for NIST computational and storage support, which will be required to expand in future years

- As planning takes place, resources need to be put in place to support NIST users that are currently performing computationally intensive tasks both on premises as well as externally.

- As data volumes increase, and research collaborations involve more outside parties beyond the NIST boundaries, it is recommended that NIST build on a set of tools to assist with data mobility.

- It is recommended that NIST consider increasing staffing levels to directly assist with cyberinfrastructure technologies, as well as creating a program where proactive assistance is given to researchers that have data-intensive use cases.

- In the time since the initiation of this deep-dive, NIST Networking has acted on the anticipated recommendation regarding establishment of a peering relationship with NOAA N-Wave wherein NOAA is making the NIST Gaithersburg campus a 400G Point-of-Presence in the DC Metro ring. In addition hardware dedicated to the establishment of perfSonar nodes at key points on the NIST network is being procured to enable better monitoring of network throughput for research data transfers.

- Some NIST use cases, for example remote sensors and other forms of field work, can benefit from non-traditional networking approaches. This may include emerging technologies provided via wireless edge and satellite-based networking

(e.g., Starlink). It is recommended that NIST investigate ways this may integrate to some research use cases in a secure and performant manner.

- It is recommended that NIST build on existing capabilities to integrate laboratory instruments into the infrastructure. Currently there are teams in the OISM and in the NIST Labs that have expertise and missions to perform lab automation and move data from instruments to storage. It would be worthwhile to re-examine resource allocations across NIST to see whether there are more optimal ways of satisfying NIST's laboratory instrument integration needs.

2 Deep Dive Findings & Recommendations

The deep dive process helps to identify important facts and opportunities from the profiled use cases. The following sections outline a set of findings and recommendations from the NIST Deep Dive that summarize important information gathered during the discussions surrounding case studies, and possible ways that could improve the cyberinfrastructure support posture for the campus.

The NIST research environment is heterogeneous by necessity to support varied research missions. Each project individually prioritizes IT needs, including the need for high-speed networking, modern computing, data collection, IT security, or high-capacity storage. The NIST Laboratories use a variety of methods to achieve these needs, including some combination of reliance on Office of Information Systems Management (OISM) for IT needs common to most NIST staff, using local IT groups within Laboratories for more specialized needs if available, using reimbursable services from OISM for more specialized needs, or individual research projects performing their own specialized IT when necessary. As a result of these ad hoc methods, research data is spread out across the institution, IT security is labor-intensive, and research IT capabilities are not consistent across the organization.

2.1 Findings

The following sections outline a set of findings EPOC would like to highlight after completing the review process. These are meant to draw attention to positive and negative experiences that research users have had within NIST, and offer a framework for future improvements. Overall, NIST as an organization offers a number of technology support areas for research use cases that meet and exceed capabilities at other organizations. The areas where NIST falls short are mostly due to being behind in terms of adapting to technological changes, as well as preparing to scale to the next level of requirements that will come to fruition in both the near and long terms.These are broken into the following rough categories:

- Planning and Investment
- Non-Technical Challenges
- Technical Challenges
    - Storage
    - HPC
    - Data Mobility
    - Staff Expertise
    - Networking Challenges

2.1.1 Planning and Investment

The Research Computing Advisory Committee (RCAC) is a new research focused IT governance organization that has started to define some of the research IT challenges that NIST Laboratories face. The RCAC includes a member from each of the research Laboratories and one from OISM. The RCAC can create working groups to tackle specific topics or issues. They plan to create a roadmap for research IT, but this will require support from senior management as well as funding to execute on the roadmap. Some of the challenges to the success of the RCAC include the following:

- RCAC is new and doesn't have the recognition across NIST that other established groups do.
- Though many of the research IT challenges were identified as part of NIST's current strategic plan, other challenges facing NIST were prioritized more highly. The RCAC is positioned to focus senior management attention on the research IT challenges and to advocate for the funding needed.
- There should be an effort to look at similar institutions in order to adopt modern practices instead of reinventing the wheel relating to networking, computation, storage, and data mobility.

2.1.2 Non-Technical Challenges

NIST's organizational approach to research IT support is challenging and uneven. Some of the factors that influence NIST's ability to provide first class research IT support to the Laboratories include:

- Relying solely on OISM's overhead-funded support to develop solutions for the Laboratories creates gaps that must be addressed by individual Laboratories.
    - The NIST research environment is heterogeneous, by necessity, which makes it difficult to provide solutions at economies of scale affordable to many research projects.
    - There is a mismatch between researchers' IT needs and expectations and the ability of the OISM to deliver on those needs, as the majority of institutional-level funding is focused on maintaining all of NIST's IT needs including research IT, with relatively little effort available to address specialized needs specific to Laboratory projects.
    - To provide research IT support, the OISM offers fee-for-service solutions that many research projects cannot afford. Local IT support is available within some Laboratories to provide research IT assistance to either augment OISM offerings or to provide additional services.
    - Some researchers undertake their own IT support in order to accomplish the research mission, even when support is available, because the team may lack of awareness that services exist, or the existing service offerings may not fully meet the needs of the project, or there is a distrust of centralized services, or researchers may believe they can provide their own support for lower cost.
- There are multiple communication channels announcing new IT service offerings and service changes, however NIST staff frequently rely on word-of-mouth to learn about services, so many staff may not realize a service exists or is available for use.
- Some of the services OISM provides as fee-for-service to the NIST Laboratories include:
    - Providing special purpose lab support, for example LabVIEW coding, networking, hardware support (e.g., National Instruments hardware, data acquisition, sensors, lab design/experiment design),
    - Supporting many of the HPC clusters for NIST,
    - Managing a governance process for scientific software license purchases and providing support for license availability,

    - Providing application development support and software development infrastructure for Lab projects,
- Fee-for-service approaches can result in Laboratory projects opting out of those services, which creates additional unsupported heterogeneity and further IT support gaps. In addition, other institutions have found that fee-for-service support approaches actually cost the institution more in charging and tracking overhead than the funding that is brought in.
- It is not clear upper management is fully aware of the burden of IT support taken on by research staff, so there has been little higher-level support for fixing problems not seen at the organizational level.
- There is a critical lack of staffing and a growing lack of IT expertise that will only grow worse over time:
    - Resources in the form of key staff members are leaving the institution due to retirement and attrition.
    - Keeping up with cyberinfrastructure trends is challenging.
    - The workload to adapt technology for research use cases is increasing.
    - There are research data storage and data management issues that OISM lacks the resources to address.
    - Some IT expertise is distributed across NIST, but these members of the community are often already overloaded.
    - Recruiting new IT experts is challenging (as is simply attracting qualified candidates), in part due to industry pay scales being much higher.
    - New recruits can find the NIST environment very challenging, due to the culture of allowing many disparate solutions to flourish.
    - Hiring and maintaining senior staff is challenging owing to salary compression and lower pay scales compared to industry.
- There is a need for better communication with the research community about expectations and realities of technology operation.
- A more transparent process for how technology (software, hardware, etc.) is evaluated by OISM would be beneficial to the research community. A number of use cases expressed frustration that tools or services were disallowed, despite their value to the process of science.
- The decisions made for IT investments haven't always captured the needs of research. When money is given to OISM to implement a research IT solution, it is often one-time funding that doesn't address staffing, maintenance, and refresh needs over time.
- OISM can be heavily bureaucratic, which can lead to less agility for fast-paced research IT needs.
- If a researcher doesn't like existing solutions, they can create their own. Sometimes, those same researchers assume that OISM (or local Laboratory IT support) will support the new service in perpetuity without a prior agreement.
- Because of the lack of expert involvement, non-experts take on the role of research IT support, which leads to additional problems, such as the security of the services. This in turn causes reputational risk for NIST.

- Use of short-term support (e.g., post docs) can lead to lack of continuing institutional knowledge for technical approaches, which then increases the burden for existing staff.

2.1.3 Technical Challenges

2.1.3.1 Storage

NIST storage solutions are a mixture of different technologies and approaches. Currently there is a wide range of storage sizes and mount types, as well as usage expectations. A given storage solution may depend heavily on the location of a research team and the tools that specific research team are trying to use.

- NIST supports several different storage systems.
    - The OISM-supported on-premise file servers in Boulder and Gaithersburg are provided for research teams to use. Space is allocated a few TB at a time via a quota system. When a quota is reached, more can be requested, but there is a delay in fulfilling requests and a manual approval process. Requesting large amounts of storage on this system is extremely challenging.
    - NIST has access to several cloud storage solutions, such as AWS S3, Google Drive, Box, and OneDrive/Sharepoint. Each has a different cost model, however there is not a storage quota.
        - NIST currently has high-speed, direct connections to only AWS; storage available via the other cloud services is accessible at NIST's public internet speeds.
        - Moving large volumes (TB or above) in and out of those services is constrained both by NIST public internet bandwidth limitations as well as by the service provider's data exchange throughput.
        - OISM's marketing of these cloud solutions diminished after their introductions; the costs and availability are not universally known. Researchers new to NIST may not realize these services are available.
        - Other than cloud storage for dissemination of datasets to the public, cloud storage solutions are not centrally funded, so researchers who want to use this approach must allocate research money to cover the expenses.
- Research staff are expected to identify their research data storage location, preservation, archiving plans in a formal data management plan. This may or may not be appropriate for a researcher to do in isolation. NIST's Open Access to Research project in conjunction with the OISM and the NIST Library offer help with data management planning and provide a software tool that guides the creation of the Data Management Plan (DMP). Laboratory IT Security Officers and Laboratory-specific IT groups may also offer help with the selection of research data storage/backup mechanisms that go into data management planning. Nevertheless research data management practices are a work in progress, and there is a large variance in practice from group to group. Data security may not be

part of a researcher's priorities. Currently a researcher can state in their project DMP that the data is stored on a local hard drive in the lab and that is an acceptable answer.

- There is some OISM support for more advanced data transport and storage management tools, such as Globus and Starfish but the OISM lacks the resources needed to foster widespread adoption hence they are only used by a few projects. There is no well-defined set of data movement tools or portal-based systems that currently interface with the diverse set of backend storage existing in the NIST Labs.

2.1.3.2 HPC

NIST does not have a coherent, organization-wide HPC strategy that provides a low-barrier-to-entry resource for researchers. In reality, HPC could benefit many research projects across NIST if there was a lower barrier-to-entry for using HPC resources at NIST.

- There is an interest in creating an "off-ramp" to more sophisticated HPC capabilities. In practice this could mean:
    - Forming relationships with off-site HPC centers
    - Establishing network paths to make data transfer seamless
    - Having tools/software available to ease deployment and analysis
    - Hiring personnel to consult with researchers and help them with code development
- NIST computational solutions are a mixture of different technologies and approaches.
- Currently researchers may have access to and attempt to use a number of approaches:
    - Workstations and laptops, located within the laboratory environment, to perform rudimentary analysis
    - Lab/team clusters and/or high-end workstations/servers maintained by research groups
    - Lab/team clusters maintained by OISM as a service to research groups
    - Institutional HPC maintained by OISM for NIST-wide use but paid for by a subset of NIST Laboratories
    - Cloud computing via approved institutional sources (e.g., AWS)
    - Use of R&E-based HPC, and HTC outside of NIST (e.g. XSEDE/ACCESS, DOE HPC, NOAA, OSG, Chameleon, Jetstream)
- Currently, no single team provides support to researchers/staff who are new to HPC and may need help figuring out which HPC resource to use and how to get started.
- Documentation on how (and where) NIST users utilize external computation resources would help a number of research groups.
- Current computational capabilities at NIST are effective for some cases, but are unlikely to be able to scale to meet future needs. In particular, the use cases profiled in this document shows potential growth and a need to increase

technology support. Computational capabilities are also needed to meet the needs of initiatives, such as CHIPS (https://www.nist.gov/chips). It is suggested that NIST invest in programs to ensure:

- Additional assistance and documentation in getting new models running on available resources;
- Additional software development resources, for example to assist with increasing parallelism;
- Availability of more cores/CPUs for parallelism in workload;
- Availability of faster CPUs and GPUs to handle more intense workloads;
- Availability of GPUs for emerging work in AI/ML;
- More storage at NIST to support data ingest during computation, and larger outputs from processing workloads;
- Ability to share results from computation with internal and external collaborators.

2.1.3.3 Data Mobility

Data mobility in or out of NIST is felt to be the responsibility of the user, and some users waste time trying to solve technical problems on their own without reaching out to OISM or local Lab IT support to see if there are existing solutions available. Some IT resources are devoted to maintaining high performance use cases (e.g., Globus), and these solutions are not specific to an area of research. Globus is available across the NIST ecosystem, but many users are not aware of this as an option.

- NIST could benefit from creating a “standard” data ingest pipeline approach that can be adopted by a number of use cases. In practice, this may include:
  - Understanding the interface between the instrument/sensor/data generation component
  - Linking this to an efficient research network
  - Using standard interfaces to high-speed storage and data transfer tools
  - Creating an efficient migration path to storage and computation
  - Building an ecosystem of data sharing on top of the storage layer that is separate from the tools of acquisition and analysis.
- An important aspect of this is the 'Data Plumbing' project, which is a collection of tools MML has written to provide automated data transfer from laboratory instruments (or intermediate data collection locations) to centrally managed storage. Data Plumbing is an effective solution to the barriers that prevent data transfer from individual laboratories to a central storage solution, but it is a manual and very labor-intensive process to connect each individual instrument. To date, MML has connected about 80 instruments, which represents less than 20% of the MML instrument inventory. There are many other parts of NIST that could benefit from this service, if it was available more widely.
- IT security remains a source of some friction in planning data workflows. NIST could benefit from creating enclaves that deal with “secure” computing and storage that are separate from the more general research or enterprise environments.

- Researchers often use portable drives sent through the mail to share data (GB, TB) with external collaborators instead of using the network due to performance problems.
- Data made available to the public suffers from several key problems:
    - Data longevity is increasing, where old data is still valuable and must be retained, cataloged, and made available much longer than in past research flows.
    - Data volumes, and experimental dataset sizes, are increasing. This means additional storage will be needed to support research, either supplied by NIST or through cloud providers.
    - Backups, even on cold storage, are still necessary.
    - With increased data volumes, and a population of external users looking to access this data, the need for robust and scalable systems will increase. This will mean creating new approaches to serve the data, ways to streamline delivery across the wide area network, and adapting the existing tools to changing technologies.
- NIST laboratories have on-going issues with data transport. As HPC gets more powerful and scientific instruments produce more data, it is increasingly difficult to manage data sets, including for basic activities such as data protection. NIST, as an institution, has an excellent network roadmap with plans for more capacity in the near future, which will certainly help with this.

2.1.3.4 Staff Expertise

There is a general lack of technical staff that can serve in research software engineering[2] roles that span NIST. This can be defined as staff to address the following gaps:

- Cybersecurity reviews of data, to better classify some or all of research data sets, ideally through the existing ITSO role;
- Research software engineering support, that has expertise in writing, running, and maintaining HPC/HTC code, as well as porting software to environments at NIST or to environments external to NIST; and
- Science engagement support that offers assistance in automating workflows, e.g., better integrating instrumentation, data acquisition, analysis, and data transport from research networks to other networks.
- Addressing researchers' lack of expertise in code development that improves research workflows, or in data wrangling, or in integrating open source software components, or in creating web-based services, etc.

This role should be normalized to have similar responsibilities across the organization.

2.1.3.5 Networking Challenges

Network challenges negatively impact research projects. This includes:

- OISM's networking can struggle when supporting data leaving the NIST ecosystem. For example:

[2] https://www.computer.org/publications/tech-news/research/us-rse-supporting-the-research-software-engineer

    - Some users have resorted to physical delivery of data to or from remote collaborators; it is not clear whether this is due to a lack of knowledge about existing NIST services that can facilitate those data transfers (e.g., Globus, Box, Google Drive), whether there are shortcomings with NIST's configuration of those services, or whether the issue is the first (last) 100 meters of networking on a NIST campus.
    - Integrating new data-intensive devices is possible, but would be better engineered if researchers could consult with lab and/or central IT support before deployment.
    - Mixing enterprise/administrative use cases with those that are more aligned with research on the same network can cause friction with data transfers. It would be useful to investigate if separating the enterprise and research-use networks is a possible solution.
    - VPN latency reduces responsiveness for remote users.
- OISM's network and information security adopts approaches that do not follow practices aligned with others in the research and education (R&E) community. For example, there are several layers of firewalls that are utilized for network traffic that enters or leaves the NIST perimeter as required for Federal civilian agencies.
- OISM staff are interested in implementing commonly used R&E cyberinfrastructure technologies, including perfSONAR, Science DMZs, and Data Transfer Nodes.
- OISM staff indicate that we have the technical means to increase peering to support R&E and cloud traffic external to the NIST campuses but not the funding or the evidence that these things are required yet.

2.1.3.6 Science Use Case Findings

The Science Uses Cases highlighted a number of specific findings:

- Several use cases regularly use computation (HPC, Cloud) for research activities. The use takes the form of:
    - NIST/private workstations and clusters
    - NIST institutional computing
    - Government-supported R&E computational resources (DOD, DOE, NSF)
    - Cloud (AWS, etc.)
- Computational use takes the form of simulation (done to validate models, or train AI/ML), and analysis (e.g., running models against data).
- CTCMS was interested in identifying a partnership with an "off-site HPC center" so they could transition some users to that resource more easily and free up institutional resources for other users. This might be an approach to be evaluated with the availability of ACCESS (XSEDE) resources and others.
- Remote access to NIST resources only available through the VPN can be problematic with users reporting slowdowns for interactive uses, for example uploads from the field and visualization.
- OISM does not offer assistance to create automated data transfer pipelines. The MML IT Service Team created a tool for use in MML but this solution is not available in other organizations.

- The majority of the NIST use cases studied experienced shortcomings in network connectivity performance. There are large gaps in supporting the data transport between instruments and analysis resources, as well as in sharing data with collaborators. Examples where current networking were insufficient included:
    - Several teams had created manual workarounds, instead of using the networks, that were adding significant delays to their research progress.
    - High-speed network links between NIST and external collaborators were needed to support several use cases. Globus may be a solution but many researchers do not know that NIST has Globus.
    - Some projects resort to sneaker-net (using portable storage devices) to move data, both inside of NIST and to external collaborators.
    - Data mobility is still being handled by older tools (e.g., SCP, FTP), and could benefit from adoption of high-performance solutions such as Globus - which is already available at NIST.
- Field research could benefit from mobile networking solutions that work in more locations than commercial cellular connectivity. The wider availability of satellite solutions (e.g., StarLink) has made it more feasible to equip field researchers with portable networking that is capable of reaching networking speeds beyond 50Mbps in some cases.

2.2 Recommendations

The EPOC Deep Dive process helps to identify important facts and opportunities from the profiled use cases. The following outlines a set of findings from the NIST Deep Dive that summarize important information gathered during the discussions surrounding case studies, and possible ways that could improve the cyberinfrastructure support posture for the campus.These are broken into the following rough categories:

- Planning and Investment Recommendations
- Non-Technical Recommendations
- Technical Recommendations
    - Storage Recommendations
    - HPC Recommendations
    - Data Mobility Recommendations
    - Staff Expertise Recommendations
    - Networking Recommendations
    - Science Use Case Recommendations

The following sections outline a set of recommendations EPOC suggests after completing the review process. These are non-binding, but can assist NIST in preparing for the future of their research activities based on the trajectory of the science and technology discussions. Many of the recommendations require investment from high levels of the organization, while others are less resource intensive but will require commitment to change procedure and culture both from the top down and the bottom up. Overall, NIST as an organization offers a number of technology support areas for research use cases that meet and exceed capabilities at other organizations. The areas where NIST falls short are in adapting to technological changes and in preparing to scale to the next level of requirements that will emerge in the near and longer terms.

2.2.1 Planning and Investment Recommendations

EPOC recommends that NIST invest heavily in developing an architecture for research IT support (sometimes called "data architecture" in R&E parlance) for the research ecosystem specifically, separate from enterprise IT support. Some potential areas this would address include:

- An externally facing Science DMZ.
- Generalizing the approach to onboarding instruments (e.g., via "data plumbing") so that they can efficiently use computing and storage resources.
- Architecture that explicitly accounts for workflows from data acquisition to storage to modeling/analysis (compute) to collaboration to dissemination.
- Potential expansion of computing capabilities, either on premise or through external partnerships, as appropriate.
- Creating example workflows for HPC for researchers to follow.

The Research Computing Advisory Committee (RCAC) presents an opportunity for the NIST labs to get more involved in planning and investment for research IT that impacts their projects and programs. The RCAC should implement a communications campaign across NIST to raise awareness of their efforts. Along with that, they should also do the following:

- Create a roadmap for research IT at NIST that can include priorities and identify funding needs.
- Work with OISM to identify solutions that can be built together to satisfy the needs of NIST researchers.
- EPOC recommends that RCAC create a comprehensive program to engage with the users of NIST services:
    - Creating a system to evaluate current and future technology needs;
    - Setting regular check ins with NIST research areas;
    - Publishing results for others to view and comment on; and
    - Holding regular meetings to share findings.

2.2.2 Non-Technical Recommendations

- It is a widely held perception by researchers that security requirements are opaque and immutable. Each laboratory has an IT Security Officer (ITSO) who should be able to help them find solutions that balance security and mission requirements, however researchers may not know to ask their ITSO for help with security that impedes their work.
- OISM and the RCAC should build upon the current capabilities and adjust approaches used to communicate capabilities to the NIST research community, taking into account the hybrid nature of work post-COVID. This may include:
    - Updated and adapted documentation;
    - Direct engagement with researchers on a semi-regular basis;
    - "Office hours", where people can ask questions or get help; and
    - "Show and tell", where researchers can advocate for the ways that technology support has addressed their workflow needs/concerns.

- Users of NIST resources often are unaware of services offered. It would be beneficial to revisit the methods used to document the resources NIST provides (or doesn't provide). This should include:
    - ***Storage:*** An overview of the services NIST provides at the institution level, or remotely through partnerships, and the primary ways users can interact (e.g., mounts, APIs, portals).
    - ***Computation:*** An overview of the services NIST provides at the institution level, or remotely through partnerships, and the primary ways users can interact (e.g., ways to get accounts, how to access, how to find documentation to provision and use computational resources).
    - ***Data Mobility:*** Suggested tools and procedures to migrate data within and external to the NIST campus.
    - ***Information and Network Security:*** Ways to request reviews of data and tools, setting expectations on timelines and ways the NIST team can assist researchers in evaluating and addressing concerns surrounding controlled data and tools.
    - **Software and Workflow Adaptation:** Information on how NIST staff can assist researchers with adapting tools and workflows to use the advanced services that are available (e.g., storage, computation, and networking).
- Protection of NIST resources (systems, information, etc.) is important, and as such there is a defined process for reviewing the adoption of new technologies. Instruments and control systems being added to networks, as well as software that integrates analysis capabilities, must be reviewed for risks before being fully implemented. It is recommended that OISM better convey the process of this review, the expectations of time, and how the research community can assist to streamline the activities.

2.2.3 Technical Recommendations

2.2.3.1 Storage Recommendations

- NIST will need to prioritize investment in storage solutions for research data in future years.
- The EPOC team recommends adopting an on-premise model with a broad set of capabilities, as shown in the Figure 2.2.3.1.1:
    - Archival storage (e.g., tape) where seldom used data sets can be backed up. NIST should evaluate the cost benefit analysis if this should be on premises or in the cloud.
    - Hot storage (e.g., disk, group / parallel filesystem mounts) where data that should be accessible can live. This should be readily available to users via LAN/VPN mounts for easy access, but also an option for sensors or other instrumentation to send data to directly.
    - Fast storage (e.g., DTN and applications) that can facilitate fast writing directly from the WAN or LAN. This does not need to be as large as fast storage, with the intention that it can be a buffer between use cases and the hot storage.

- In addition, NIST researchers will also need improved support for those that choose to store in commercial clouds.

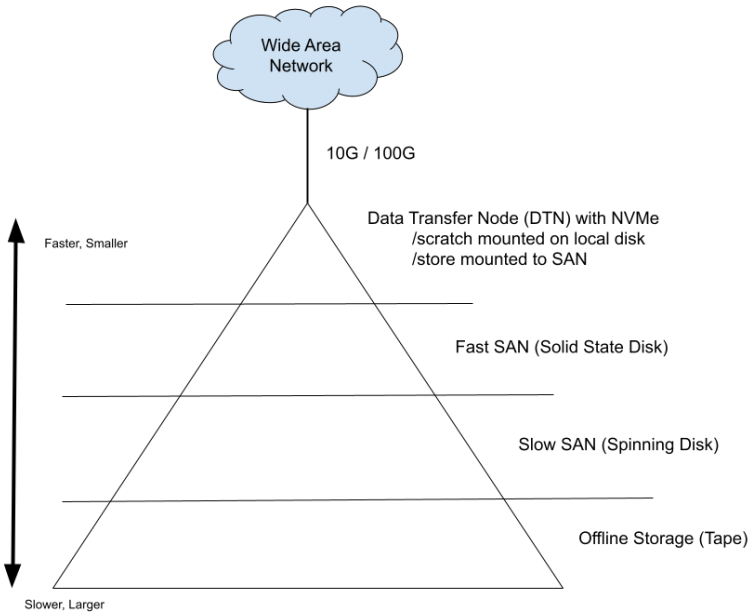

Figure 2.2.3.1.1: Data Storage Pyramid

2.2.3.2 HPC Recommendations

- NIST computational support will be required to expand in future years. This should include:
    - Assessing existing clusters to determine the right mix of central and local;
    - Moving to an institutional model to support mid-scale computational jobs on premise;
    - Providing support for researchers who need compute but do not know how to get started or have trouble taking advantage of existing clusters;
    - Adopting approaches that partner with HPC and HTC centers provided by the R&E community; and
    - Having support for those that choose to compute in commercial clouds.
- NIST users currently perform a number of computationally intensive tasks that are done on premises as well as externally. A future strategy for computation should involve:
    - Evaluating the major forms of computational needs at NIST and remotely:
        - Single and multi-processor use cases best performed on a workstation where fast-feedback and/or visualization is required.
        - High Performance computing (HPC) use cases that require clusters of machines and can be adapted to use tools such as MPI.
        - High Throughput computing (HTC) use cases that can use network-connected computers (e.g., grids) to split workloads that are not highly parallelizable.

        - Compute environments tailorable to individual researcher needs provided by external sources facilitated centrally for NIST-wide use (for example, JetStream2, which is an NSF-funded cloud with machines located at IU,TACC, etc.).
        - Off premise resources not operated by NIST.
    - Figuring out a NIST-wide strategy to provide the resources at NIST, or via partnerships with other entities (E.g., NSF funded facilities like TACC or ACCESS (formerly XSEDE), the DOE HPC facilities, or commercial providers), or a combination of several offerings.
    - Providing a workforce that can adapt computational workloads to the aforementioned systems to unburden the researchers from having to do these tasks.

2.2.3.3 Data Mobility Recommendations

- EPOC can work with NIST on data mobility approaches, including measuring performance and adopting data portals.
- As data volumes increase, and research collaborations involve more outside parties beyond the NIST boundaries, it is recommended that NIST build on a set of tools to assist with data mobility. These tools should be able to integrate into existing research workflows with powerful APIs, utilize existing and planned storage resources, and be capable of reaching other R&E based repositories, including HPC centers and instrumentation facilities.

2.2.3.4 Staff Expertise Recommendations

- It is recommended that NIST consider increasing staffing levels to directly assist with cyberinfrastructure technologies, as well as creating a program where proactive assistance is given to researchers that have data-intensive use cases. Possible use cases include:
    - Engineers capable of taking in requirements, as well as offering focused IT support to unburden research staff from having to perform some aspects of software integration between instruments, computation, storage, and data mobility aspects of a workflow.
    - Staff who can assist researchers in “picking the right tool for the job” and assisting them in starting use of those tools.
    - Researchers need coding support. In order to advance their research, they may need systems support to adapt existing workflows to operate on HPC resources at large national facilities (DOE, NSF, etc.), or cloud computing providers.
    - Gathering requirements for field research, and offering solutions that can scale for the desired workflows.
    - Ongoing technical assistance for problems related to data transfer and storage. This should focus on areas of friction, including:
        - Data transfer between collaborators using modern hardware and software tools;
        - Discussions about security profiles matching research use cases; and

    - Working to establish remote connectivity for field sites.

2.2.3.5 Networking Recommendations

- NIST Networking should continue to monitor WAN demand and continue to leverage its existing relationship with NOAA's N-Wave organization to provide WAN capacity in anticipation of evolving demand.
- EPOC will work with NIST to educate and adopt perfSONAR monitoring for network performance.
- Some NIST use cases, for example remote sensors and other forms of field work, can benefit from non-traditional networking approaches. This may include emerging technologies provided via wireless edge and satellite-based networking (e.g., Starlink). It is recommended that NIST investigate ways this may integrate to some research use cases in a secure and performant manner.

2.2.3.6 Science Use Case Recommendations

- It is recommended that OISM engage with the research community on solutions that can scale with the Virtual Private Network (VPN) infrastructure present on campus. This engagement should acknowledge the performance and accessibility challenges that have been identified, and offer other technical solutions that may facilitate research workflows.
- It is recommended that OISM work with the RCAC and other stakeholders to improve on existing communication approaches for alerting the research community to pending technology upgrades. This includes, but is not limited to, networks, computation, storage, and network peering.
- It is recommended that OISM and ADLP work together to adopt a regular refresh strategy with funding for shared IT resources and staffing supporting the NIST scientific mission, including but not limited to networking, scientific software, compute, and storage, so that they can better prepare for increases in capability and capacity that are driven by the research use cases. RCAC can be a good conduit for OISM to work with ADLP on needs.
- It is recommended that ADLP and OISM work jointly to develop abilities to better integrate laboratory instruments. This could include generalizing the "Data Plumbing" approach to a wider number of use cases at NIST. This may require expanding the scope, hardware and software support, and team that is able to support this service.

3 Process Overview and Summary

3.1 Campus-Wide Deep Dive Background

Over the last decade, the scientific community has experienced an unprecedented shift in the way research is performed and how discoveries are made. Highly sophisticated experimental instruments are creating massive datasets for diverse scientific communities and hold the potential for new insights that will have long-lasting impacts on society. However, scientists cannot make effective use of this data if they are unable to move, store, and analyze it.

The Engagement and Performance Operations Center (EPOC) uses the Deep Dives process as an essential tool as part of a holistic approach to understand end-to-end research data use. By considering the full end-to-end research data movement pipeline, EPOC is uniquely able to support collaborative science, allowing researchers to make the most effective use of shared data, computing, and storage resources to accelerate the discovery process.

EPOC supports five main activities

- Roadside Assistance via a coordinated Operations Center to resolve network performance problems with end-to-end data transfers reactively;
- Application Deep Dives to work more closely with application communities to understand full workflows for diverse research teams in order to evaluate bottlenecks and potential capacity issues;
- Network Analysis enabled by the NetSage monitoring suite to proactively discover and resolve performance issues;
- The Data Mobility Exhibition and associated work with our simplified portal to check transfer times against known performant end points;
- Coordinated Training to ensure effective use of network tools and science support.

Whereas the Roadside Assistance portion of EPOC can be likened to calling someone for help when a car breaks down, the Deep Dive process offers an opportunity for broader understanding of the longer term needs of a researcher. The Deep Dive process aims to understand the full science pipeline for research teams and suggest alternative approaches for the scientists, institutional IT support, and national networking partners as relevant to achieve the long-term research goals via workflow analysis, storage/computational tuning, identification of network bottlenecks, etc.

The Deep Dive process is based on an almost 15-year practice used by ESnet to understand the growth requirements of Department of Energy (DOE) facilities[3]. The EPOC team adapted this approach to work with individual science groups through a set of structured data-centric conversations and questionnaires.

3.2 Campus-Wide Deep Dive Structure

The Deep Dive process involves structured conversations between a research group and relevant IT professionals to understand at a broad level the goals of the research team and how their infrastructure needs are changing over time.

The researcher team representatives are asked to communicate and document their requirements in a case-study format that includes a data-centric narrative describing the science, instruments, and facilities currently used or anticipated for future programs; the advanced technology services needed; and how they can be used. Participants considered three timescales on the topics enumerated below: the near-term (immediately and up to

[3] https://fasterdata.es.net/science-dmz/science-and-network-requirements-review

two years in the future); the medium-term (two to five years in the future); and the long-term (greater than five years in the future).

The case study process tries to answer essential questions about the following aspects of a workflow:

- ***Research & Scientific Background***—an overview description of the site, facility, or collaboration described in the Case Study.
- ***Collaborators***—a list or description of key collaborators for the science or facility described in the Case Study (the list need not be exhaustive).
- ***Instruments and Facilities: Institutional & Non-Institutional***—a description of the network, compute, instruments, and storage resources used for the science collaboration/program/project, or a description of the resources made available to the facility users, or resources that users deploy at the facility or use at partner facilities.
- ***Process of Science***—a description of the way the instruments and facilities are used for knowledge discovery. Examples might include workflows, data analysis, data reduction, integration of experimental data with simulation data, etc.
- ***Computation & Storage Infrastructure: Institutional & Non-Institutional***—The infrastructure that is used to support analysis of research workflow needs: this may be institutional storage and computation, it may be private, it may be shared, or it may be public (commercial or non—commercial).
- ***Software Infrastructure***—a discussion focused on the software used in daily activities of the scientific process including tools that are used at the institution or remotely to manage data resources, facilitate the transfer of data sets from or to remote collaborators, or process the raw results into final and intermediate formats.
- ***Network and Data Architecture***—description of the network and/or data architecture for the science or facility. This is meant to understand how data moves in and out of the facility or laboratory focusing on institutional infrastructure configuration, bandwidth speed(s), hardware, etc.
- ***Resource Constraints***—non-exhaustive list of factors (external or internal) that will constrain scientific progress. This can be related to funding, personnel, technology, or process.
- ***Outstanding Issues***—Listing of any additional problems, questions, concerns, or comments not addressed in the aforementioned sections.

At a physical or virtual meeting, this documentation is walked through with the research team (and usually cyberinfrastructure or IT representatives for the organization or region), and an additional discussion takes place that may range beyond the scope of the original document. At the end of the interaction with the research team, the goal is to ensure that EPOC and the associated CI/IT staff have a solid understanding of the research, data movement, who's using what pieces, dependencies, and time frames involved in the Case Study, as well as additional related cyberinfrastructure needs and concerns at the organization. This enables the teams to identify possible bottlenecks or areas that may not scale in the coming years, and to pair research teams with existing resources that can be leveraged to more effectively reach their goals.

3.3 NIST Deep Dive Background

In October of 2022, staff members from the Engagement and Performance Operations Center (EPOC) met with researchers and staff from the National Institute of Standards and Technology (NIST) for the purpose of a Deep Dive into scientific and research drivers. The goal of this activity was to help characterize the requirements for a number of campus use cases, and to enable cyberinfrastructure support staff to better understand the needs of the researchers within the community.

This review includes case studies from the following campus stakeholder groups:

- Fire Modeling and the, National Fire Research Laboratory (NFRL)
- Disaster and Failure Studies (DFS) Program (includes National Construction Safety Team (NCST) and National Windstorm Impact Reduction Program (NWIRP))
- The Remote Sensing Laboratory (RSL)
- The Joint Automated Repository for Various Integrated Simulations (JARVIS)
- The National Advanced Spectrum and Communications Test Network (NASCTN)
- Genome in a Bottle (GIAB)
- EL Data, Security, and Technology (ELDST)
- Office of Information Systems Management (OISM)
- Material Measurement Laboratory IT Service Team
- Center for Theoretical and Computational Materials Science (CTCMS)

Material for this event included the written documentation from each of the profiled research areas, documentation about the current state of technology support, and a write-up of the discussion that took place via e-mail and video conferencing.

The case studies highlighted the ongoing challenges and opportunities that NIST has in supporting a cross-section of established and emerging research use cases. Each case study mentioned unique challenges which were summarized into common needs.

On October 31st and November 3rd 2022, staff from NIST and EPOC participated in virtual discussion on the use cases and potential next steps to develop a set of sustainable approaches to provide technological support.

Several themes came out of the discussion, in part stemming from the way NIST is organized and traditional boundaries (often silos) between organizations at NIST:

- A primary concern was the apparent disconnect between OISM's broad array of services to NIST and the research IT needs of the NIST Laboratories. OISM's mandate is to support NIST overall - which results in many service offerings that are not tailored towards research IT.
- OISM's resources are not primarily tasked to address the needs of research staff at this time.
- Staffing levels for IT support, both support for enterprise and research IT, were acknowledged as a problem.

EPOC and NIST staff spent a lot of time discussing the nature of storage in the role of a research workflow, and how an architecture can be defined to address some of the high-level needs:

- Localized storage (typically maintained by a researcher or group) that connects to instruments and computers. Meant to be fast, but limited, and facilitates immediate research needs (e.g., not for long-term storage, or sharing, maybe not backed up, who knows this is where the data resides when someone retires)
- Institutional storage maintained for all users. Slower than a localized solution, but scalable to support internal sharing, archiving, and mobility to other resources (e.g., instruments, computation, etc.), and can be backed up, is discoverable
- Shareable storage, either maintained at the institution or by a cloud provider to facilitate collaboration with external entities.

3.4 Organizations Involved

The Engagement and Performance Operations Center (EPOC) was established in 2018 as a collaborative focal point for operational expertise and analysis and is jointly led by the Texas Advanced Computing Center (TACC) and the Energy Sciences Network (ESnet). EPOC provides researchers with a holistic set of tools and services needed to debug performance issues and enable reliable and robust data transfers. By considering the full end-to-end data movement pipeline, EPOC is uniquely able to support collaborative science, allowing researchers to make the most effective use of shared data, computing, and storage resources to accelerate the discovery process.

The Energy Sciences Network (ESnet) is the primary provider of network connectivity for the U.S. Department of Energy (DOE) Office of Science (SC), the single largest supporter of basic research in the physical sciences in the United States. In support of the Office of Science programs, ESnet regularly updates and refreshes its understanding of the networking requirements of the instruments, facilities, scientists, and science programs that it serves. This focus has helped ESnet to be a highly successful enabler of scientific discovery for over 25 years.

The Texas Advanced Computing Center (TACC) at the University of Texas at Austin designs and deploys the world's most powerful advanced computing technologies and innovative software solutions to enable researchers to answer complex questions to help them gain insights and make discoveries that change the world. TACC's environment includes a comprehensive cyberinfrastructure ecosystem of leading-edge resources in high performance computing (HPC), visualization, data analysis, storage, archive, cloud, data-driven computing, connectivity, tools, APIs, algorithms, consulting, and software.

National Institute of Standards and Technology (NIST) was founded in 1901, and serves as a non-regulatory federal agency within the U.S. Department of Commerce. NIST's mission is to promote U.S. innovation and industrial competitiveness by advancing measurement science, standards, and technology in ways that enhance economic security and improve our quality of life.

4 NIST Case Studies

NIST presented a number of use cases during this review. These are as follows:

- Fire Modeling and the, National Fire Research Laboratory (NFRL)
- Disaster and Failure Studies (DFS) Program (includes National Construction Safety Team (NCST) and National Windstorm Impact Reduction Program (NWIRP))
- The Remote Sensing Laboratory (RSL)
- The Joint Automated Repository for Various Integrated Simulations (JARVIS)
- The National Advanced Spectrum and Communications Test Network (NASCTN)
- Genome in a Bottle (GIAB)
- EL Data, Security, and Technology (ELDST)
- Office of Information Systems Management (OISM)
- Material Measurement Laboratory IT Service Team
- Center for Theoretical and Computational Materials Science (CTCMS)

Each of these Case Studies provides a glance at research activities, the use of experimental methods and devices, the reliance on technology, and the scope of collaborations. It is important to note that these views are primarily limited to current needs, with only occasional views into the event horizon for specific projects and needs into the future. Estimates on data volumes, technology needs, and external drivers are discussed where relevant.

4.1 NIST Engineering Lab, Fire Modeling and the National Fire Research Laboratory (NFRL)

Content in this section authored by Randall McDermott and Artur Chernovsky, NIST

4.1.1 Use Case Summary

The National Fire Research Laboratory (NFRL) is involved in the development of computational fluid dynamics (CFD) models for fire protection, forensics, and research applications in buildings and wildlands. Also, we perform large scale fire tests and collect measurement data as well as video of the experiments.

Our use case can be thought of as two parts of a whole:

- First, we are developing a computational fluid dynamics (CFD) model of fire. The model is used by fire protection engineers, building designers, and fire investigators. This model runs on HPC clusters, usually around 10-100 cores, using the Message Passing Interface (MPI) for internal communication. The core set of CFD codes have scaled up to 10,000 cores on Titan at Oak Ridge (every few years), when allocation time was available because NIST doesn't have a resource on this scale. We would like to make this a more routine endeavor. Note that gaining access to leadership class machines, like Titan, often requires a laborious proposal and reporting process that acts as a disincentive to using these machines unless they are absolutely necessary.
- Second, Validation of the fire model requires large-scale experiments. At NIST we have a large-scale facility, the National Fire Research Laboratory (NFRL), where we can run experiments ranging from kitchen stove top fires up to single story house fires. We collect calorimetry data (how big is the fire, in terms of megawatts). We collect many channels of temperature and heat flux data and also take video, often from several angles, of the experiments. All this data must be archived, backed up, and readily retrievable by our fire scientists and model developers. The data life cycle can be decades. For example, we still use data from the 1950s to validate fire models. The need to be able to reproduce experiments requires careful storage of metadata for each experiment.

4.1.2 Collaboration Space

Domestic and international universities, government agencies, and scientific institutions (e.g., University of Maryland, [Bureau of Alcohol, Tobacco, Firearms and Explosions], and Juelich Supercomputing Center in Germany). We also have an active international user forum with about 3000 members who need access to documentation of the fire model. This group may also be interested in doing their own validation studies, in which case they might need access to the NFRL experimental data. We participate in an international collaboration sponsored by the International Association of Fire Safety Science (IAFSS) that aims to improve the models and guide best practices in modeling. The collaborators require access to model source code (through GitHub) and data (currently also provided on GitHub).

4.1.3 Instruments & Facilities

The needs will be described in terms of the two parts of the use case.

Computational
We maintain a small department level Linux cluster with about 1000 cores that provides computing to 10-20 people. This machine is now 10 years old and we plan to replace it in the coming year with only modest upgrades in size (limited by power, cooling, funding). I do not know the specs of the internal network between our desktop and the Linux machine, but it is currently adequate for data visualization (running visualization tools from the desktop with a mounted NFS drive to the Linux machine). When accessing the machine remotely, we use an NX client that allows us to run our vis tools from the Linux machine. The machine is running 24/7 for validation and/or continuous integration. Power shutdowns in our building (seems to be happening more and more) are extremely disrupting.

Experimental
10-20 Windows OS machines for data acquisition. 20+ video digitizers, including infrared cameras, stereo image digital correlation systems, 5-10 large visual display systems for real-time data monitoring. Typical measurements in fire experiments will be temperature from thermocouples, heat flux, force measurements, stress/strain gauges, gas analyzers, anemometry, smoke density via optical absorption, X-ray spectroscopy to check materials for lead contamination. Frequency of use depends on the scale of the experiment. Small experiments (fires of less than 1 megawatt) run monthly. Large experiments (up to 20 MW) may take a year or more.

4.1.4 Data Narrative
A typical validation exercise will start with planning (which often includes preliminary modeling). The experimental data is collected and stored in raw form (usually on a server within NFRL), and analyzed and reduced by the experimentalists. Usually, this results in a published report of an experiment. The modeling work generally follows and a publication may result.

We consider the integration of data and the model to be part of the model “validation suite”. The modelers work with the experimentalists to reduce the raw data into simple csv formats that are amenable to comparisons with the model results. We have a continuous integration framework that runs daily on a modest sized Linux cluster (100 node) to compare model results with stored experimental results. Any anomalies are flagged. If the results indicate a problem with either the model or the data, then we work to correct the problem, if possible, within the scope of the project.

One major consideration for model fidelity is grid resolution. We often run into the constraint of not having easy and fast access to large-scale computational resources that would help us test higher grid resolutions. But it is also not that simple. Running extremely high resolution requires a scalable code and this requires very knowledgeable computer scientists to develop. Our team consists more of engineers and mathematicians who are admittedly not HPC experts. The point is, we could use help to improve code performance for HPC.

4.1.4.1 Data Volume & Frequency Analysis

The use case has indicated that current data volumes approach the GB level for the computational use case and the TB level for the experimental use case, on a daily periodicity (when actively engaged in either use case; computations run continuously for half the year, while experimental data is collected less frequently, perhaps one quarter of the year). The primary data volume factor for experimentation is the collection of videos.

4.1.4.2 Data Sensitivity

There are two situations in which data sensitivity may become a factor for this research:

1. All documents go through an internal review process involving several layers of data review to ensure standards are being met.
2. If a proprietary product was tested, it may also require outside review by project partners external to NIST. This is exceedingly rare, but is available as an option if required.

4.1.4.3 Future Data Volume & Frequency Analysis

The use case has indicated that future data volumes approach the TB levels for both use cases, on a daily periodicity (when actively engaged in either use case).

4.1.5 Technology Support

The following subsections outline specific components of the use of technology to support the use case.

4.1.5.1 Software Infrastructure

The needs will be described in terms of the two parts of the use case.

Computational

- Compilers
- scripting tools (Matlab, Python, Bash)
- vis tools (Smokview [in house], Paraview, VisIt)
- profilers (we are not very experienced with these)
- NoMachine (NX client for vis tools across network, runs on Linux cluster in group, no data transfer)
- geometry preprocessors (PyroSim, BlenderFDS)

Experimental

- Matlab
- Python
- LabView
- Visual Studio
- Adobe Premium (video editing)
- ffmpeg (video streaming)
- instrument specific analytical software
- AutoCAD

4.1.5.2 Network Infrastructure

The primary network is the NIST enterprise network, offering default 1Gbps speeds. Higher speed networking is available for specific cases. When accessing resources from off-site, the VPN must be used, which imparts particular challenges for responsiveness.

4.1.5.3 Computation and Storage Infrastructure

The needs will be described in terms of the two parts of the use case.

Computational

We have a three-tier plan. First, for our day-to-day use, we will maintain our division level cluster (with aforementioned upgrades). This will handle HPC up to about 1000 cores. Next tier is that we hope that NIST will build a central system that will allow scaling up to 10k cores. Beyond this, we will expect to use NSF level machines (used to be XSEDE, now ACCESS) to scale up to 100k cores. For storage, we have yet to run into a bottleneck on our Linux machine. We just purchased 32 TB for $1000. It's cheap. For massive HPC runs, we would expect to analyze the data in-place and reduce it to whatever is needed for the project, then it would not need to be stored indefinitely.

Experimental

Currently have the storage we need (30 TB used out of 80 TB capacity) with RAID backup. The limitation is the gigabit network (even on campus) for video instrumentation and user access. For example, to stream full resolution of a FLIR research camera would require 10Gbps network capacity for uncompressed raw data. Users will benefit from faster network connection, say, scaling access from 10 to 1000 users. Data is being accessed from the NFRL data servers from a public facing website or over a Samba connection inside the NIST firewall.

4.1.5.4 Data Transfer Capabilities

The most typical use case for this is visualizing simulation data on desktop machines (anything besides the compute server where the data lives). Internally, the network can handle this. Over VPN it is very slow, takes hours. We use scp for transfer.

4.1.6 Internal & External Funding Sources

We are funded through NIST Scientific and Technical Research and Services (STRS).

4.1.7 Resource Constraints

Staff constraints are the largest barrier to productivity currently, namely through retirements of critical staff members that maintain portions of the infrastructure.

4.1.8 Ideal Data Architecture

Increasing the core capacity to some of the NFRL components to 10 Gbps in order to handle the anticipated need for higher bandwidth with more users (video downloads) in the future would be a good first step to addressing some of the bottlenecks from a technology perspective. We are working on this now but waiting on fiber pulls will be the biggest delay. Adding knowledgeable staff that can help maintain the infrastructure, such as the HPC cluster.

4.1.9 Outstanding Issues

The aforementioned use of the VPN when off-campus poses considerable challenges to some of the interactive use cases.

The code is a computational fluid dynamics code. It has a particle solver and a chemistry solver and a pressure solver. None of these currently exploit GPU in any way which means we cannot take advantage of some of the campus-wide HPC clusters.

Could data visualization be improved (made cleaner and more efficient)? Often these types of activities require a modern understanding of new tools, libraries, or methods.

4.2 NIST Engineering Lab, Materials and Structural Systems Division, Disaster and Failure Studies Program

Content in this section authored by Tanya Brown-Giammanco, NIST

4.2.1 Use Case Summary

The Disaster and Failure Studies Program coordinates all field deployments to investigate building failures caused by earthquakes, wildfires, hurricanes, tornadoes, or other causes. Disciplines include engineering (civil, mechanical, structural, earthquake), materials science, meteorology, metrology, modeling, and social sciences.

Our deployment teams collect perishable data in the immediate aftermath of disasters and conduct analyses of data and/or do forensic investigations to determine the technical cause of building failures and assess emergency response and communications procedures. Findings from our investigations result in recommendations for changes to building codes, standards, and practices, with the goal of reducing the likelihood and severity of future similar damaging events.

4.2.2 Collaboration Space

Each investigation is different, and may require collaboration with a variety of public, private, and academic institutions. We accept, and encourage, submission of data from public citizens, but also from collaborations through contracts or agreements, which require the exchange of data or generate new data from analyses. We typically exchange data with our collaborators through cloud services, like NIST-supported Box and Google Drive. We do not share data publicly during an active investigation, only with those on a need-to-know basis supporting the investigation. At the conclusion of an investigation, we share all data unless it poses a direct risk to the public to do so. In the past, we have shared data at the conclusion of an investigation via websites or specific data portals, but going forward we are looking to leverage NSF's DesignSafe platform (https://www.designsafe-ci.org/).

4.2.3 Instruments & Facilities

We use a variety of equipment based on the needs of a specific investigation. These could be located on the NIST campus, another location, or could be located at a collaborator's lab. We typically use NIST approved laptops, drones, mobile devices, and scanning, video, or photographic devices for some data collection and processing, and rely heavily on NIST-supported cloud-based resources to preserve and share data with other staff and collaborators.

4.2.4 Data Narrative

This will be dictated by the needs of each investigation. In nearly all cases, some perishable data will first be collected in a field setting using laptops, mobile devices, scanning, video, or photographic devices. These data are perishable because the area changes quickly after a disaster. People in the area may collect photos, videos, and other data during and immediately after a disaster. These people may include first responders who will only be in the area during and immediately after a disaster so it is imperative to connect with these people quickly in order to collect this data. It is also important to

collect images, videos, and measurements of the disaster area as close to the disaster as possible without impeding rescue efforts. These raw data will all be uploaded to a cloud-based platform for storage and sharing. Analysis of the raw data and generation of new data, plus comparison of these, will typically occur on NIST devices, with end-products uploaded to cloud-based platform. One challenge is ensuring the data has the proper metadata so that the team can respond to Freedom of Information Act (FOIA) requests after the investigation and more easily make public postings of the appropriate data. The team also needs to be able to identify and protect potential sensitive data such as PII.

4.2.4.1 Data Volume & Frequency Analysis

The use case has indicated that the current data volumes approach the TB level, on a daily periodicity when the use case is active. Field data collection is on the order of weeks in most instances, but there may be additional data collected via retrieval of historic design documents, interviews of witnesses/survivors/emergency responders after that time. Data will also be generated as team members conduct lab testing, modeling, or analysis. Data collection usually continues for several years before an investigation is completed. By nature, this is a "bursty" activity, thus there are periods where no activity may occur. Figure 1 describes the data flow.

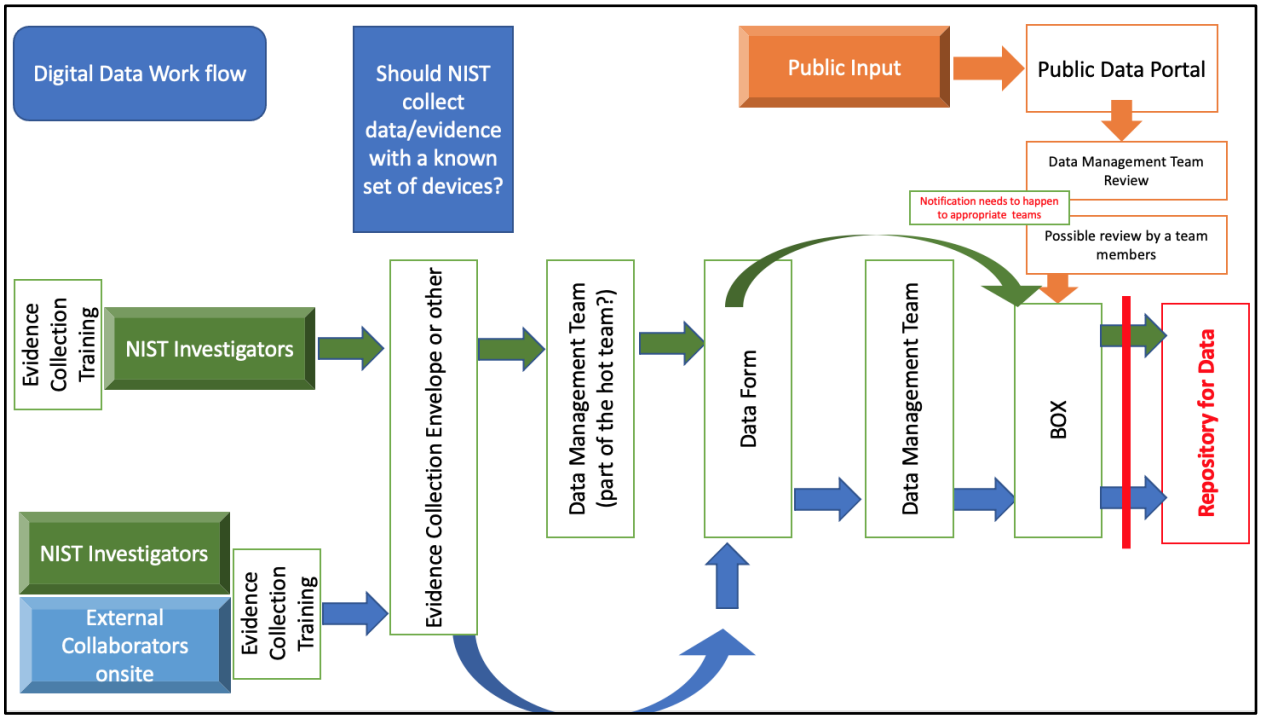

Figure 1: DFS Data Flowchart

4.2.4.2 Data Sensitivity

Due to the active investigatory nature of this data, almost all data is sensitive by default until it can be redacted. There is currently a challenge in ensuring that the data is organized in such a way that at the end of an investigation, FOIA requests can be

responded to quickly with the appropriate level of data and redaction, and that there are public postings of the appropriate data. This is also where identification of PII is critical.

4.2.4.3 Future Data Volume & Frequency Analysis

The use case has indicated that the future data volumes approach the TB level and is stored in NIST-supported cloud services (Google Drive, Box), on a daily periodicity when the use case is active.

4.2.5 Technology Support

The following subsections outline specific components of the use of technology to support the use case.

4.2.5.1 Software Infrastructure

ArcGIS, MatLab, ATENA, LS-Dyna, GriffEye, I'm sure there are others that I'm not aware of. We are able to purchase licenses for many of these, but our on-site capabilities for complex modeling with LS-Dyna are not sufficient so the simulation was too slow. The team considered using resources at TACC (no response), UIUC (couldn't meet needs), and the Ohio supercomputer center, but the DOD resource was free.

4.2.5.2 Network Infrastructure

Our teams need general internet access when working in the laboratory to ensure data can be stored and exchanged with collaborators, external team members, contractors, and other official investigation personnel at other agencies. In the field environment, this is often extremely difficult, as we typically have to rely on the cellular network, which may or may not be functional after a disaster, and even if it is functional, the network is usually overloaded because of the many people on the ground assisting with rescue, recovery, etc. When we can establish a location for long term investigative work, the EL Data Security and Technology (ELDST) group has established a local network to allow us to collect and move data as needed.

4.2.5.3 Computation and Storage Infrastructure

As mentioned above, we are currently relying on the Department of Defense supercomputing resources to run complex structural models using LS-Dyna. While EL has LS-Dyna available, there were not enough licenses to run the simulations in a timely manner.

4.2.5.4 Data Transfer Capabilities

We have transferred data to/from collaborators, with sizes ranging from GBs to TBs. We have used portable hard drives, NIST Google Drive and Box, as well as our Disaster Portal (https://www.nist.gov/disaster-failure-studies/data-submission-portal) to do this. This has taken hours to days, and is dependent on the VPN connection to NIST, data/internet connection (which can be low or non-existent in disaster areas where we only have a cell network to rely on, if that), location, size, file types. For disasters where there is no cellular network (e.g., Wildfires), researchers have to use portable devices and spend time uploading data after-hours from hotel networks.

When acquiring datasets from collaborators out in the field, there are limitations. Datasets are dropped into a file storage solution but are not automatically integrated with other data. When transferring large datasets to NIST Box and Google it takes patience, focus, and time, which can be difficult in a high-pressure field setting. The uploads are manual and any mistakes by the people performing the upload could mean starting the upload process over.

4.2.6 Internal & External Funding Sources

NIST Science and Technical Research and Services (STRS) funds (annual programmatic funds), and in some instances Congressional appropriations are provided for specific technical investigations.

4.2.7 Resource Constraints

I believe one of our biggest challenges is difficulty finding data at the end of an investigation so that we are able to publicly release some of the data and respond to Freedom of Information Act (FOIA) requests. Of the three investigations that have been completed under the National Construction Safety Team Act, only one of them (World Trade Center) has data posted publicly.

4.2.8 Ideal Data Architecture

One central location for all disasters where you log in, pull up the disaster, look at datasets relating to that disaster, and display data. It would be good to have the ability to compare data between disasters or perform overlays. GIS mapping of data would be helpful (time lapse view of disaster location pinning photos, videos, other records to the site based on time). Some ideals: avoid duplicate data, or at least allow researchers to find original data when data is duplicated and modified (e.g. images cropped, data redacted), manage access to specific data, which might include external collaborators or contractors having access to some data, custom data collection interfaces per disaster, some automated metadata tagging, a data inventory, faster uploads/downloads, reliable data connections in the field for collection, tagging, and lookups, and easy identification of sensitive data with access controls.

4.2.9 Outstanding Issues

Our biggest challenges are that we must respond to disasters very quickly and can only have so much "prep" work done in advance, and each event is unique and requires different data to be collected. Once we have a clearer picture of specific needs, we must stand up data collection efforts very quickly before the perishable data disappears. In some cases, this might just be a folder in Box where we can upload data as we collect it. We also have to operate in environments that are less-than-ideal, since we have extremely limited on-site support, are remote, and often have poor connectivity to resources and have to stand up our own network solutions. We need solutions that can be implemented quickly, used by many, and can operate in poor network situations. There may be several years between disasters so equipment and tools purchases for one disaster may be stale by the time we need to use them again.

4.3 Physical Measurement Lab, Sensor Science Division, Remote Sensing Laboratory

Content in this section authored by B. Carol Johnson, NIST

4.3.1 Use Case Summary

The Remote Sensing Laboratory (RSL) comprises spectral radiance and irradiance sources, filter and hyperspectral radiometers, and an environmental chamber that are utilized to perform calibrations, characterizations, and intercomparisons of radiometric artifacts (radiometers, sources, reflectance standards).

The Remote Sensing Laboratory (RSL) is designed to address Earth-oriented remote sensing radiometry for the spectral region from the ultraviolet to the short-wave infrared. The facility is used to calibrate, validate, or characterize test sources using the RSL radiometers, and to perform these same functions on test radiometers using the RSL spectral radiance and irradiance sources (https://www.nist.gov/programs-projects/remote-sensing-laboratory). The lab curates a multi-decade time series, acquiring and integrating data from a wide variety of radiometers and sources. New instruments are often integrated into the system. Data are incorporated from both on-campus acquisitions as well as from fieldwork. The IT challenges encountered on this project, including both hardware/software and long-term data management, can be considered consistent with a typical Physical Measurement Laboratory (PML) project.

4.3.2 Collaboration Space

The multi-decade RSL time series acquired from a wide variety of instruments is archived using the OISM file servers. Portions of the data are shared via FTP and SFTP with servers at San Jose State University Moss Landing Marine Laboratories (MLML) in Moss Landing, California (https://mlml.sjsu.edu/). When the RSL is used for calibration services and collaborations within PML, the data are shared internally. When external sharing is required, data is manually moved from a NIST archive into cloud services such as Sharepoint, or Google Drive. With the Marine Optical Buoy (MOBY) project, we may download relevant data from the internal project websites, which are hosted on a Moss Landing Marine Laboratories (MLML) server.

4.3.3 Instruments & Facilities

The Remote Sensing Laboratory has standard sources of spectral radiance and radiance and irradiance measuring filter radiometers or portable spectrographs. We have a thermal environmental chamber for temperature and relative humidity characterization of sources, radiometers, or critical optical components. We have an amplifier calibration facility. We have a 1kHz OPO tunable laser system and associated electronics. We have standards of spectral irradiance and diffuse reflectance standards. Currently we are implementing validation measurements between other facilities in the division in order to assess uncertainties. In the next year, we plan to develop a lamp/plaque facility for spectral radiance of hyperspectral radiometers and detector characterizations using the tunable laser.

RSL instruments are controlled over a private network that is controlled by a computer that has dual network interfaces so that it can connect to the NIST network as well.

4.3.4 Data Narrative

The primary workflows for data to publication are twofold:

- customized data pipelines from instrumentation
- calibration reports from instrumentation

For the former, each instrument acquires data during the experiments. The DAQ is either from RSL custom programs or COTS software from the instrument vendor. The data represent instrument results or instrument health (housekeeping) or both. These data files are archived on the OISM file servers; they are saved there automatically by the DAQ program. They reside there until we want to process, at which point the scientist would most likely copy the files to a new location for processing particular experiments. For the MOBY part of what RSL does, all of this is backed up onto MLML servers, where it may reside in an archive or be subjected to further processing by collation with additional instrument files. These data pipelines were developed in collaboration with MLML using the MLML_DBASE structure created by William W. Broenkow, Richard Reaves, and Stephanie Flora ("Introduction to MLML_DBASE Programs, 1993"). Key components in the MLDBASE structure include the processing programs, the raw data, the processed data, configuration files for the instrument, configuration files for a particular raw data file, PNG files for graphics, and html files for web display.

For the latter, there is a long-term archive of calibration data for spectrographs stored in comma or tab delimited format. These data are stored in a folder on the OISM file servers. The record for these is stored in a handwritten laboratory notebook.

4.3.4.1 Data Volume & Frequency Analysis

The use case has indicated that the current data volumes approach the MB level, weekly.

4.3.4.2 Data Sensitivity

Some data are public domain as it is funded by NASA or NOAA under public access guidelines. Other data are protected by NIST's Measurement Services or Cooperative Research and Development Agreement (CRADA) agreements between NIST and a non-federal party.

4.3.4.3 Future Data Volume & Frequency Analysis

Future data volumes approach the MB level, weekly.

4.3.5 Technology Support

The following subsections outline specific components of the use of technology to support the use case.

4.3.5.1 Software Infrastructure

LabView is used to control the instruments. MATLAB programs are used to analyze and report results. Some acquisition has been done via python and other equipment will likely

be migrated to control via python. Often the control programs are one-off and specific only to this application.

4.3.5.2 Network Infrastructure

We primarily use Windows but more recently a Linux machine, to control the instruments and archive the raw data on the OISM file servers. We archive the data in its original form, and analyze it separately, keeping backups on the OISM file servers. The laptops connect wirelessly to the NIST network, and connect to the instruments using a local area network.

We have severe problems in our ability to keep these interfaces operational, and this can be traced to a disconnect in understanding between the groups who design and modify the network (OISM), and the applications that are used for the scientific workflow. The simplest way to explain this is to use an analogy: When NIST infrastructure, such as roads or buildings, will be closed for a repair, we all get several emails in advance warning us of the situation and offering information on how to deal with the closure. Obviously, we could take corrective action in the absence of these emails. However, when a change to the network, domain settings, firewall rules, etc., occurs, we aren't notified, we have no way of discovering for ourselves, and we are left with computers that mysteriously stop communicating with our instruments.

Further, the people who manage the network (OISM) fail to appreciate that modern instruments, like everything else, require internet access for data acquisition, updates, and general control. The only network option for research equipment does not allow access to the internet except for specifically requested addresses. The internal network seems to be set up more for general business use, than for research. I understand NIST has to be careful, have a good firewall, etc. But please let us in on the conversation!

Another example of failure was the recent mandate from OISM to migrate our "Documents" folder to the cloud. Our functional account data was scheduled to be moved to the cloud automatically. We contacted OISM to create new folders on the OISM file servers and move all the files over. It was not a smooth migration because the tools in Windows are not smart enough to cover all contingencies in copy/paste or move – for example the tools stop on "can't move this file" so we had to get additional help from the OISM experts to complete the migration.

4.3.5.3 Computation and Storage Infrastructure

I don't think we have any special requirements other than central file services to support research and improved IT support.

4.3.5.4 Data Transfer Capabilities

We use FTP or SFTP for file transfer. The files are not huge, but transferring the whole database would take quite a while, especially using the VPN.

4.3.6 Internal & External Funding Sources

NOAA, Interagency agreement and NASA, Interagency agreement.

4.3.7 Resource Constraints

The biggest impediment to our work is a lack of research-focused networking and the lack of central storage for research data. It isn't always clear where to find IT support. There may be local IT groups to provide support or support may be provided by another researcher.

4.3.8 Ideal Data Architecture

Staff to develop and maintain instrument data acquisition and processes would be useful. A network architecture that allows researchers the flexibility to control instruments (and IoT devices) over Ethernet without requiring either a huge amount of overhead managing the security settings (and figuring out which requirements don't apply to our instruments!), trying to figure out who can help us keep the instruments functioning, while still allowing installation of additional packages/features/updates. Centralized storage that allows us to automatically store research data.

Ideally, we would have a research network with fewer IT security restrictions. Expertise would be provided to ensure that we are able to keep the resources secure, and there would be less reliance on enterprise tools intended to secure resources in a non-research enterprise environment. Services need to persist when changes are made to the rest of the network.

4.3.9 Outstanding Issues

There are no other issues to report.

4.4 Materials Measurement Lab, Joint Automated Repository for Various Integrated Simulations (JARVIS)

Content in this section authored by Francesca Tavazza, NIST

4.4.1 Use Case Summary

JARVIS (Joint Automated Repository for Various Integrated Simulations) is a repository designed to automate materials discovery and optimization using classical force-field (JARVIS-FF), density functional theory (DFT) (JARVIS_DFT), machine learning (JARVIS_ML) calculations, and experiments. Research focused areas include, but are not limited to, DFT modeling of energetics, elastic, electronic, optoelectronic, topological, and thermoelectric properties, superconductivity, phonon, infrared, and Raman spectra for crystalline solids. Jarvis also includes Quantum Monte Carlo modeling, extensive Machine Learning modeling of the same properties computed using DFT as well of STEM/STM images and XANES spectra, Quantum Computing capabilities and Tight Binding modeling. The focus of JARVIS-FF is to compare a large number of force fields for key properties like energetics, point defect and surface energies. A ML model for fitting such classical force-fields is also part of JARVIS.

The Materials Genome Initiative (MGI) highlighted how accelerating research in the field of materials requires data and algorithms to be easily shareable with the whole scientific community. The JARVIS project is part of this effort, with three databases (two of material properties and one of ML models) and all codes, workflows and various tools needed to acquire, process and output data (JARVIS-Tools: an open-source software package for data-driven atomistic materials design) publicly available through its websites (https://jarvis.nist.gov/, https://github.com/usnistgov/jarvis). In other words, the high-level goal of the JARVIS project is to facilitate material research at the atomistic level by providing systematically computed data on material properties as well as ways to compute and analyze such data. Such tools are designed to be easily modified by the stakeholders, to be applicable to the specific research they are interested in.

Our stakeholders are any other person or group in the global atomistic modeling or experimental community. Such a community is clearly well aware of the JARVIS project, as, to this day, JARVIS tool has been downloaded more than 200K times. Currently, it is getting 1.3K downloads/month, which is an extremely significant number given the size of the atomistic community. The repositories have had more than 140K items downloads, with many groups/institutions asking to be provided with the whole database(s) at once.

Data life cycle: Once created, our data remains constantly relevant as it expresses materials properties and such properties do not change and are of interest to the modeling and experimental community without time limit. The data generated in the JARVIS project in the past is still useful for present and future works.

4.4.2 Collaboration Space

The JARVIS project creates its own data. All data are shared among collaborators and to the stakeholders in electronic format using the JARVIS web-apps (https://jarvis.nist.gov/,) FigShare

(https://figshare.com/authors/Kamal_Choudhary/4445539) and Jupyter/Google colab notebooks (https://github.com/JARVIS-Materials-Design/jarvis-tools-notebooks). These datasets are downloaded all over the globe. Major users are from the USA, China, and India. As our data is available online, the geographic locations of the stakeholders do not impose any barriers.

The datasets are created through density functional theory, molecular dynamics, and calculations and their post processing, scholarly articles and a few experiments.
In future, we might want to have a REST API and Globus connections to our dataset for exchanging information. The data size is very large and constantly growing (~100 TB per year). In the future, larger storage capabilities will be necessary.

4.4.3 Instruments & Facilities

We use a number of super-computing facilities at NIST such as the CTCMS, and OISM-managed compute clusters (i.e., Raritan, Enki and Nisaba) for performing CPU as well GPU intensive calculations. The CPU usage has a limit up to 1000 cores per user (such as on the Raritan cluster). For GPU usage, similar constraints are in place. Additionally, several high memory calculations, such as beyond density functional theory methods and large atomic structures, are not possible on the current platforms because of memory bottlenecks.

As mentioned above, we are in dire need of storage (up to 1 petabyte) to meet ongoing as well as future demands. Metadata such as electronic wavefunctions during a quantum calculation cannot be stored right now because of memory storage limitations. Access to such metadata would drastically enhance future subsequent calculations.

4.4.4 Data Narrative

One of the goals of JARVIS is to discover suitable materials for targeted applications and we use screening based on available materials data for finding most suitable candidates. We use python-language based workflows which are part of JARVIS-tools to automate DFT, MD, ML calculations. After the calculations, materials properties are analyzed, and webpages are produced (following a predefined XML schema) to show relevant information, only discarding all the raw data. We use experimental data as much as possible to evaluate error and uncertainty in predictions.

4.4.4.1 Data Volume & Frequency Analysis

The use case has indicated that the current data volumes approach the TB level, hourly, when the experimentation is running.

4.4.4.2 Data Sensitivity

There are no sensitive aspects to this research data.

4.4.4.3 Future Data Volume & Frequency Analysis

Data volumes will approach the PB level, hourly, when the experimentation is running.

4.4.5 Technology Support

The following subsections outline specific components of the use of technology to support the use case.

4.4.5.1 Software Infrastructure

The following software packages are regularly used:

- Electronic structure (VASP, Quantum espresso, QMCPack, Wannier90)
- Force-field (LAMMPS, ASE)
- Machine learning (Scikit-learn, PyTorch, LightGBM, Scikit-image, DeepGraphLibrary)
- Quantum computation (Qiskit, PennyLane, Tequila, Cirq)
- Web-app (Django, HTML, Javascript)
- Data mobility tools (Rsync, GDrive, Figshare API)
- VASP is a commercial software, others are publicly available packages.
- Github

4.4.5.2 Network Infrastructure

Figshare, Google Drive, OneDrive, Globus

4.4.5.3 Computation and Storage Infrastructure

We are currently using several High-Performance Computing Clusters internal to NIST. We are planning to continue using those plus whatever new systems NIST will deploy in the future, in addition to High Throughput Computing (e.g., Grid, etc.), Cloud (e.g., Commercial, or R&E operated).

4.4.5.4 Data Transfer Capabilities

We do perform data transfer to/from collaborators but, currently and in the foreseeable future, they do not constitute bottlenecks for our project. This is mostly due to using Figshare to share data with anyone who needs access to the processed data in the databases.

Things would be different if we were able to store all the raw data/metadata and wanted to share them with collaborators. As pointed out many times in this document, we currently can’t do that because of storage limitations.

4.4.6 Internal & External Funding Sources

Our funding comes directly from NIST, most of it because of MGI

4.4.7 Resource Constraints

Lack of availability to sufficient computing nodes, that contain sufficient speed and memory, constrains our scientific productivity. In the future, we expect such constraints to become a serious threat to the project if more resources are not made available.

Data storage is another serious limiting factor for our work, as the databases are in the 330 TB range, and that is without saving important metadata (because of current space limitations) that would have been useful to the user (for instance: it would have permitted

to start calculations of new properties from converged configurations, therefore saving days of computations).

4.4.8 Ideal Data Architecture

The perfect data architecture would resemble something like a NIST version of XSEDE where there is dedicated staff to install and maintain software and hardware as well as availability of high-performance CPU/GPU/Storage. We can choose XSEDE comet system configs for example:

- Intel Haswell Standard Compute Nodes
  - ***Node count***: 1,944
  - ***Clock speed***: 2.5 GHz
  - ***Cores/node***: 24
  - ***DRAM/node***: 128 GB
  - ***SSD memory/node***: 320 GB
- NVIDIA Kepler K80 GPU Nodes
  - ***Node count***: 36
  - ***CPU cores:GPUs/node***: 24:4
  - ***CPU:GPU DRAM/node***: 128 GB:48 GB
- NVIDIA Pascal P100 GPU Nodes
  - ***Node count***: 36
  - ***CPU cores:GPUs/node***: 28:4
  - ***CPU:GPU DRAM/node***: 128 GB:64 GB
- Large-memory Haswell Nodes
  - ***Node count***: 4
  - ***Clock speed***: 2.2 GHz
  - ***Cores/node***: 64
  - ***DRAM/node***: 1.5 TB
  - ***SSD memory/node***: 400 GB
- Storage Systems
  - ***File systems***: Lustre, NFS
  - ***Performance Storage***: 7.6 PB
  - ***Home file system***: 280 TB

4.4.9 Outstanding Issues

Our data needs to be publicly available on-demand through various web services. For large storage amounts, institutions sometimes recommend a mix of "hot" (readily-accessible) and "cold" (barriers to access but cheaper to operate) storage. The "always current" character of our data has an implication that it also should be always available for human and machine consumption.

4.5 Communications Technology Lab, National Advanced Spectrum and Communications Test Network (NASCTN)

Content in this section authored by Jason Coder, NIST, with contributions from Duncan McGillivray, NIST, Aric Sanders, NIST, and Adam Wunderlich, NIST.

4.5.1 Use Case Summary

This use case represents the portion of the National Advanced Spectrum and Communications Test Network (NASCTN) that resides in the Communication Technology Laboratory at NIST. This includes projects that are primarily funded by other government agencies, typically the Department of Defense. These projects utilize NIST facilities such as the National Broadband Interoperability Testbed and the infrastructure that surrounds it.

The National Advanced Spectrum and Communications Test Network (NASCTN) is a multi-agency chartered partnership that organizes a national network of Federal, academic, and commercial test facilities. It provides testing, modeling, and analysis necessary to develop and deploy spectrum-sharing technologies and inform future spectrum policy and regulations. Members include DoD, NASA, NIST, NOAA, NSF, and NTIA. NASCTN is hosted within the Communications Technology Laboratory (CTL) at NIST.

This research primarily focuses on the measurement, modeling, and analysis of wireless communications. To perform this work, we rely heavily on expertise from the following domains: statistics, mathematics, machine learning/artificial intelligence, radio-frequency metrology, cellular network engineering. Through our research, we aim to develop new ways of measuring the performance of wireless communications systems, their potential interference with other systems, and their ability to coexist with other systems.

The work that NASCTN conducts varies greatly in the degree of interagency collaboration. Some projects are led by a single agency, where work products are fostered within a single agency's network, publication guidance, and dissemination avenues. At other times, NASCTN projects necessitate leveraging multiple agencies where project collaboration, tests, data, and analyses need to be negotiated and coordinated across interagency boundaries. These interagency avenues provide significant challenges as agencies contained within a cabinet level department have disparate rules, regulations, and tools available to their disposal.

The data lifecycle described here is broken up into working data, production data, and publication data. Working data is generated by networked test equipment and consolidated with metadata into a database accessible to the test team. This data is stored locally before processing and retention is governed by agency guidelines. In practice, this data is typically sunset with the project conclusion.

Production data is presented as wrangled raw data elements, typically published along with technical reports. The retention of production data is in line with that of technical reports and agency guidelines.

Publication data refers to processed and analyzed data, which is published along with technical reports and is retained along agency guidelines.

4.5.2 Collaboration Space

Primary collaborators include members of the NASCTN network: DoD, NASA, NIST, NOAA, NSF, and NTIA and their contractors (chiefly, MITRE). We share working data, production data, and publication data with these collaborators.

Publication data is also made publicly available at the conclusion of the project. Collaborators are geographically located across the United States. Publicly released data may be downloaded anywhere in the world.

At present, we share data with collaborators through explicit file transfer services, ranging from shipping a hard drive to an individual to virtual sharing using a secure FTP service.

In the future, having sharing services more tightly integrated with our data acquisition and processing software would save considerable time and enable collaborators outside of NIST to play a larger role. There has also been interest from our collaborators in tapping into our cellular network infrastructure through a robust VPN interface.

4.5.3 Instruments & Facilities

Each NASCTN project utilizes a different combination of facilities and scientific equipment. Most utilize some of CTL's measurement chambers (e.g., National Broadband Interoperability Testbed, 5G Coexistence Testbed), it's commercial-grade cellular networks, and scientific instrumentation (e.g., signal generators, signal analyzers, network analyzers, network load generators). Each piece of instrumentation is connected to an IT network depending on each network's individual policies and where the data need to go. The four possible networks include:

- the NIST enterprise network,
- NIST research equipment network (REN),
- the Public Safety Communications Research (PSCR) Demonstration Network,
- or a local private network.

The NIST enterprise network is available in every lab and office and is centrally managed by Office of Information Systems Management (OISM).

The NIST research equipment network (REN) is a fee-based network segment also provided by OISM. The basic block of IP addresses behind the REN firewalls is provided by OISM, but we have invested significant resources over the past five years to implement the ports and network switching capacity to meet our needs. This equipment and portions of the REN network is managed by scientists or network engineers outside of OISM. Our goal is to have 10 Gbps copper or fiber connections between any two points in our lab (achieved through our own network switching gear).

The PSCR Demonstration network has a separate authority to operate (ATO) from DOC and is managed by the PSCR team of network engineers and security professionals, who operate the PSCR network independent of NASCTN and OISM. We have invested some of our own resources to ensure we have the connectivity we need to this network. This network hosts our cellular deployments. Our connections to the PSCR network are on the order of 10 Gbps, achieved through dedicated fiber connections between labs.

Finally, local private networks may be used to form small, temporary connections between instrumentation on specific projects. These are managed by local research staff.

Generally, the network resources are located on the same site (but not necessarily in the same building). The specific computational resources used are described in our other responses.

There is discussion about rebuilding and relocating some of the measurement chamber facilities. In those discussions, the existing network structure is maintained, but an additional classified-ready network is proposed.

4.5.4 Data Narrative

The research endeavor begins with a period of programmatic planning between multiple agencies. In this period, several key documents are proposed, edited, and finalized. This requires input to a single document across agency boundaries, shared research on previously completed projects, and joint team meetings. The important parameters of the experiment are defined and simulations are conducted to provide more information about the problem space. Out of this planning period, a test plan is created specifically targeting the problem of interest. If one or more of the test partners has instruments, computational assets, or infrastructure that can be leveraged, this is taken into consideration. If not, the required product is purchased.

For example, a test was run on a communications network consisting of a cell phone exchanging messages with a cell phone base station. This test leveraged a large chamber with radio frequency absorbing foam (inside of a NIST facility), but required the purchase of the cell phone and cell phone base station. The cell phone base station required a set of industry specific software (e.g., network core) that a collaborator on a managed network outside of NIST owned and was leveraged. To acquire data, specialty diagnostic software was loaded to the cell phone, and the configuration of the cell phone base station was performed by vendor specific software. The electromagnetic radiation from the cell phone was monitored using an instrument configured for the NIST internal network. The specialty software that cell phone networks use was only available on the collaborator's wireless network, requiring the experiment to be conducted inside of their network. Once data was generated, it had to be physically transported from the collaborators network to the NIST internal network via solid state drives. The data from those drives was injected into a database inside of the NIST network. The data, which was 100's of GB in total, was then processed inside of the NIST network.

4.5.4.1 Data Volume & Frequency Analysis
Current data volumes approach the GB level, daily.

4.5.4.2 Data Sensitivity
Some of the data generated in this use case may qualify as CUI, be proprietary, or be a "deemed" export and subject to export control depending on the frequencies of interest, the modulation format of the signals, the measurement tools used, or the measurement of specific vendor devices (which then becomes BII). In general, the sensitivity level is established before measurement data is acquired so all staff know how to handle it. If data is deemed higher in sensitivity, special drives provided by OISM are used to encrypt the data.

4.5.4.3 Future Data Volume & Frequency Analysis
Future data volumes may approach the GB level, daily.

4.5.5 Technology Support
The following subsections outline specific components of the use of technology to support the use case.

4.5.5.1 Software Infrastructure
- Analysis and Visualization
    - Matlab/Simulink – provided by NIST
    - Python – open source
    - JMP – commercial requires purchase
    - R – open source
    - Java

- Databasing
    - mySQL – open source
    - mongoDB – open source

4.5.5.2 Network Infrastructure
Our connections to the NIST enterprise network are Shielded Cat6 copper[4]. Connections to the NIST REN are 50% shielded Cat6 copper and 50% multimode fiber. NIST REN connections are managed by Cisco Nexus switch gear purchased with our funding and each connection is capable of 10 Gbps. REN connections outside the building are significantly slower, and we try to avoid them when moving data. To aid in this, we've placed a server near the lab as an intermediate data storage location and to pre-process data before moving it out for analysis. The server is on the NIST REN but managed by local research staff.

[4] This shielding is a special requirement for our lab. Given that we do precision measurements on wireless signals, we are cognizant of other sources of signals and attempt to shield them when possible.

Connections to the PSCR network are through dedicated single mode fiber between labs and buildings on the boulder site. In each lab, we have a switch to break out to individual connections. Connections to the labs are 10 Gbps, but connectivity may vary experiment to experiment.

Computational resources are located in the Boulder Computing Facility (BCF), connected via fiber.

We share data with external collaborators through secure FTP or postal mail.

Moving data between the NIST enterprise network and the NIST REN is accomplished through firewall policy management and/or manually moving data with external storage (i.e., sneaker net). Moving data between the PSCR network and the NIST networks is only possible through sneaker net (see outstanding issues below). Our sneaker net resources will typically handle terabytes of data per project.

4.5.5.3 Computation and Storage Infrastructure

Current resources include dedicated CPU/GPU compute nodes and network attached storage (NAS) purchased and supported by NIST/CTL in Boulder. NASCTN operates a 400 TB NAS system, two NVIDIA GPU machines with a total of over 69,000 GPU cores, and two Dell servers with a total of 80 CPU cores. Several machines are housed in the Boulder Computing Facility (BCF) and others are housed in the NASCTN laboratory. We see potential advantages to making greater use of cost-effective cloud-based storage and computing in the future, e.g., to facilitate data transfer and data processing from sensors fielded outside the NIST campus, and to reduce the need to procure computing and storage hardware.

4.5.5.4 Data Transfer Capabilities

Yes, NASCTN has previously transferred sensitive data (100's of gigabytes) to outside collaborators by sending hard drives through conventional postal mail and by utilizing collaborators' file transfer capabilities. We need better options to share and receive large, sensitive (access controlled) datasets with collaborators, both inside and outside of the federal government.

4.5.6 Internal & External Funding Sources

The NASCTN program is currently funded by NIST, the hosting agency on STRS funds. However, individual projects of the NASCTN program are typically funded by a sponsor. These sponsors could be industry stakeholders or other Federal agencies. In practice, Federal agency sponsors of NASCTN work leverage "Spectrum Relocation Funds" that are derived from FCC led spectrum auctions.

4.5.7 Resource Constraints

Scientific productivity in NASCTN projects is heavily affected by data transfer choke points due to test architectures having to traverse internal network boundaries or even inter-agency network boundaries. Future work will include elements of data generation

occurring in the field outside of a readily available network architecture (IOT solutions). We anticipate cloud storage and computing to be leveraged more heavily in the future to alleviate some of the network boundary challenges. In practice, however, not all test collaborators have equitable access to cloud solutions due to test agency policies, or even procurement limitations. Furthermore, scientific productivity is impacted by burdensome agency policies that compartmentalize collaboration tools to internal access even though the collaboration tools are intended to overcome compartmentalization. A specific example of limits on collaboration tools is the MS Teams implementation at NIST when teaming with other agency collaborators. Here, at best, collaborators are furnished with guest accounts, which renders seamless Teams integration on the collaborators end moot. In practice, collaboration tools and data transfer tools are not implemented on a government wide level, which has the practical effect of segmentation of scientific work and results in significant overhead and opportunity costs.

4.5.8 Ideal Data Architecture

There are several pieces that we need for an ideal data architecture, the first piece is access. Secure, simple access for any participant in any experiment / test. These participants should have access to data, instruments, produced software and subscribed software with any connection to the internet. Next, we need a way to manage and communicate with hardware that is addressable, secure, easy and requires no additional software from the host agency. Once this hardware was attached to the networkable space it would create and deliver data to a centrally located database (e.g., NOSQL) with the ability to receive data at Giga-Bits per second. Flags could be set in the data for automatic analysis. The analysis software would live on a flexible computer that could increase its number of allocated cores with a very fast connection to the host computer(s) of the database. The software for analysis would be managed in git repositories that have full continuous development / continuous integration tools available. Collaborators would be able to access the data, visualize the data, and store linked analysis on the centralized database. Fleet deployment, configuration, and management of hardware would require only registering a device.

- Wish list
    - Remote login for all collaborators
    - Networked research hardware that doesn't require all of the same "enterprise" security agent software that are installed on OISM managed desktop computers.
    - Large fast database available to all, but with permissions managed
    - Fast load balancing computational node attached with high throughput connection to database
    - Fleet deployment of hardware, i.e., one click image install, life signs analysis, automatic updates
    - No computational overhead for security software

4.5.9 Outstanding Issues

Data Diode

In FY21, we tried to solve the problem of using sneaker net to move data from the PSCR network to the NIST enterprise network. The best solution for this appeared to be a data diode. A data diode is a network appliance that only allows for the transfer of information in one direction (i.e., a valve). Such devices are used at other Government agencies, including intelligence agencies. We did the market research to identify a solution that appeared to meet all relevant security standards, prepared a procurement to purchase the data diode at our cost (~$150k), and received approval to deploy the device on the PSCR network. However, we had to cancel this process because OISM wouldn't agree to having it deployed on their network. Thus, sneaker net remains in operation with high costs.

Sharing data with collaborators at other federal agencies

If this process were made easier, it would enable collaborators at other Federal agencies to play a larger role in some of our projects. This opens the door to new opportunities not possible with today's solutions/policies.

Policies/procedures for CUI (and higher)

It isn't clear how large data sets need to be marked to indicate if they contain CUI. Do we need to mark individual files, folders, raw data, processed data, metadata, etc.? Policies on this aren't clear, but as we generate and work with an increasing amount of controlled data, this will become more increasingly important and guidance would be helpful. It also needs to be made more clear which IT systems are suitable to handle what types of controlled data. Looking ahead, there are discussions about developing an IT network capable of classified work. It isn't clear who (e.g., OISM, research staff) would manage a classified network or how it will operate.

Collection of data at remote sites

In addition to sending data out to collaborators, we're also seeing an increasing need for bringing data into our IT networks from remote sites. As we deploy sensors in the field, we need to develop efficient ways to bring those data sets back to our IT networks for processing. These in-field sensors may operate for several months without being serviced in person, but we need to pull data from them on a regular basis through remote access. It isn't clear how we can achieve this under the current network/policy structure.

4.6 Materials Measurement Laboratory (MML) Biomarker and Genomic Sciences Group (BGSG), Genome in a Bottle (GiaB)

Content in this section authored by Nathan Olson, NIST

4.6.1 Use Case Summary

The Biomarker and Genomic Sciences Group develops standards and methods for improving confidence in fundamental measurements in biology through new and improved techniques, methodologies, and standards based on optical and genomic methods. They focus on improving the quantitative measurements of biological markers of gene expression in eukaryotic cells and biological fluids.

Primarily, they provide materials and resources for stakeholders to use in validating genome sequencing and data analysis methods. These stakeholders include NIST researchers in the MML-Biosystems and Biomaterials Division as well as collaborators in the Information Technology Laboratory (ITL). Additional stakeholders may include researchers from government (national and international), academic, or industry, specifically research and clinical laboratories performing human whole genome sequencing.

Whole genome sequencing data are first generated from the reference materials. While some of the sequencing data is generated internally, most of the data is generated by collaborators or contracted out. The raw sequencing datasets tend to range from hundreds of GB to TB of data. The raw data is processed using bioinformatics pipelines to generate genome assemblies or variant calls. These assemblies and variant calls are then used to generate the reference characterizations. Collaborators then used these characterizations, and publicly available raw sequencing data, to validate their methods. The characterization and sequencing data are made publicly available from a NIH hosted ftp site.

4.6.2 Collaboration Space

- Google deepvariant team - Bay Area CA, collaborator contracts out some data generation, does not do primary data analysis, main focus is on bioinformatic methods development.
- Baylor team - Houston TX, collaborator generates data sets and does some primary analysis and bioinformatic methods development
- Miten - Boston MA (Northeastern), collaborator generates data sets and does some primary analysis
- UCSC team - Santa Cruz California, collaborator generates data sets and does some primary analysis and bioinformatic methods development

4.6.3 Instruments & Facilities

- NIST clusters - OISM-managed nisaba primarily for AI/ML work, also used OISM-managed enki & raritan clusters, and CMCS clusters, with mixed experiences. The nisaba cluster is currently being used daily to multiple times a week.

- AWS - S3 buckets, EC2 instances; Cloud storage - used daily; compute - used monthly
- Personal laptops and desktops - mostly macbooks and Dell desktops and workstations running linux (ubuntu and arch)
- Team managed QNAP NAS (24Tb storage, raid 10) - likely phasing out will transition big data storage and data backup copies to Isilon
- MML Isilon storage - used regularly as a local copy of large datasets.

Future plans include migrating data analysis pipeline and exploratory analyses to AWS using pipeline execution / orchestration tools and Web based IDE for exploratory data analysis.

4.6.4 Data Narrative

Raw sequence data are processed to generate secondary data used for exploratory data analysis and material characterization. Snakemake (https://snakemake.readthedocs.io/en/stable/), a python based pipeline development language and execution engine, is used to process the raw data and document the data analysis process. The snakemake pipelines are mainly run on NIST user desktops/ workstations and NIST clusters. We are exploring methods and tools for orchestrating/ running snakemake pipelines in the cloud using headless services offered by AWS such as lambda and batch functions (e.g., Tibanna https://tibanna.readthedocs.io/en/latest/ and Amazon Genomics CLI - https://aws.amazon.com/genomics-cli/).

For exploratory data analysis, primarily generating results and figures for manuscripts our group used the R programming language and the Rstudio IDE run on personal laptops and desktops. We are also looking into using AWS SageMaker as a cloud-based alternative to running Rstudio, with the idea that cloud-based solutions will make it easier for team members to work more collaboratively on projects.

4.6.4.1 Data Volume & Frequency Analysis

The current data volumes approach the GB level, weekly.

4.6.4.2 Data Sensitivity

There are no sensitive aspects to the research data.

4.6.4.3 Future Data Volume & Frequency Analysis

Future data volumes should approach the GB level, daily.

4.6.5 Technology Support

The following subsections outline specific components of the use of technology to support the use case.

4.6.5.1 Software Infrastructure

- gitlab - for collaborating on internal code projects
- Globus/ aspera for data transfers
- Snakemake - for defining and executing bioinformatic data analysis pipelines

- Visual code studio - General IDE for scripting and development
- Rstudio - IDE primarily for the statistical programming language. Primarily use hard-wired Ethernet connections to desktops and workstations along with Wi-Fi for laptops. Team mostly works remotely connecting to NIST networks through the VPN for workstation and cluster access. For cloud access the NIST / NOAA n-wave proxy user for data transfers. R.
- Open-source software used for bioinformatic analyses (all publicly available and free of charge)

4.6.5.2 Network Infrastructure

Primarily use hard-wired Ethernet connections to desktops and workstations along with Wi-Fi for laptops. Team mostly works remotely connecting to NIST networks through the VPN for workstation and cluster access. For cloud access, the NIST / NOAA n-wave proxy user for data transfers.

4.6.5.3 Computation and Storage Infrastructure

We plan to migrate most of our analysis workflow to AWS cloud while using a mix of OU and NIST level data storage for backup and NIST clusters for GPU and informatic analyses with high data read/write volumes.

4.6.5.4 Data Transfer Capabilities

Yes, we regularly share datasets with collaborators as well as transfer datasets from collaborators to NIST. Dataset sizes can range from GBs to TBs, with the largest dataset being 26TB. Larger data transfers can take multiple days primarily using S3 buckets with the AWS cli tools or ftps sites using Aspera or curl. Sharing and transferring large datasets was challenging, primarily due to issues related to requirements for granting write access to NIST S3 buckets as well as allowing access to the S3 buckets by collaborators outside the NIST firewall. Occasionally, mailing hard drives is easier than transferring data over the network.

4.6.6 Internal & External Funding Sources

Funding is provided through NIST.

4.6.7 Resource Constraints

There are no additional resource constraints to list at this time.

4.6.8 Ideal Data Architecture

Our ideal data architecture is one where team members can collaboratively develop and run snakemake based bioinformatic pipelines on NIST local desktops/workstations, NIST clusters, or AWS as appropriate. Similarly, for exploratory data analysis a cloud-based IDE for data analysis such as Rstudio workbench, Jupyter Lab, and Google CoLab. These IDE along with the visual studio IDE provide a platform for more collaborative research allowing for remote collaboration and pair programming.

4.6.9 Outstanding Issues

Our biggest challenge so far has been getting access and setting up different IT resources for our work. It was not straightforward to get access to the NIST clusters. The clusters have varying levels of documentation and are not user friendly. Additionally, the clusters are running at full capacity and the queue can be very long.We have found AWS a viable option but getting things setup in the cloud has been slow, partly due to the learning curve and familiarizing ourselves with AWS offerings and getting appropriate access and permissions to test out and implement different services. We currently have an engagement with AWS Pro-services to facilitate getting our ideal data architecture setup.

4.7 Engineering Laboratory Data, Security & Technology (ELDST)

Content in this section authored by Carolyn Rowland, NIST

4.7.1 Use Case Summary

The Engineering Laboratory Data, Security & Technology (ELDST) group provides IT services and software development to the Engineering Laboratory (EL), one of the NIST research labs:

- Assist staff in navigating IT security rules for purchasing, contracts and lab implementations,
- Assign and manage devices for the research equipment network in EL,
- Help staff with datasets (e.g., visualization, organizing datasets for public release)
- Write web tools, research tools, and help staff with services like Github or pages.nist.gov
- Consult with staff on the various platforms and tools available and help them choose the best one for their needs
- Work with OISM on research use cases to help provide specialized support
- Help set up hybrid workshops and meetings and provide moderation services for complex hybrid events

4.7.2 Capabilities or Special Facilities Offered

We provide small-scale hosting of public and non-public research applications and datasets as well as servers that run scientific software application license servers (e.g., ArcGIS, Oxygen). We manage a file server for EL research data to augment the central resources that are available for administrative data. We manage the public server pages.nist.gov for NIST that allows NIST Github users to leverage the Github wiki capabilities while meeting Federal security requirements.

4.7.3 Technology Narrative

The following sections outline the technology footprint for this use case.

4.7.3.1 Network Infrastructure

This group provides special purpose network capabilities (outside of institutional network infrastructure) to support research such as wireless backhaul in remote campus locations or temporary network facilities in field work and disaster locations.

4.7.3.2 Computation and Storage Infrastructure

There is no computation provided by this group. Some local groups have their own HPC clusters. System administration support for these clusters is provided by researchers or through fee-for-service by OISM.

There is a small storage solution provided for research data. This enables researchers to put their data on a lab-wide shared resource for internal collaboration. Any external collaborations with associated storage would use Box or Google due to the collaborative abilities of those tools.

4.7.3.3 Network & Information Security

The IT security officers for EL reside in this group. They can consult on any IT security technology issue that arises within EL. On occasion the IT security officers have supported standalone networks for disaster and failure studies or localized networking needs.

4.7.4 Organizational Structures & Engagement Strategies

The following sections outline the organization structure, and engagement with the research community.

4.7.4.1 Organizational Structure

This group reports to the EL Directory/Deputy.

4.7.4.2 Engagement Strategies

We commonly engage with the EL researchers for various reasons:

1. Need to do some kind of data collection, management, analysis, visualization, or publishing
2. Need a software tool written
3. Need consulting on existing tools (e.g., collaboration using Google, Box, meeting platforms, streaming platforms, virtual conferences, using AWS or pages.nist.gov publishing)
4. Need special purpose lab support (e.g., LabVIEW coding, networking, hardware support (e.g., National Instruments hardware, data acquisition, sensors, lab design/experiment design)
5. Need off-site services or collaborations (e.g., non-standard mobile device support, databases, remote network solutions)

4.7.5 Internal & External Funding Sources

Funding is provided by EL.

4.8 Office of Information Systems Management

Content in this section authored by James Fowler, NIST

4.8.1 Use Case Summary

NIST's Office of Information Systems Management (OISM) plans, deploys, and manages information technology (IT) resources and infrastructure, promulgates Department of Commerce IT/Security policies, implements NIST-specific policies, and assesses IT security controls for NIST's IT systems underlying NIST's intramural, extramural, and administrative programs. Hence OISM's responsibilities support all NIST staff across all organizational boundaries and encompass all NIST facilities.

The NIST OISM organization is responsible for the following customer-facing services:

- Getting Help & Information: These services include operation of IT help desks at the NIST Boulder and Gaithersburg campuses as well as operation of the web-based NIST Service Portal.
- Getting Connected: These services include wired/wireless network access, remote access to the NIST network, network security, user authentication, and management of user access.
- Application Development & Support: These services include custom and fee-based software development services along with enhancement and maintenance of existing software developed for supporting and executing NIST's mission.
- Desktop & Mobile Computing Services: These services include centralized management of NIST's traditional computing endpoints (i.e., Windows and MacOS computers), mobile endpoints (iOS and Android phones and tablets), central data storage and backups for NIST users, as well as purchasing services.
- Communications & Collaboration: These services include user and resource provisioning for NIST's primary communication/collaboration platform, Microsoft's Office365 environment as well as for G-Suite, Box, Video Collaboration Platforms such as Zoom, WebEx, and BlueJeans, and for legacy telecommunications (plain old telephone system) infrastructure on the Boulder and Gaithersburg campuses, provisioning of handsets, and audio-conferencing services.
- Hosting & Co-Location Services: These services include the OISM offerings for customer-managed virtual machines (on-premise and in Amazon Web Services), OISM-managed virtual machines, and co-location in OISM computing facilities.
- IT Security: These services include the IT security assessment efforts that OISM offers to NIST IT System Owners, vulnerability and compliance testing services, security incident response and investigation services, and other services related to OISM's IT security program.
- Research Services: These services include OISM's shared, centrally-managed software licenses for technical staff, for OISM-developed custom laboratory automation systems, for system administration/hosting of High-Performance Computing clusters, for support (along with partners in NIST's Labs) of

NIST's Open Access for Research tools, and for nascent research data infrastructure offerings.

4.8.2 Collaboration Space

OISM does not have any additional collaborations to list in this section.

4.8.3 Capabilities or Special Facilities Offered

The OISM manages computing facilities on the Gaithersburg and Boulder campuses that accept Lab-owned IT equipment meeting specific eligibility requirements. The most common need from NIST Labs that necessitates co-location in OISM computing facilities is public internet access, i.e., a Lab's use case requires a public-facing connection. One well-known example is the Physical Measurement Lab's Internet Time Service.

The OISM also manages peering/cross-connects with NOAA's N-Wave network infrastructure, with Internet2 through the Mid-Atlantic Crossroads (MAX), and hosts Globus infrastructure for high-speed, large-volume data exchange with external partners.

NIST provides unique user facilities, such as the NIST Center for Neutron Research, the National Cybersecurity Center of Excellence, the Center for Nanoscale Science & Technology's Nanofab facility, the Hollings Marine Laboratory (HML) in Charleston, SC, the NIST Fort Collins Facility, in Fort Collins Colorado, and Kehaka, Hawaii NIST campus located on the island of Kauai (HI). All of these facilities are supported by the OISM.

4.8.4 Technology Narrative

The following sections outline the technology footprint for this use case

4.8.4.1 Network Infrastructure

The NIST network consists of two main sites, Gaithersburg, MD and Boulder, CO, that were built to mostly mirror each other and connect internally. There is a small group of NIST staff stationed at NIST Ft. Collins. The network at Ft. Collins is a logical extension of the Boulder network using an Internet VPN.

There is a small group of NIST staff stationed at the NOAA HML facility in Charleston, SC. Users at HML connect to NOAA switches that offer a VLAN that is a logical extension of the Gaithersburg network.

Over the past couple of years, NIST has switched WAN carriers to NOAA NWAVE for most WAN services.

Internal LAN (behind the main firewall)

The campus internal LANs consist of Cisco equipment organized in the typical 3-layer model. The campus core layer consists of a pair of Cisco Catalyst 9K multi-layer switches. The campus core switches uplink to the main firewalls with 10G connections through a set of L2 FireEye appliances that are positioned inline. The distribution layer consists of Cisco Catalyst 9K multi-layer switches that uplink to both campus core

switches at 40G (Gaithersburg) or 100G (Boulder). The access layer consists of Cisco Catalyst 9K switches that uplink to two distribution switches at 10G. A few switches still uplink at 1G. Data center access switches consist of Nexus 5K (Gaithersburg) and Nexus 9K (Boulder). Users connect to access switches with 1G copper connections. Servers in the data center connect to access switches at 1G, 10G, and 40G. The wireless access layer consists of Cisco 5520 Wireless LAN Controllers with Cisco 9120 LWAPP access points. The wireless network offers several WLANs for staff, devices, and guests.
NIST has recently deployed a "Science Network" to provide higher speed connectivity for scientific servers and storage systems deployed in labs or the data center. The Science distribution layer consists of a pair of Nexus 9K switches that uplink to campus core at 40G. Science access switches are also Nexus 9K switches that are positioned in buildings that need them. Science switches uplink to the Science distribution routers at 100G. End users and data center servers connect to the Science switches over single mode fiber at 10G, 40G, or 100G. The Science Network is configured for jumbo frames.

External LAN (in front of the main firewall)
The campus external LANs have a second pair of Cisco Catalyst core switches positioned just outside the main firewalls. Secondary external firewalls also connect to the external core including the E-NIST DMZ firewall, VPN firewall, and the Shared Services firewall. There are also several perimeter routers that connect to the external core.

Internal WAN
NIST Internal WAN is built on the NOAA NWAVE network. All three sites connect to the NOAA NWAVE network with a pair of 10G connections. Different classes of traffic logically separated and routed through separate VRFs. NIST recently migrated its path to cloud CSPs (AWS, Azure, and Google) to the NWAVE network. Using the same internal WAN links, the sites can reach each other as well as the CSPs. NWAVE connects NIST to its CSP via the Internet2 Cloud Broker service. NIST currently pays for a 5G link to AWS East Commercial, and 2G links to AWS East GovCloud, AWS West Commercial, and AWS West GovCloud, respectively. We are just starting to build into Azure and Google for a NIST project.

External WAN
Gaithersburg and Boulder connect to the Internet with a pair of Cisco routers (ASRs in Gaithersburg, Catalyst in Boulder). The Internet routers connect to a separate pair of 10G NOAA NWAVE links that route Internet traffic through the NOAA TICAP service. There is also a separate VRF for routing traffic to other DoC bureaus over NWAVE. Gaithersburg has a third perimeter router connecting to the DoC TLS backbone at 10M. Gaithersburg, Boulder and Ft. Collins each have a separate "Time Routers" to support the NIST Time Service. These Time Routers connect the NTP Time Servers to the Internet through separate, 1G, non-MTIPS Internet links from Verizon.
Ft. Collins connects back to Boulder via an IP-SEC VPN running on top of a 100M, non-MTIPS Internet link from Verizon. Ft. Collins Internet traffic is routed out the Boulder Internet links.

Internet traffic for NIST HML is routed back to Gaithersburg over NWAVE first.

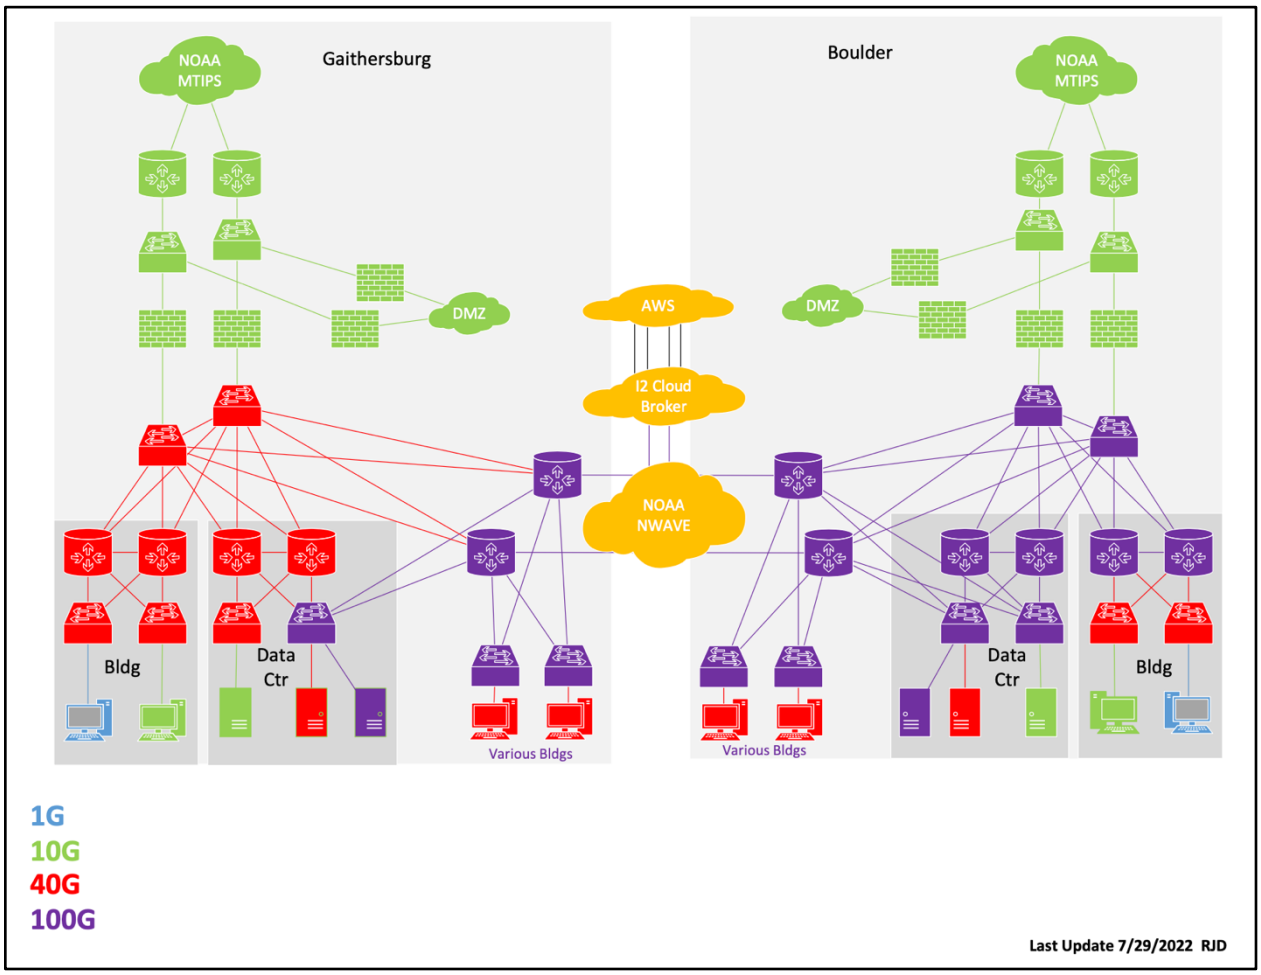

Figure 2: NIST Network

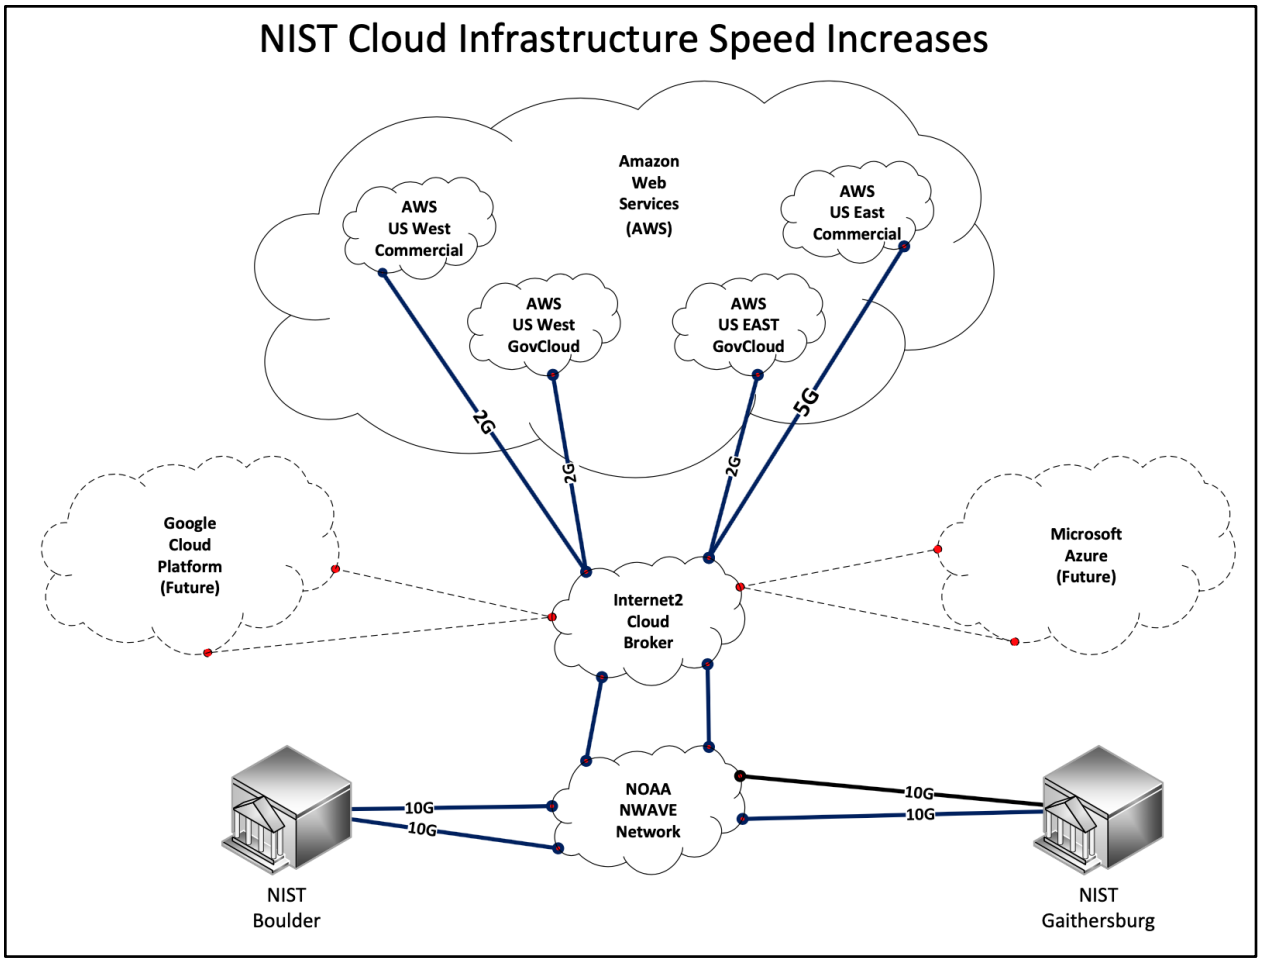

Figure 3: NIST EXT Links

4.8.4.2 Computation and Storage Infrastructure

The following HPC systems are managed by OISM on behalf of the NIST Labs: Raritan, Enki, Nisaba, Baker-Jarvis, and Simba. The following systems are managed by OISM on behalf of the Engineering Lab and their use is restricted to those approved by that Lab: Hercules, Blaze, and Burn. These are listed in Figure 4. The abbreviations used below are GCF (e.g., Gaithersburg Computing Facility), BCF (e.g., Boulder Computing Facility), and 224 (e.g., a building on the Gaithersburg campus because these systems are not inside of either of the aforementioned datacenter environments).

Available To...	Name	Cores/GPUs	Location	Approx. Rack Count	Managed By	Paid By
All NIST Staff	Raritan	18028 Cores, (36056 threads) / 0 GPUs	GCF	37	OISM/RSO	MML/PML/ITL
All NIST Staff	Enki	520 Cores, (2080 threads) / 52 GPUs	GCF	4	OISM/RSO	MML/ITL
All NIST Staff	Nisaba	220 Cores, (4400 threads) / 41 GPUs	GCF	4	OISM/RSO	ITL
All Boulder Staff	Baker-Jarvis	896 Cores, (1792 threads) /0 GPUs	BCF	4	OISM/RSO	CTL
All Boulder Staff	Simba	384 Cores, (768 threads) / 24 GPUs	BCF	1	OISM/RSO	MML
EL & SPO Staff	Hercules	3424 Cores, (4192 threads) / 0 GPUs	GCF	5	OISM/RSO	EL
EL/733	Blaze	1144 Cores/ 0 GPUs	224		OISM/RSO	EL
EL/733	Burn	432 Cores/ 0 GPUs	224		OISM/RSO	EL

Figure 4: NIST Centrally-Managed High Performance Computing Systems

In addition to on-premise compute systems, OISM also supports a NIST environment in Amazon Web Services (AWS). In AWS, OISM offers NIST customers the ability to self-service their own AWS Elastic Cloud Compute (EC2), Elastic Block Storage (EBS) instances, to store files in Simple Storage Service (S3) buckets, and to take advantage of many other (but not all) AWS commercial service offerings.

For on-premise storage, OISM supports a 1-PB Isilon research data storage system in Boulder and is in the process of deploying a similar system in Gaithersburg. There are also general-purpose NetApp data storage systems available for project and organizational directories. NIST recently completed "home drive" migration from on-premise Network Attached Storage to Microsoft's OneDrive for the vast majority of users (those who were not migrated to OneDrive were instead migrated to Google Drive for technical reasons). In addition to OneDrive and Google Drive, OISM supports Kiteworks Secure File Transfer, Box and AWS S3 for data exchange.

For off-premise storage, OISM supports the aforementioned NIST enclave in AWS allowing NIST projects to choose among AWS' S3 offerings according to each project's needs. In addition, OISM makes Microsoft OneDrive and Google Drive available. Finally, the OISM offers a network-based backup service from CommVault enabling self-management of backups for NIST endpoints; those backups are ultimately stored in AWS. For NIST's key financial, human resource, grants, and acquisitions systems and data, backups are stored in a secure off-site facility (i.e., neither on-premise nor in a cloud service provider's environment).

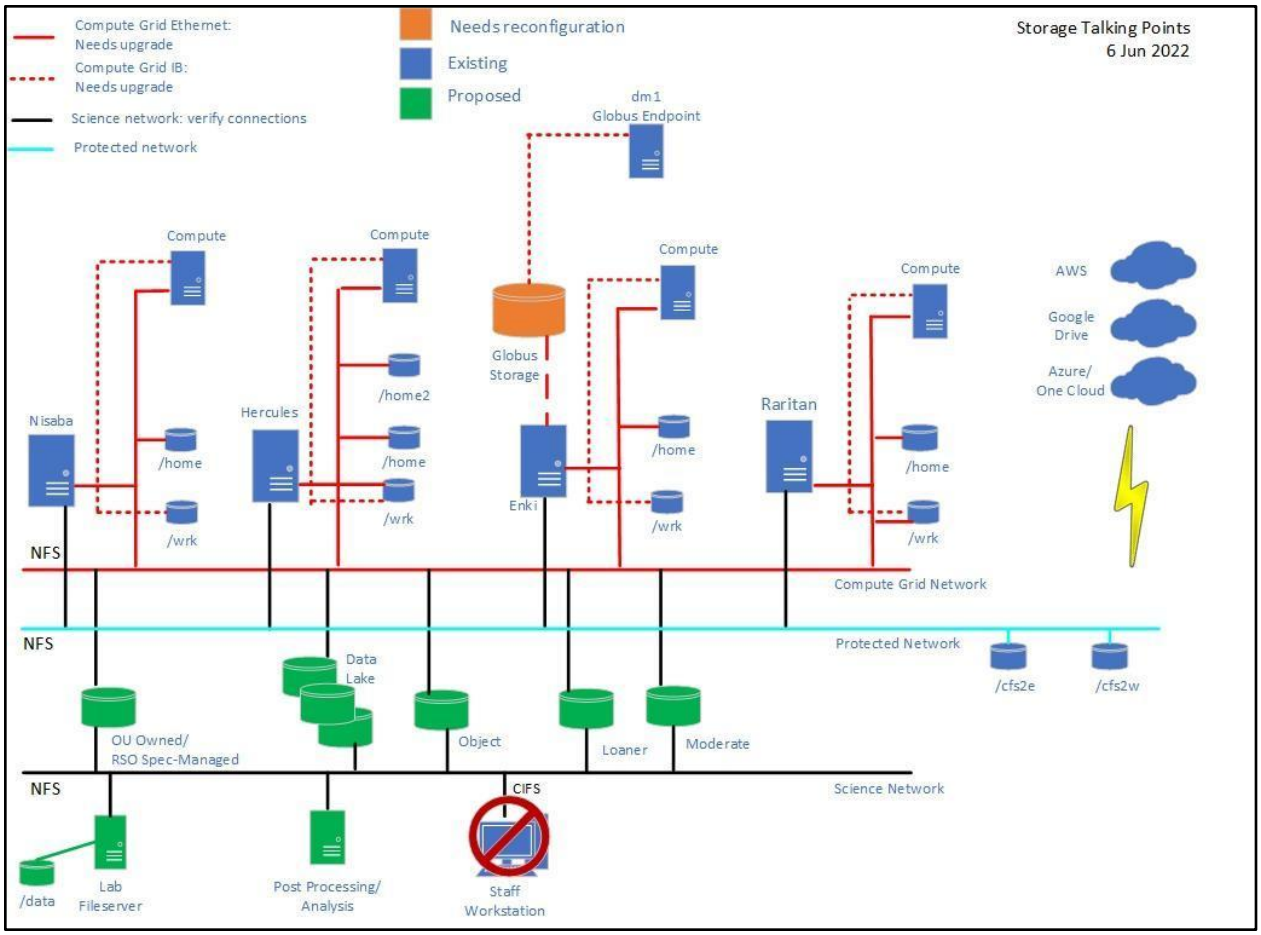

Figure 5: NIST Storage

4.8.4.3 Network & Information Security

Perimeter Security

The first line of defense from Internet attacks are inbound/outbound ACLs at the perimeter routers. These ACLs provide the basic, best practice packet filters that are recommended on any Internet routers.

Public Firewalls

- The second line of defense comes from several public firewalls
- The E-NIST firewall protects the NIST public-facing DMZ subnets
- The VPN firewall protects the remote VPN connections coming into NIST
- The SSDoC firewall protects the Share Service infrastructure that includes Enterprise Continuous Diagnostic and Mitigation (ECDM) service that NIST hosts for the entire DoC
- The Main firewall protects the internal LAN

Internal Firewalls

- The third line of defense comes from several internal firewalls

- The infrastructure firewall protects the production traffic to the infrastructure servers like email and file servers
- The Management firewall protects the privileged traffic to the infrastructure
- The Admin/Financial firewall protects the central administrative and financial application servers
- The Database firewall protects central database servers
- The REN (Research Equipment Network) firewall protects sensitive research equipment that are isolated on separate VRFs and VLANs
- The Cloud firewall protects traffic to and from our CSPs

Security Appliances

- Proxy Servers
- IDS/IPS
- Full Packet Capture
- Network Access Control

REN Networks

REN stands for Research Equipment Networks. This is a network for devices that can no longer be managed due to obsolescence or because of specialized research software that requires a fixed configuration. For example, NIST has many pieces of expensive scientific equipment that are controlled by a PC. These PCs have special hardware and software drivers to control the scientific equipment. In many cases, the PCs cannot be upgraded or patched because of the special hardware and driver code. Over time, these PCs become insecure and need additional security controls to prevent attacks.

In response, NIST built special "REN" networks that are isolated from the campus at L3 with VRFs and isolated at L2 with private VLANs. These REN networks have a default route across the campus backbone to the REN firewall that serves as a single choke point for all REN networks.

4.8.4.4 Monitoring Infrastructure

NIST performs monitoring across the organization in order to detect different types of issues. For example, within the HPC environment several monitors are in use to provide insight to the compute environment for health, performance/capacity planning, security, and compliance. Users have access to performance/capacity and a health dashboard for systems. Additional monitors are being deployed as time allows to provide increased insight.

- XDMoD – metrics for user jobs and hardware utilization based on parsed logs from SLURM.
- Observium – SNMP based polling to gather/display system performance data for CPU, memory use, network I/O, disk I/O, etc.
- SLURM Reports – Customized reports using SLURM tools to provide information on job efficiency.

- Slurmy – Home grown software that provides a GUI view to cluster node allocation/job status based on output from SLURM commands with hooks into Observium to provide job specific performance information.
- Icinga – Nagios fork/improvement to monitor node health and provide alerts to incidents.
- Splunk – log file scanning/notification for system health and security events.
- Starfish – Provides user storage (i.e., quota) and volume information for capacity planning.
- Tenable – Regularly scheduled scans performed by NIST security.
- Ansible – NIST standard playbooks regularly run to ensure compliance.
- The NIST Network Team uses 3 primary network monitoring tools; WhatsUp Gold, MRTG, and Splunk. WhatsUp Gold is a commercial network monitoring platform. It is used for by the network team for network mapping/visualization, performance monitoring, alarms/alerts, NetFlow monitoring, and configuration backups. Multi Router Traffic Grapher is made available to the NIST users for general network health and real time graphs of critical network interfaces. Splunk is used by the network team for logging and alerting.
- Homegrown software is used to identify logical/physical locations of specific MAC/IP addresses anywhere on the network.
- Solarwinds Network Engineers Toolkit is also used for network troubleshooting.

4.8.4.5 Software Infrastructure

The OISM provides access to a number of commercial scientific software applications for all NIST users. The NIST Scientific Computing Steering Group (SCSG) determines what scientific software applications are provided by the OISM and how many licenses for each application are to be made available to NIST.

Most of the scientific software applications supported by the OISM do not have site licenses - instead, a fixed number of licenses are provided (the exceptions are NIST's enterprise licenses for The MathWorks products and for National Instruments' LabVIEW software).

These are the commercial scientific software applications currently maintained by the OISM as determined by the SCSG: Ansys, Autodesk, COMSOL Multiphysics, Intel Composer, L3Harris Geospatial ENVI and IDL, Maplesoft Maple, Mathworks MATLAB & Toolboxes, Molpro, National Instruments LabVIEW, OriginPro, SolidWorks, Scienomics MAPS, & Wolfram Mathematica.

In addition, the OISM supports Globus for NIST users; Globus is a secure reliable service for research data management as well as a platform to build data focused applications and services. With Globus, NIST researchers can easily move, share, and manage data sets of any size.

The following software is provided on the Raritan HPC system:

- Anaconda Python (2 and 3)
- CUDA (only support on the GPU nodes n[4000-4002])

- Gaussian: 09 and 16 (with Linda)
- Open MPI
- Nvidia HPC SDK
- VASSP (License restrictions apply)

The following software is provided on the Enki HPC system:

- CP2K-v7.1 with GPU support
- CPU-only psmp
- GPU-accelerated psmp
- LAMMPS-v3Mar20 with GPU support

The following software is provided on the Nisaba system:

- Make (GNU make 4.1)
- Cmake (3.16.4)
- GCC (9.3.0)
- Open MPI (4.1.1 Environment Modules)
- NVIDIA CUDA Compiler driver (11.0 Environment Modules)
- NVIDIA CUDA Deep Neural Network Library (8.3.2)
- NVIDIA Collective Communications Library (2.11.4)
- MATLAB (2021a)
- Mumax (3.1.0)
- ParaView (5.9.1)
- GROMACS (2020.5)

4.8.5 Organizational Structures & Engagement Strategies

The following sections outline the organization structure, and engagement with the research community.

4.8.5.1 Organizational Structure

The NIST Chief Information Officer (CIO) is the principal advisor to the NIST Director on the effective application of information technology. The CIO's responsibilities include planning, directing, and implementing information technology services for NIST in pursuit of the NIST mission. The NIST CIO implements the provisions of the Federal Information Technology Acquisition Reform Act (FITARA) for NIST through authority formally delegated by the Department of Commerce's (DOC) CIO on an annual basis. The NIST CIO is also responsible for compliance with the Federal Information Security Modernization Act (FISMA) and for compliance with numerous mandates issued by the Office of Management & Budget and by the Department of Homeland Security. The NIST CIO organization routinely reports on policy compliance and on investment management through DOC. The NIST CIO reports to the NIST Associate Director for Management Resources – whose purview also includes NIST's financial, human resource, and facilities organization – thereby providing a unified approach to management of the services supporting NIST's operations. The NIST CIO has line authority for all centralized NIST IT services in Gaithersburg and Boulder and is responsible for implementing an effective IT security program at NIST.

See https://www.nist.gov/oism-organization for an organizational chart and https://www.nist.gov/oism for a description of organizational functions.

Several individual NIST Operating Units supplement the OISM's Information Technology services with OU-specific services. These services span the spectrum from Linux desktop support to full-fledged public-facing services. In between are examples of research computing support, e.g., management and maintenance of OU-owned HPC clusters, laboratory information management efforts, and so on.

4.8.5.2 Engagement Strategies

- NIST HPC Survey (~2018)
  - OISM, in conjunction with several HPC stakeholders, conducted a NIST-wide survey intended to assess NIST interest in HPC and to gather feedback on the state of NIST's offerings.
- NIST Research Computing Survey & Focus Group Meetings (~2020)
  - OISM, again in conjunction with several stakeholders, conducted a NIST-wide survey intended to assess NIST interest and gather feedback on multiple research computing components (i.e., HPC, scientific software, research data storage, and consulting/software development).
- NIST Lab Automation Survey & Focus Group Meetings (~2022)
  - OISM conducted a NIST-wide survey intended to assess NIST technical staff preferences regarding LabVIEW licensing, other lab automation enablers, and consulting/software development.
- SCSG Governance (since ~2005)
  - A Scientific Computing Steering Group comprised of a representative from each NIST Laboratory along with a representative from OISM serving as de facto facilitator meets ~monthly. Though originally conceived with a scope inclusive of HPC, in practice the group has focused entirely on scientific software after support models for HPC were devised in the first year. The group determines what software comprises the NIST-wide, shared scientific software portfolio.
- Collaboration Tools Steering Group (since 2016)
  - In March 2016, the ADLP established a Collaborative Tools Steering Group to determine whether collaboration tools that were being assessed for an OU-specific use case should be deployed for NIST-wide use. Subsequently the group has convened periodically to prioritize the security assessment/deployment of collaboration tools to be made available NIST-wide. A sampling of the tools that have been "authorized" by the group includes Box, Github, Overleaf, Slack, and Trello. The complete list is available if EPOC is interested.
- Research Computing Infrastructure Team (2017-2019)
  - A team comprised of representatives from each of the NIST Labs with an interest in NIST-wide HPC and high-speed networking was convened explicitly to plan for the acquisition, configuration, deployment, and support for the Enki AI/ML cluster, for the storage

that was purchased in conjunction, and the (limited) 10G LAN endpoint connections. Once the Enki AI/ML cluster was operational, and an informal support structure was devised along with network connections to specified locations on the Gaithersburg campus were contracted, interest in and purpose for the group waned.

- RSO/SST Meetings with HPC PoCs (monthly)
    - OISM's Research Services Office's Scientific Systems Team holds monthly meetings with PoCs for Raritan, Hercules, etc. to discuss issues, maintenance needs, and the like.
- IT Security Working Group Meetings (Monthly)
    - The NIST CISO (OISM IT Security & Networking Division Chief) holds monthly meetings with NIST IT Security Officers to discuss IT topics that impact or are impacted by NIST IT Security practices, policies, etc.
- SCMMR/Network Roadmap Steering Group (2020 - Present)
    - A subcategory of NIST's appropriated funding is specifically legislatively purposed for improvements/repairs to NIST's infrastructure. NIST leadership has successfully made the case that several areas of NIST's IT infrastructure should be treated no differently than NIST's plumbing or electrical infrastructure and hence can be improved through use of these funds. NIST's Facilities Management leadership, the OISM Network architect, and representatives from NIST Laboratories meet periodically to prioritize which needs in NIST's network roadmap should be funded in each annual cycle.
- Research Computing Advisory Committee Governance (2022)
    - NIST's Research Computing Advisory Committee was formed as part of an overall effort to re-vamp/rekindle IT Governance at NIST. This committee provides strategic leadership and vision to accelerate NIST research through advanced research computing initiatives. This committee also serves as a consultative body to the Lab Directors and ADLP. Committee members collaborate with OISM to leverage technology advances and expand the impact of NIST research infrastructure. Committee scope includes but is not limited to advanced networking, high performance computing, big data, specialized research applications, and computationally-intensive research support across a wide range of disciplines.
- NIST 2020 Strategic Plan
    - With participation of staff and management across NIST over the course of a year, a new Strategic Plan was released in 2020. Among the objectives identified were "Upgrades to NIST's IT Infrastructure" and "Facilitating Next-Gen Research Data Infrastructure." Though NIST Senior Management did not select those objectives for the initial work plan, OISM has attempted to make progress on those objectives as resources allowed.

4.8.6 Internal & External Funding Sources

NIST as a whole is funded through congressional appropriations, through collections of fees for services it sells, and through reimbursable agreements with other agencies. OISM (like the rest of the Management Resources organization) is funded via a tax on related to the funds identified above. In addition, OISM receives a relatively small congressional appropriation and collects fees for services. Nearly all of the OISM's Research Services Office's work is funded via fee-for-service.

4.8.7 Resource Constraints

OISM management recognizes that there has been a period of time where NIST underinvested in the network infrastructure in place at both campuses. The result of that underinvestment impacted NIST's researchers' ability to move data from lab to office to collaborator, move data across campuses, and exchange large volumes of data with external collaborators. The network physical plant in Boulder (i.e., wiring/cabling) was particularly out-of-date as compared to the Gaithersburg campus. Network hardware at both campuses was not being refreshed in a timely manner as well. In approximately 2017 NIST Senior Leadership began providing funding to address network infrastructure deficiencies and since that time OISM has been "catching up" as funding allows. Still there are too many areas on both campuses where network performance issues still need to be remediated. Similarly OISM has made significant strides in addressing NIST's network connectivity between Boulder & Gaithersburg and from each campus to external networks as well as to cloud service providers. NIST researchers are still frustrated by issues they perceive to be network-related; OISM's ability to diagnose and resolve the many granular issues that impact specific use cases is still quite limited owing to lack of staff resources - particularly staff resources who can diagnose and resolve technology issues from end-to-end.

There are also resource data storage and data management issues that OISM lacks the resources to address in the way they should be addressed. This is less of a storage system/network/software management problem and instead more of a staffing issue; OISM lacks the staff resources to perform coordinated planning with the NIST Labs, define satisficing research data management services in conjunction with the Labs, to define a service sustainment funding model, and to support operational maintenance of research data storage services that are sorely needed at NIST.

Another impediment that OISM faces is its ability to attract and retain qualified staff in both the Boulder and Gaithersburg areas. Both locations are replete with private sector entities that can offer far more compensation than the government, in newer facilities, likely with more flexibility, and perhaps better opportunities for career growth than NIST can. The one notable exception to this phenomena is in OISM's ability to attract IT security assessors; it appears that candidates for these positions view working for a few years on OISM's Assessment team as a useful stepping stone to greater opportunities in the private sector. So here the problem is more of retaining staff and not so much in attracting staff.

Two more constraints that are worth noting are (1) NIST's deteriorating physical facilities and (2) NIST's funding model for IT services that directly support NIST's Lab mission. On (1) trying to deploy and maintain IT infrastructure in facilities that were designed/constructed in the 1950's and 60's is very difficult. Plumbing leaks, roof leaks, lack of HVAC control, and inadequate service access eat into OISM's resources' productivity. On (2) most of the research IT services offered by OISM are fee-for-service meaning that individual NIST Labs and/or projects must pay OISM for the specific services utilized. This greatly limits the scope of the services provided and results in no ability whatsoever to plan for future improvements common to a larger population of the NIST Labs.

4.8.8 Outstanding Issues

No additional issues to report.

4.9 Material Measurement Laboratory (MML) IT Service Team

Content in this section authored by Ann Leith, NIST

4.9.1 Use Case Summary

The MML IT Service Team provides specialized research-focused IT support to the Material Measurement Laboratory (MML). MML is one of the largest organizations within NIST, it has ~800 staff members who are located at both the Gaithersburg and Boulder campuses as well as a number of other smaller satellite locations.

The team is funded solely by MML. We provided specialized support that is outside of the boundaries of what is provided by NIST's central IT organization, OISM.

4.9.2 Collaboration Space

MML researchers collaborate with numerous external organizations and individuals - too many to list. Some of the more notable and long-standing relationships are with Argonne National Laboratory, Brookhaven National Laboratory, and the Institute for Bioscience and Biotechnology Research at the University of Maryland.

4.9.3 Capabilities or Special Facilities Offered

In addition to general support, the MML IT Service Team manages a Dell-EMC Isilon storage array with 1.2 PB of available storage. We have created a service in conjunction with this storage called 'Data Plumbing', which is a collection of tools we have written that provides automated data transfer from laboratory instruments or intermediate data collection locations to the Isilon. Many of our instruments are located on the Research Equipment Network (REN), which is an internal network with very restricted connectivity to other NIST resources like file storage. Although it is possible to move data from the REN to the general internal network, many researchers found this process difficult and would resort to thumb drives and other sneakernet type solutions.

Data Plumbing is an effective solution to the barriers that prevent data transfer from individual laboratories to our central storage solution, but it is a manual and very labor-intensive process to connect each individual instrument. To date, we have connected about 80 instruments. In MML alone, that represents less than 20% of our instrument inventory. There are many other parts of NIST that would benefit from this service, but to my knowledge no other support group has sufficient resources to offer this service.

4.9.4 Technology Narrative

The following sections outline the technology footprint for this use case

4.9.4.1 Network Infrastructure

MML does not own or manage any networks. See OISM response.

4.9.4.2 Computation and Storage Infrastructure

MML does own a small cluster, CTCMS, that is covered in Section 4.10 Center for Theoretical and Computational Materials Science (CTCMS). There are also a number of NIST-level clusters that should be covered in the OISM response.

4.9.4.3 Network & Information Security
MML does not own or manage any networks or information security outside of the standard NIST resources. See OISM response.

4.9.4.4 Monitoring Infrastructure
None within the boundaries of MML

4.9.4.5 Software Infrastructure
Data storage or transfer tools in use in MML include Google, Box, Microsoft OneDrive, Globus (use is not widespread - this service offering is not mature), Data Plumbing is based on Rsync and custom code, we also are participating in the OISM-led Starfish pilot

4.9.5 Organizational Structures & Engagement Strategies
The following sections outline the organization structure, and engagement with the research community.

4.9.5.1 Organizational Structure
We are a small six-member team located within the Office of Operations, part of the Laboratory Office of MML. Most of the team's interactions with other support groups at NIST are informal, team members are encouraged to seek out and nurture relationships with key personnel in OISM. The team also always has representation on NIST-level working groups, like the home drive migration team or other project teams as appropriate. A relatively small part of our support work comes in the form of tickets in ServiceNow, the NIST IT ticketing system.

4.9.5.2 Engagement Strategies
Each team member is assigned to support one or more of the research divisions within MML. Requirements from individual users and projects bubble up through the team.

The team lead of the MML IT Services team is also the senior IT advisor for MML and sits on the MML Management Team. Each year, MML identifies a number of focus areas. IT is not a primary focus of these activities but occasionally there is a clear relationship between a program focus and IT capabilities. The Data Plumbing project grew out of a more general data management focus area.

4.9.6 Internal & External Funding Sources
Funding is provided by NIST.

4.9.7 Resource Constraints
The team is resource limited - Data Plumbing is a very labor-intensive project and we don't have sufficient staff to roll it out quickly. We are also constrained by the generally poor state of NIST IT infrastructure. Poor network performance, lack of compute, and lack of a cohesive storage strategy, all impact our effectiveness. Researchers need help with locating appropriate storage locations, as MML staff not infrequently discover that individuals have purchased a desktop NAS to store their data and then have problems with that device. One of the constraints in identifying a suitable storage location can be

network performance. NIST's external links to general internet and cloud services are often too slow to handle very large data transfers reliably. MML management is very good about providing funding for shared IT resources within our organization, but many of the problems we face are too large to be solved within a single operating unit and need to be tackled at the NIST level.

4.9.8 Outstanding Issues

No additional issues to report.

4.10 Center for Theoretical and Computational Materials Science (CTCMS)

Content in this section authored by Andrew Reid, NIST

4.10.1 Use Case Summary

The primary mission of the CTCMS is to reduce friction and provide a high degree of operational flexibility in support of high-performance and research computing for its users. In terms of user traffic, the principal use cases are the running of high-throughput codes for density functional theory and molecular dynamics. We also support a number of less common use-cases, where hands-on attention from a knowledgeable system administrator who is also a practicing scientist adds value. These include building and installing custom software, helping users stand up novel workflows to support new use-cases, and providing a venue for novel interfaces to research computing. We were also early adopters of high-speed network connectivity.

In terms of data architecture, we are custodians rather than owners of data, and provide several large-capacity storage systems and data protection services. The CTCMS has two primary file systems, both NFS-mounted and shared across almost all CTCMS resources (cluster nodes, workstations, service nodes, etc.), with a two-tiered data-protection scheme, including both a daily snapshot of the full file system, to protect against failure of the primary system, and a daily incremental back-up capable of point-in-time restoration, to protect against accidental deletion. On a smaller scale, we operate a public-facing Globus end-point, and provide a means for users to move data to a public-facing web server, which is well-adapted to high-volume or machine-generated (or both) web content.

4.10.2 Collaboration Space

The CTCMS is a mostly inward-looking facility, seeking to meet team HPC users' needs for compute and data storage. We have some collaborations with the Platform for Network Innovation, which provides high-speed optical networking infrastructure off the NIST site, which we use to support our Globus end-point, which is currently in use in support of an additive-manufacturing benchmarking program as a data transfer and aggregation venue. The PNI is currently operated out of the Communications Technology Laboratory.

We also have some collaborations with the Information Technology Laboratory on innovative use-cases for HPC, in which users can assemble image-processing workflows (“pipelines”) on a NIST-facing web platform, and then push those pipelines to the CTCMS public cluster for execution.

Closer to the thrust of the question, an “off-ramp” to more sophisticated HPC capabilities is a long-standing wish-list item for us, we think it would take the form of a relationship with an HPC center outside of NIST, where they might support higher-capacity HPC systems with better interconnect fabrics, clusters that are able to handle larger jobs using more total memory, or systems able to support a much higher degree of parallelism. This type of relationship would enable us to possibly “pre-vet” our users so that they can scale

up/out without as much friction as if they were complete newcomers to HPC or to the scale-out resource.

4.10.3 Capabilities or Special Facilities Offered

The CTCMS has a number of capabilities above and beyond the core functionality of HPC cluster operations and data storage and protection.

These arise from saying “yes” when an MML researcher says, “I wonder if research computing can do X”, for any X. Use cases to date have involved novel workflows (web-based dispatch to the HPC cluster, for instance) or special-purpose workstations optimized for certain workloads. Criteria for taking on a project are informal, but would be influenced by the value to the researcher and the degree of difficulty and amount of new equipment that would be required.

As already mentioned, we operate a number of systems on the PNI, including a public-facing Globus end-point.

We operate a number of “private” clusters, consisting of head nodes, modest data storage, and compute nodes and interconnection fabric dedicated to a small, well-defined group of users. These clusters support e.g. some emerging computational efforts for soft materials, and the JARVIS high-throughput DFT computations.

We also operate a public-facing web server that is well-adapted to high-volume or machine-generated web content, and also supports the long-term validity of published URLs. This contrasts with NIST's institutional approach to public web services, which revolves around manual point-and-click construction of pages through the Drupal interface, and frequently deprecates public URLs.

The CTCMS public cluster has an attached MongoDB instance that supports the Interatomic Potentials Repository project. Jobs run on the cluster can automatically connect to the MongoDB instance on completion and record their findings in a machine-readable format which facilitates later analysis and dissemination.

We also run some NIST-facing web applications, notably an instance of the Configurable Data Curation System in support of the tracking of physical samples used in experiments. This is a nod towards complementing the NIST Laboratory Information Management System or future electronic lab notebooks.

We also run license servers for some licensed commercial software, notably the Abaqus finite-element package, used for solid mechanics modeling, and the ThermoCalc software, used in support of CALPHAD phase-based thermodynamic and kinetic property computations.

We also provide informal consultation for users wishing to acquire their own HPC capability -- this was the starting point for both of the private clusters operated by the CTCMS.

CTCMS staff are also heavily involved in the HPC Carpentry educational effort, a program for up-skilling novice HPC users, using the pedagogical methods and web infrastructure of the broader Carpentries (Software Carpentry, etc.) organization.

4.10.4 Technology Narrative

The following sections outline the technology footprint for this use case

4.10.4.1 Network Infrastructure

There are two major networks of importance to the CTCMS. They are the NIST campus protected network, operated by NIST OISM, and the CTCMS machine-room network, which we operate. The NIST protected network connects CTCMS workstations in user's offices, and administrative systems (data servers, etc.) to the main CTCMS infrastructure, and is firewalled off from the internet at large, but is still generally considered relatively untrusted. The CTCMS machine room network lives entirely within one locked room, and is highly trusted, allowing host-based authentication over SSH between cluster nodes, and ease of access and set-up for CTCMS admins. Additionally, private clusters, and each of the seven racks of the CTCMS public cluster, hosts an Infiniband interconnection fabric, providing high-bandwidth, low-latency connectivity within the clusters or racks.

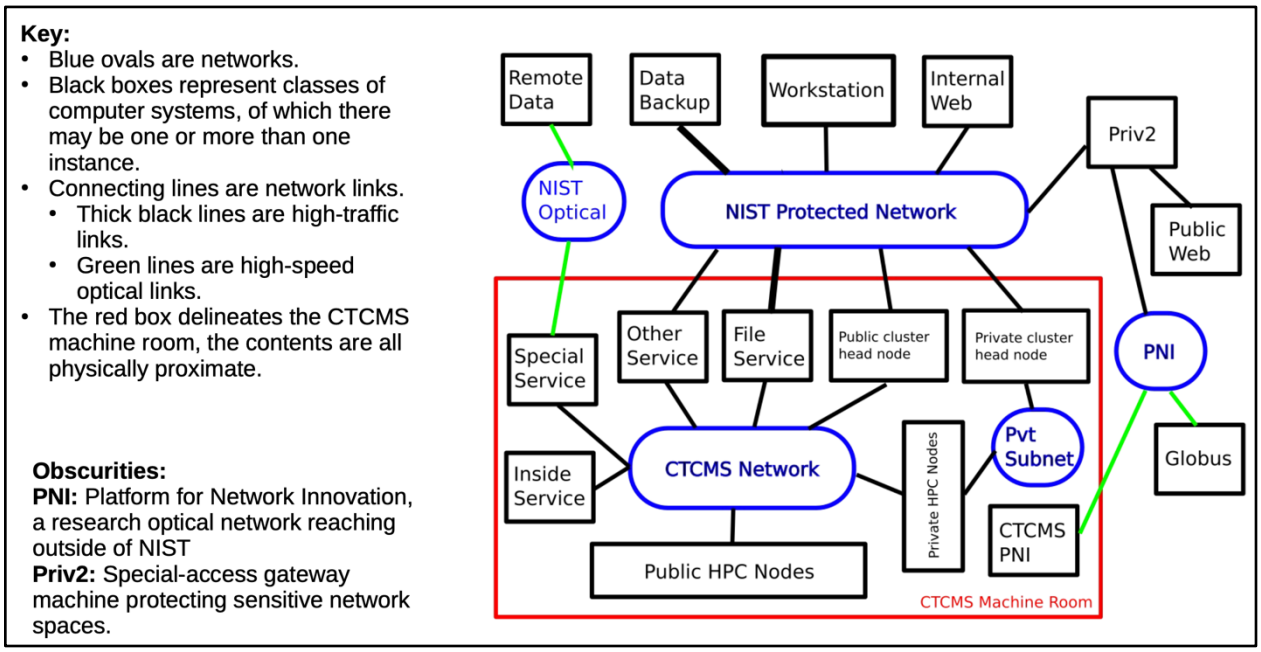

Figure 6: CTCMS Network Diagram

4.10.4.2 Computation and Storage Infrastructure

The CTCMS public cluster is divided into 7 racks, each hosting around computational nodes, and sharing an Infiniband interconnection fabric. This is the primary HPC workhorse of the system. In addition, private clusters in the CTCMS machine room provide computational services to smaller user communities within NIST.

CTCMS data is mainly stored on two ~100 TB capacity primary servers, and made available over NFS system-wide. This data is protected by nightly copies to “mirror” servers in the CTCMS machine room itself, and incremental back-ups to servers in another location on campus.

The CTCMS PNI infrastructure consists of the Globus end-point, and several supporting servers, all with high capacity (10s of TB to ~200 TB) storage services. The precise use-case of this infrastructure is still evolving because it is not yet clear how to arrange the resources to meet their particular need. It's clear that the equipment and capability NIST has stood up is a move in the right direction, but how users will access the data, which users need what kind of access, and how/whether to transport the data elsewhere at NIST are still open questions.

4.10.4.3 Network & Information Security

A reasonably sharp boundary is drawn between the CTCMS machine-room network and the NIST protected network. Inbound TCP connections to the machine room network are not permitted. We also do not allow password authentication over SSH at all on CTCMS systems.

CTCMS operates its own Kerberos and LDAP system, for user authentication and account management. This actually simplifies management of the system for administrators and provides the convenience of “single sign-on” for users. This system does not inter-operate with NIST OISM supported account management. This is a minor issue for users, but probably a bigger issue for CTCMS system operators, who have to do a bit of account management on top of the more science-focused issues.

Our weakest security component is the widespread use of NFS for file-system sharing. We have previously investigated promising replacement candidates for NFS file systems, for both security and performance reasons, but have never found one that meets all of our file-system needs robustly. This is a critical concern to some staff.

4.10.4.4 Monitoring Infrastructure

All CTCMS systems aggregate log data on a central log server, which filters out routine events and forwards non-routine events to admins via e-mail.

For HPC system operations, the CTCMS public cluster machines have Prometheus monitoring software on them, which monitors basic health parameters (CPU load, memory usage, etc.), and displays all of the info on a Grafana dashboard. Some private clusters previously used the Ganglia monitoring tool, but this has fallen out of favor.

Real data-driven introspection on HPC operations is a wish-list item, and we are aware of a number of interesting and useful free software tools from TACC and others to do some sophisticated monitoring of cluster usage, including more detail on network traffic, load by software package, network and disk operations, and so forth, but we have not found the time to deploy any of these tools. We are also aware of OISM ORS's use of xdmod to good effect, and we might benefit from following their lead.

4.10.4.5 Software Infrastructure

We principally run scientific codes, many of which are built from open or commercial source, and some fully commercial closed-source codes. Virtually all of the CTCMS infrastructure is open-source, relying heavily on Linux operating systems and tooling.

4.10.5 Organizational Structures & Engagement Strategies

The following sections outline the organization structure, and engagement with the research community.

4.10.5.1 Organizational Structure

In our facility, we take pride in saying "yes" to requests for new capabilities, but also have a strong preference for incremental approaches, building up capabilities step-wise with new sub-capabilities that build towards the goal. We have this flexibility because of the relatively small scale at which we operate -- system operators (there are two) are fully aware of all aspects of the system at all times, and so can pivot quickly to new operational schemes or workflows.

4.10.5.2 Engagement Strategies

We had good engagement with our research community in two noteworthy instances recently.

The first of these is in standing up an HPC capability for the soft-materials scientists in our division. Initially, our role was to just host the equipment, but it quickly evolved into an HPC consulting role, helping the investigators build up good workflows in the HPC environment, and consulting with the system owner on hardware upgrades and changes in network connectivity.

The second was standing up the web-based workflow pipeline for cell-imaging investigations. This system uses a web-based system to graphically construct workflows and then dispatch them to the HPC resource. CTCMS hosts the web server and the HPC resource, and worked closely with colleagues in the Information Technology Laboratory to build the bridge over which workflows built in the web tool can be dispatched to the HPC system via a Kubernetes-based series of intermediate micro-services.

4.10.6 Internal & External Funding Sources

CTCMS relies on the largess of Division 642 at NIST, and additional Materials Genome Initiative funding from the MML lab office, in equal measure. Currently, 3/4 of an FTE is committed to address CTCMS operations but in practice the time spent exceeds that level on a regular basis.

4.10.7 Resource Constraints

We have an on-going issue with data transport. As HPC gets more powerful, and data volumes get larger, it's harder and harder to manage data sets, including for such basic activities as data protection. NIST as an institution has an excellent network roadmap and will have fatter pipes in the near future, which will certainly help with this. “Data protection,” in this context, means protection against loss due to equipment failure or

accidental deletion, it does not refer to protection against unauthorized access. Our means of data protection is copying data to back-up servers, which requires transport. Data volumes have gotten larger over the years, but our transport channel has been gigabit ethernet for many years, so the time it takes to move data has gotten larger. Also, the character of data sets has changed -- some HPC software stores data in large structured files like HDF5. To the file system, these look like big binary blobs, and a file-based back-up system needs to make a fresh copy of the whole HDF5 file whenever any change is made to it, since file-based systems can't interrogate the structure of the file. There are two ways out. Fatter pipes with more bandwidth relieve the strain by making transport faster. More sophisticated back-up systems that examine data at the block level might be better at protecting HDF5 files.

In addition, HPC systems are getting denser, packing more capabilities and also more watts into smaller and smaller footprints. It's possible that we will outgrow our machine room's ability to power and cool all the systems we want to run in the near future. When this occurs will depend to some degree on what future HPC systems look like. If all HPC nodes become liquid-cooled, then we will have a problem sooner, but if the market continues to offer a wide range of useful air-cooled systems, we can run those and still be an asset to our users. We do not currently have a liquid-cooling capability, and so cannot operate HPC systems which require this.

4.10.8 Outstanding Issues

IT security remains a source of some friction in planning data workflows. For the Additive Materials AMBench project, for instance, for which we use our Globus endpoint, the Globus data touches down on a machine in a public-facing DMZ of the PNI network, and so is not available for users to view or manipulate other than through the Globus interface. Users have a requirement to integrate this data with some database and data management systems, and do that by working with collaborators at the Johns Hopkins University, rather than internally through NIST or the CTCMS. The security concerns with respect to this data are legitimate, but the solution seems to lead to a high-friction workflow. This is principally an IT security constraint. Our Globus end-point is on the PNI, which is isolated from the main NIST network, and cannot access general internet-connected systems. This makes it hard for users to upload these data sets to internet-based databases or data management tools like the Materials Project, or even the data.gov data repository, because the data is "trapped" on PNI-attached systems, and generally only accessible via the Globus interface.

The Johns Hopkins solution is, they have their own Globus end-point which is internally attached to their data management system. This is not a bad solution, but it means we are dependent on JHU to maintain the system. What we would really like is to "liberate" the data so that users can use a wider variety of tools to manipulate and manage it. The large volume of this data makes this harder, but this is related to the earlier comment about how the PNI use-case is still evolving -- figuring out the most useful way for users to access the data will be a big step forward in solving this issue.