






ESnet

ENERGY SCIENCES NETWORK

Nuclear Physics Network Requirements Review: One-Year Update

LBLN-2001381

September 2, 2020



BERKELEY LAB



U.S. DEPARTMENT OF
ENERGY

Office of Science

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 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. 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.

Nuclear Physics Network Requirements Review: One-Year Update

Office of Nuclear Physics, U.S. Department of Energy (DOE) Office of Science
Energy Sciences Network (ESnet)
September 2, 2020

ESnet is funded by the DOE Office of Science, Office of Advanced Scientific Computing Research. Benjamin Brown is the ESnet Program Manager.

ESnet is operated by Lawrence Berkeley National Laboratory (LBNL), which is operated by the University of California for the DOE under contract DE-AC02-05CH11231.

This work was supported by the Directors of the Office of Science, Office of Advanced Scientific Computing Research, Facilities Division, and the Office of Nuclear Physics.

This is LBNL report number LBNL-2001381.

Participants and Contributors

- John Arrington, Argonne National Laboratory
- Ben Brown, Department of Energy Office of Science, Advanced Scientific Computing Research
- Steve Beher, Michigan State University, Facility for Rare Isotope Beams
- Mario Cromaz, Lawrence Berkeley National Laboratory
- Eli Dart, Lawrence Berkeley National Laboratory, Energy Sciences Network (ESnet)
- Robert Edwards, Thomas Jefferson National Accelerator Facility
- Bryan Hess, Thomas Jefferson National Accelerator Facility
- Graham Heyes, Thomas Jefferson National Accelerator Facility
- Clinton Jones, Michigan State University, Facility for Rare Isotope Beams
- Jerome Lauret, Brookhaven National Laboratory
- Sean Liddick, Michigan State University, Facility for Rare Isotope Beams
- Paul Mantica, Michigan State University, Facility for Rare Isotope Beams
- Andrew Melo, Vanderbilt University
- Ken Miller, Lawrence Berkeley National Laboratory, ESnet
- Inder Monga, Lawrence Berkeley National Laboratory, ESnet
- Brent Morris, Thomas Jefferson National Accelerator Facility
- Chris Pinkenburg, Brookhaven National Laboratory
- Jeff Porter, Lawrence Berkeley National Laboratory, National Energy Research Scientific Computing Center
- Gulshan Rai, Department of Energy Office of Science, Nuclear Physics
- Thomas Rockwell, Michigan State University, Facility for Rare Isotope Beams
- Lauren Rotman, Lawrence Berkeley National Laboratory, ESnet
- Richard Simon, Lawrence Berkeley National Laboratory
- Rune Stromsness, Lawrence Berkeley National Laboratory
- Paul Wefel, Lawrence Berkeley National Laboratory, ESnet
- Andrew Wiedlea, Lawrence Berkeley National Laboratory, ESnet
- Sean Wilkinson, Oak Ridge National Laboratory
- Linda Winkler, Argonne National Laboratory
- Jason Zurawski, Lawrence Berkeley National Laboratory, ESnet

Report Editors

- Ben Brown, Department of Energy Office of Science: Benjamin.Brown@science.doe.gov
- Eli Dart, ESnet: dart@es.net
- Ken Miller, ESnet: ken@es.net
- Gulshan Rai, Department of Energy Office of Science: Gulshan.Rai@science.doe.gov
- Lauren Rotman, ESnet: lauren@es.net
- Paul Wefel, ESnet: pwefel@es.net
- Andrew Wiedlea, ESnet: awiedlea@es.net
- Jason Zurawski, ESnet: zurawski@es.net

Participants and Contributors	4
Report Editors	4
1. Executive Summary	6
1.1 About ESnet	6
1.2 About the Office of Nuclear Physics	6
1.3 Requirements Review Purpose and Process	7
1.4 This Review	8
1.5 Key Findings Summary	9
1.6 Recommended Actions Summary	10
2. Introduction	11
3. Findings	12
4. Actions	14
5. Program Updates	15
5.1 ASCR	15
5.2 NP	15
5.3 ESnet	15
6. Case Study Updates	16
6.1 JLab	16
6.2 FRIB	17
6.3 GRETA	17
6.4 ANL	18
6.5 BNL	18
6.6 CMS HI	20
6.7 ALICE and ALICE-USA Computing	20
7. Synergistic Discussion	22
8. Appendix	23

1. Executive Summary

1.1 About ESnet

The Energy Sciences Network (ESnet) is the high-performance network user facility for the U.S. Department of Energy (DOE) Office of Science (SC) and delivers highly reliable data transport capabilities optimized for the requirements of data-intensive science. In essence, ESnet is the circulatory system that enables the DOE science mission by connecting all of its laboratories and facilities in the United States and abroad. ESnet is funded and stewarded by the Advanced Scientific Computing Research (ASCR) program and managed and operated by the Scientific Networking Division at Lawrence Berkeley National Laboratory (LBNL). ESnet is widely regarded as a global leader in the research and education networking community.

ESnet interconnects DOE National Laboratories, User Facilities, and major experiments so that scientists can use remote instruments and computing resources as well as share data with collaborators, transfer large data sets, and access distributed data repositories. ESnet is specifically built to provide a range of network services tailored to meet the unique requirements of the DOE's data-intensive science.

In short, ESnet's mission is to enable and accelerate scientific discovery by delivering unparalleled network infrastructure, capabilities, and tools. ESnet's vision is summarized by these three points:

1. Scientific progress will be completely unconstrained by the physical location of instruments, people, computational resources, or data.
2. Collaborations at every scale, in every domain, will have the information and tools they need to achieve maximum benefit from scientific facilities, global networks, and emerging network capabilities.
3. ESnet will foster the partnerships and pioneer the technologies necessary to ensure that these transformations occur.

1.2 About the Office of Nuclear Physics

The mission of the DOE Office of Nuclear Physics (NP) program is to discover, explore, and understand all forms of nuclear matter. Nuclear science began by studying the structure and properties of atomic nuclei as assemblages of protons and neutrons. At first, research focused on nuclear reactions, the nature of radioactivity, and the synthesis of new isotopes and new elements heavier than uranium. Today, the reach of nuclear science extends from the quarks and gluons that form the substructure of protons and neutrons, once viewed as elementary particles, to the most dramatic of cosmic events: supernovae.

At its heart, nuclear physics attempts to understand the composition, structure, and properties of atomic nuclei; discover new forms of nuclear matter, including that of the early universe; measure the quark structure of the proton and neutron; and study the mysterious and important neutrino. Rapid advances in large-scale integration electronics, computing, and superconducting technologies have enabled the construction of powerful accelerator, detector, and computing facilities. These provide the experimental and theoretical means to investigate nuclear systems ranging from tiny nucleons to stars and supernovae.

NP provides most of the federal support for nuclear physics research in the United States. Approximately 1,620 scientists, including 880 graduate students and postdoctoral research associates, receive support from NP. In addition, the program supports three National Scientific User Facilities. Other agencies also use these NP facilities for their own research.

NP helps the United States maintain a leading role in nuclear physics research, which has been central to the development of various technologies, including nuclear energy, nuclear medicine, space exploration, and the nuclear stockpile. The program produces highly trained scientists who help to ensure that the DOE and the United States have a sustained pipeline of highly skilled and diverse science, technology, engineering, and mathematics workers who are knowledgeable in nuclear science.

1.3 Requirements Review Purpose and Process

ESnet and ASCR use requirements reviews to discuss and analyze current and planned science use cases and anticipated data output of a particular program, user facility, or project to inform ESnet's strategic planning, including network operations, capacity upgrades, and other service investments. A review comprehensively surveys major science 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?

Requirements reviews help ensure that key stakeholders have a common understanding of the issues and the actions that ESnet may need to undertake to offer solutions. The ESnet Science Engagement Team leads the effort with collaboration from departments across the organization, including Software Engineering, Network Engineering, Infrastructure, and others. This team meets with each individual program office within the DOE SC every three years, with intermediate updates scheduled every off year. ESnet collaborates with the relevant program managers to identify the appropriate principal investigators, and their information technology partners, to participate in the review process. ESnet organizes, convenes, executes, and shares the outcomes of the review with all stakeholders.

1.4 This Review

Following the review held on May 8 and 9, 2019, between ESnet and NP, an interim review was held on September 2, 2020, to update the status of NP-supported scientific programs and facilities that were originally profiled. Preparation for this event included communication with key stakeholders (e.g., program and facility management, research groups, technology providers, and a number of external observers) to seek updates to the case studies that were prepared in 2019. Each group was asked to prepare a verbal or written briefing that updated the current, near-term, and long-term status, expectations, and processes supporting its science going forward.

This interim review included case studies from the following NP stakeholder groups:

- Thomas Jefferson National Accelerator Facility (JLab):
 - Measurement of a Lepton-Lepton Electroweak Reaction (MOLLER).
 - The Solenoidal Large Intensity Device (SoLID).
 - Theory Group and Lattice Quantum Chromodynamics.
 - The Electron-Ion Collider (EIC), a Brookhaven National Laboratory (BNL) and JLab joint project.
- The Facility for Rare Isotope Beams (FRIB).
- The Gamma-Ray Energy Tracking Array (GRETA).
- Argonne National Laboratory (ANL):
 - Gammasphere and the Argonne Tandem Linear Accelerator System (ATLAS).
 - CLAS12 and EIC.
- BNL:
 - The Scientific Data and Computing Center (SDCC), formerly known as the Relativistic Heavy Ion Collider (RHIC) and ATLAS Computing Facility (RACF).
 - The Solenoidal Tracker at RHIC (STAR).
 - The Pioneering High-Energy Nuclear Interaction eXperiment (PHENIX) and sPHENIX.
 - EIC, a BNL and JLab joint project.
- Compact Muon Solenoid (CMS) Heavy Ion (HI) research program.
- The ALICE (A Large Ion Collider Experiment) Project and ALICE-USA Computing.

The review participants spanned the following roles:

- Subject-matter experts from the NP activities listed previously.
- ESnet Site Coordinators Committee members (ESnet's liaison at each lab) whose lab hosts major NP activities. These include ANL; BNL; JLab; LBNL, including the National Energy Research Scientific Computing Center (NERSC), its supercomputing center; and Oak Ridge National Laboratory.
- DOE SC staff from ASCR and NP.
- ESnet staff supporting positions related to facility leadership, scientific engagement, and networking.

1.5 Key Findings Summary

Several key findings emerged from the review:

- Since the May 2019 ESnet NP Requirements Report, the DOE selected BNL as the site for EIC. This new NP facility will allow scientists from across the nation and around the globe to peer inside protons and atomic nuclei to reveal secrets of the strongest force in nature.
- The FRIB facility at Michigan State University was officially designated an Office of Science User Facility by the DOE.
- Several NP facilities, experiments, and collaborations continue to explore the use of high-performance computing (HPC) facilities for both simulation and analysis. These usage patterns imply a stronger emphasis on network connectivity to facilitate data sharing but also new paradigms that must be supported by the DOE ASCR facilities to allow for “streaming” workflows.
- Existing NP facilities and experiments are preparing for ESnet’s next-generation network — ESnet6 — which will include enhancing connectivity where applicable, and connecting other emerging facilities, such as FRIB.
- ESnet will continue to work with GRETA to design technology support for data intensive scientific instruments, and how it can integrate into the new FRIB User Facility.
- Open Science Grid (OSG) computing and storage approaches are critical to the success of several current NP programs and collaborations for simulation and analysis use cases. The use of the OSG for future experiments is expected to continue.
- The ESnet6 data caching project has the potential to benefit several NP use cases such as the Large Hadron Collider- (LHC) affiliated experiments (CMS HI and ALICE) as well as others such as STAR and sPHENIX. ESnet will work with the NP community to define and support use cases where applicable.
- Currently, the LHC is in “Long Shutdown 2.” Due to the COVID-19 pandemic, startup scheduling has been pushed out a full year to 2022. LHC-affiliated work continues in the form of reprocessing and simulation.
- The network requirements for EIC will be formulated in the coming years in partnership with BNL, JLab, and ESnet. Plans for the EIC computing model may go beyond local network connectivity. A “federated” computing model could be envisioned as a framework that provides shared computing, storage, and networking infrastructure to meet the future needs of the scientific community.

1.6 Recommended Actions Summary

Based on the key findings, the review identified a number of actions for NP, ASCR, and ESnet to pursue.

- ESnet will work with the EIC experiment and affiliated labs to understand the networking and data requirements to support EIC operations.
- ESnet will continue to work with NP facilities and programs to provide, understand, and enhance connectivity needs, including supporting upgrades at existing locations (BNL, JLab) in the context of ESnet6, and supporting new connections and operational models (FRIB) that are designated as Scientific User Facilities. ESnet will also support next-generation science instruments and experiments (e.g., GRETA, SoLID, the EIC detector).
- ESnet will continue to work with the LHC community, in both the NP and High Energy Physics program (HEP) contexts, to understand the implications for the extended shutdown and the expected trajectory of data movement increases in the coming years after the High Luminosity (HL-LHC) upgrades.
- ESnet, the ASCR and NP communities, and affiliated efforts like the OSG will continue to work together to understand the emerging use cases of HPC use on core computational workflows (reconstruction, simulation, analysis, and reprocessing), and impacts this will have on data movement mechanisms and application programming interfaces (APIs), as well as the role of caching within the network.

2. Introduction

A Nuclear Physics Network Requirements Review, conducted by ESnet in collaboration with the DOE SC, was held on May 8 and 9, 2019. A continuation of this process involves “checkup” procedures to be conducted after one and two years have passed. ESnet and the DOE SC, in collaboration with the previous participants, held a virtual meeting on September 2, 2020, to continue the review process. Topics of discussion included the following:

- Programmatic updates from the DOE SC ASCR.
- Programmatic updates from the DOE SC NP.
- Programmatic updates from ESnet.
- Brief updates from case study authors (See [Section 6](#) for a list).
- A general discussion and Q&A on findings and actions from the review.

This report outlines the areas of discussion, new findings and actions, and updates to the overall process of science for several key NP initiatives since 2019 and looking into 2021.

3. Findings

- The DOE selected BNL as the site for EIC, with JLab as a major partner. A “federated computing” partnership, a mechanism to share computational and storage resources, is under consideration between the EIC partners. Strong network connectivity to BNL and JLab, as well as between them, is critical to the success of EIC. Additional network connectivity to support international EIC collaboration will be established in the coming years, with ESnet playing a key role in supporting the workflows.
- JLab is prepared for ESnet6 operations and has been communicating regularly with the ESnet engineering team to ensure that rack space, electrical, cooling, and connectivity services are available. Additionally, support for a 100Gbps Ethernet has been implemented between the border of the JLab network and its core network. This will facilitate extension to science use cases after upgrades are complete. Additional work to expand the infrastructure and Science Portals Network has been completed, ensuring that high-performance data transfer can be accomplished in a safe and performant manner.
- FRIB is expected to be completed by early 2022 and has recently been formally designated an Office of Science User Facility. This milestone designation will allow ESnet to deliver direct connectivity to the FRIB site.
- GRETA successfully completed DOE Critical Decision (CD) 2/3 review in August 2020 and will enter fabrication by October 2020. The GRETA collaboration expects to test the “super-facility model” operation on the ESnet infrastructure in 2021 by performing a series of data trials between the instrument and NERSC computing, via the ESnet6 infrastructure. GRETA experimental operations will be relocated to FRIB by early 2024.
- Construction of a new SDCC at BNL is on schedule and expected to complete by mid-2021. SDCC operation will span two buildings for the first few years of operation until all resources can be centralized. Local staff (computation and networking) are working to ensure the split operation does not affect scientific operations.
- STAR, and eventually sPHENIX, continues to explore the adoption of HPC resources at NERSC, namely in the area of simulation with some emerging analysis use cases. This work is done in partnership with other NP projects at JLab (CLAS12 and the Gluonic Excitations Experiment [GlueX]) and ALICE, which are also looking to use HPC resources through the OSG software stack.
- The CMS HI and ALICE are tied to the overall schedule of the LHC. Currently, the LHC project is in “Long Shutdown 2” (since 2018), with an expected start of 2021. Due to the COVID-19 pandemic, startup scheduling has been pushed to 2022. This means no live experimentation will take place during the time and only LHC-affiliated work will continue in the form of reprocessing and simulation. CMS HI would not expect to take live data until November/December 2021, prior to the traditional seasonal shutdown. ALICE has also shifted most requirements by a full year.
- CMS HI research and development work is focusing on the adoption of the new data format (mini-Analysis Object Data [AOD]). This new format is one-tenth the size of AOD and will be in use by the next run. Use of the new mini-AOD format will not replace any

other formats, but rather it will co-exist with AOD, implying a slight increase in the storage and networking needs.

- ALICE is formalizing the concept of “analysis facilities” instead of distributing analysis workloads to every participant site. When complete, there will be approximately 10 of these analysis facilities worldwide, with the U.S. hosting one at an existing ALICE participant (e.g., LBNL). The overall network requirements for such a facility are not expected to be extraordinary (1PB per week).
- ALICE strongly depends on the peering between LBNL, NERSC, and ESnet as a critical component going forward as ESnet6 is implemented, since network performance is critical to the distributed nature of the research.
- OSG computing and storage approaches are critical to the success of several NP programs and collaborations including those found at JLab (GlueX and CLAS12), BNL (STAR and sPHENIX), CMS HI, and ALICE for simulation production, but also for limited analysis use cases that can be supported with remote data streams. Emerging use is expected in other experiments and projects, including MOLLER, SoLID, and EIC. Laboratories that support these experiments and projects are implementing technology solutions to enable these, including allocating computation resources, setting up virtual organization policy, and allowing off-site access to other sites that offer OSG resources (e.g., NERSC, ALCF, and Oak Ridge Leadership Computing Facility).
- ESnet6 caching has the potential to benefit several NP use cases such as the LHC-affiliated experiments (CMS HI and ALICE) as well as others, such as STAR and sPHENIX. This service has the potential to reduce the number and frequency of long-distance data requests, thus improving performance for distributed users. Interested parties are requested to help define requirements and use cases, to provide input to ESnet as they plan this service.
- In recent years NP experiments have increased its use of HPC resources for both simulations and experimentation. An opportunity to enhance HPC adoption is related to data access. Methods to stream data work well in the high-throughput computing (HTC) context (e.g., OSG software) but do not always work on integrated HPC systems like those at NERSC, which expect data to be pre-staged (e.g., via bulk data movement). Workarounds to stage the data “near” these resources are promising, but still experimental. Future collaboration to understand the requirements and explore possible mechanisms to support the streaming data use case on HPC systems (for example, using the Cray Slingshot interconnect that will be present on future machines to provide external data access to jobs running on compute nodes) is recommended. These efforts could potentially be combined with ESnet6 caching efforts.
- The design and implementation of EIC in the coming years is a priority for NP and the partnership laboratories BNL and JLab. ESnet is committed to working with all parties to ensure that requirements are captured to support experimental operation and should be part of the future planning discussions that relate to federated computing and use of networks to support the experiment and facilities.

4. Actions

1. ESnet will work with BNL and JLab to understand the requirements to support EIC operations. This will involve learning more about the computational, storage, and networking requirements for the two facilities that must be supported.
2. ESnet will continue to work with JLab to augment site connectivity with regional providers as ESnet6 is implemented.
3. ESnet will continue to work with FRIB on formalizing a connectivity plan and operational approaches to support scientific workflows.
4. ESnet will work with GRETA to stage data trials in 2021 to ensure proper operation when the experiment is fully deployed at FRIB.
5. ESnet will continue to work with the LHC community (in both the NP and HEP contexts) to understand the implications for the extended shutdown and the expected trajectory of data movement increases in the coming years after the HL upgrades.
6. ESnet will work with ALICE at LBNL to support the emerging “analysis facility” use case as needed.
7. ESnet, the ASCR and NP communities, and affiliated efforts such as the OSG see opportunity to collaborate on synergistic activities related to the evolution of scientific workflows. These will involve reconstruction, simulation, analysis, and reprocessing efforts that have historically leveraged HTC but are exploring the incorporation of HPC facilities. Participation in existing, or the creation of new, working groups to explore these ideas is recommended. These include, but are not limited to:
 - a. Mechanisms that can be used to efficiently stage data at HPC facilities through the use of bulk data movement and storage, or through on-demand streaming mechanisms, and exploring common APIs that would facilitate streamlined software development and deployment.
 - b. The roles for “caching” of scientific data within the network and opportunities to ensure that scientific use cases are understood and supported by tooling being developed by ESnet for use in ESnet6 and the ASCR facilities as they adopt new computing architectures.

5. Program Updates

5.1 ASCR

The ESnet Requirements Review program provides strategic network insights required to support DOE SC activities. The review in 2019 with NP was the first of a revitalized effort in this space, and this checkup activity will help to precisely map any changes that have occurred in the past year.

In addition to being able to accurately map resources and barriers that may affect NP, insights can be drawn that will help to guide or inform a wider spectrum of ASCR activities, including NP's usage of the ASCR Leadership Computing Facilities (LCFs) and NERSC, networking via ESnet, and ASCR research investments. These cross-boundary insights are critical to planning the future trajectory of scientific discovery and creating a data-centric view of the landscape that will foreshadow future capabilities.

5.2 NP

Significant progress has occurred since the review in 2019. Several projects continue to pass critical design review milestones while others are progressing through scientific peer review. The partnerships that are forming around many of these experiments will rely heavily on ESnet resources to be successful. Updates include (but are not limited to) the following:

- The MOLLER experiment is progressing to CD1.
- The SoLID experiment is advancing towards an NP science review in early spring 2021.
- DOE has selected BNL as the construction site for EIC.
- FRIB at Michigan State University is working with ESnet on steps to become a directly connected Scientific User Facility.

The projects that are 5–10 years in the future, including EIC, are in the process of defining schedules and identifying resources, including institutions that may contribute computational cycles or storage space.

5.3 ESnet

The ESnet6 upgrade project began construction in 2020, and as of January 2021 the buildout of the ESnet6 optical infrastructure is nearly complete. Calendar year 2021 will continue to be a period of intense project activity for the ESnet staff. Some ESnet-connected sites have been upgraded, or will be soon, with more capacity to support scientific use cases (like those profiled in the 2019 ESnet NP Requirements Review).

6. Case Study Updates

Each of the case studies was afforded an opportunity to give an update on status. The updates for each appear in the following sections.

6.1 JLab

JLab profiled five use cases in the 2019 review:

- Facility Overview
- MOLLER
- SoLID
- EIC
- Lattice QCD

The following updates were provided by JLab to these use cases:

- There are no changes to the process of science, data expectations, or operational expectations for many of the existing experiments. This includes GlueX and the Continuous Electron Beam Accelerator Facility (CEBAF) data transfer node (DTN) Large Acceptance Spectrometer at 12 GeV (CLAS12).
- MOLLER and SoLID also do not have any design changes. MOLLER is progressing towards project CD1, and SoLID will undergo a science review.
- JLab and BNL are partners on EIC computing. Strong network connectivity between the sites will become important in the coming years.
- OSG computing is now used for the majority of GlueX and CLAS12 Monte Carlo simulation efforts. This implies a significant increase to inbound network traffic to support these efforts.
- JLab will be bringing up more OSG-affiliated computational infrastructure to its data center. Most of the computational resources are affiliated with the policy and operation of locally used virtual infrastructure, but others could affect network traffic into and out of JLab. This includes services that could attract jobs seeking OSG resources.
- JLab is prepared for ESnet6 and has been communicating regularly with the ESnet engineering team to ensure rack space, electrical, cooling, and connectivity services are in place. Additionally, support for a 100Gbps Ethernet has been implemented between the border of the network and the core. This will facilitate extension to science use cases after upgrades are complete.
- The JLab Science DMZ infrastructure and Science Portals Network are expanding. Improvements include additional DTN capabilities to support science activities, as well as portal capabilities for data sharing and analysis. Each of these enclaves features a specific security and usage profile to match user needs.
- JLab has ongoing efforts to support remote work during the COVID-19 pandemic. These include increasing capabilities such as virtual private network access, use of cloud

resources (Office 365, Amazon Web Services, Salesforce), and other business functions.

6.2 FRIB

As FRIB approaches the formal conclusion of its construction phase, ESnet's work with FRIB and NP in support of early science and the transition to full operations is ramping up.

The following updates were provided by FRIB:

- FRIB is expected to reach CD-4 by mid-2022, but early completion by early 2022 is still possible.
- FRIB has recently been designated an Office of Science User Facility; this designation will allow ESnet to deliver direct connectivity to support FRIB's instruments and experiments.
- Michigan State University will continue to support ongoing experiments (such as the National Science Foundation-funded National Superconducting Cyclotron Laboratory) through FY21 before transitioning to FRIB operations.
- GRETA remains the largest FRIB scientific driver for the immediate term, as it will be one of the first instruments that operates in the new facility and will be the largest driver for data production.
- Other FRIB experiments are in various states of planning. The High Rigidity Spectrometer is approaching CD1 and has a planned start date of 2027.
- Discussions with ESnet continue about the policy and technical implications of connectivity for FRIB. Establishing the technical aspects (network path, equipment, funding, etc.) is ongoing.
- FRIB has no other significant changes to the case study process of science, data projections, or operational aspects to report.

6.3 GRETA

GRETA is still in the design phase and has not yet started taking data. Therefore, most activities are related to benchmarking, designing, and evaluating performance once it is built. As previously discussed, GRETA will be located at FRIB, but can also operate at ANL. The instrument design facilitates a streaming data model when deployed and has functionality on and off network connectivity.

The following updates were provided by GRETA:

- The design phase for GRETA is complete, and the project successfully completed CD2/3 review in August 2020.

- GRETA fabrication will begin October 2020, and the design for the computing system is expected to be complete by December 2021. All dates have some uncertainty due to the COVID-19 pandemic.
- GRETA expects to test “super-facility” operation on ESnet infrastructure in 2021. This will be a data trial that couples instrument operation with HPC resources via a high-speed network.
- GRETA experimental operations at FRIB are planned for early 2024.
- GRETA prototyping and benchmark testing is trying to expose stress on aggregation parts of the system. Thus far, all designs have tested beyond specified levels.
- GRETA has no other significant changes to the case study process of science, data projections, or operational aspects to report.

6.4 ANL

ANL was asked to complete case studies for the Gammasphere and ATLAS instruments, as well as the user perspective for JLab’s CLAS12 and eventual use of EIC at BNL. ANL reported no updates or other significant changes to the case study process of science, data projections, or operational aspects.

6.5 BNL

BNL was tasked with three case study topics:

- The SDCC, formerly known as the RACF.
- STAR.
- PHENIX and sPHENIX.

The following updates were provided by BNL on RHIC and SDCC:

- Construction of a new SDCC on the Brookhaven site is on schedule, with only minor delays due to the COVID-19 pandemic. The project entered safe mode in April 2020 to limit staff and resumed a close to normal schedule in June of 2020. Work is expected to be complete in mid-2021, with only minor delays anticipated to the overall schedule.
- SDCC operation will span two buildings at BNL for a couple of years until all resources can be centralized. Local staff (computation and networking) are working to ensure the local distance does not affect scientific operations.

The following updates were provided by BNL on STAR:

- The EIC site selection announcement will affect SDCC (formerly RACF) in the future. It is anticipated that, near the end of life, some BNL resources (computational, storage, and potentially detectors) may be transitioned to EIC.

- STAR operations were not severely affected by the COVID-19 pandemic: RHIC was operating and not in a shutdown state.
- STAR is projecting modest changes to some of the data projections, as much as 10% beyond the original report.
- STAR will use additional HPC resources (provided by the Cori system at NERSC) and HTC resources (some provided by the LBNL Information Technology Division) in an effort to scale computational needs.
- STAR expects to upgrade data transfer hardware resources after the next run period. This is not expected to be disruptive to the workflow, and may offer performance improvement gains.
- BNL experiments are supportive of NERSC (and ESnet) participation in OSG.
- ESnet6 is experimenting with different services to support scientific projects, including data caching that can be integrated with experimental workflows. This work is still in the planning stage, and ESnet welcomes input from STAR and sPHENIX based on current use cases, as well as potential use from EIC.

The following updates were provided by BNL on PHENIX and sPHENIX:

- PHENIX is expected to complete its experimental program after the last reprocessing of a PB-scale data set. Sharing data for analysis may continue for another four to five years, but no live experimentation will occur.
- sPHENIX is expected to start taking data in 2023. At this time, no COVID-19 pandemic delays are expected. Assembly of the detector will begin in 2021.
- sPHENIX expects to use the SDCC for storage and processing, and will be working to integrate OSG-supported approaches to facilitate HTC operation that can use both local and remotely distributed resources. Investigation to integrate HPC workflows via NERSC (mainly for simulations) has started.
- sPHENIX is planning to run for three years and possibly longer depending on the EIC schedule.
- sPHENIX is exploring working with the Advanced Computing Center for Research and Education at Vanderbilt University for possible sharing of computational resources.
- sPHENIX has no other significant changes to the case study process of science, data projections, or operational aspects to report.

The following was provided by BNL on EIC:

- EIC will be the next-generation NP accelerator and detector experiment when complete. Site selection at BNL was recently announced by the DOE.
- Computation for EIC is co-managed by JLab and BNL.
- Many details of EIC are still being developed, including design, schedules, and overall strategy.
- The use of the SDCC to support EIC (post STAR/sPHENIX) is still to be determined.
- A “federated computing” model between BNL and JLab, which will share resources at both sites with an interconnection made possible via ESnet, is being considered. In

particular, the community is exploring a model of “submit anywhere, run anywhere” for computational use. OSG software and operations are expected, along with use of HPC at LCFs.

- BNL and JLab are considering the formation of a “core” of a new east-coast “super facility” as EIC computational plans come into focus. Establishing significant ESnet network capability between, and paths to, these two facilities should be discussed.
- EIC may leverage computing resources from abroad, including CNAF-INFN (the Center for the Research and Development in Information and Communication Technologies at the Italian Institute for Nuclear Physics) in Europe.
- EIC will be forming working groups to better understand scientific operations and collaboration. ESnet’s early participation in working groups can bring valuable expertise, knowledge, and innovation in the process of defining data and computational framework and networking models for EIC.

6.6 CMS HI

CMS HI is tied to the overall schedule of the LHC. Currently the LHC project is in “long shutdown” (since 2018), with an expected start of 2021. Due to the COVID-19 pandemic, startup has been pushed to 2022. This means no live experimentation, but LHC-affiliated work continues in the form of reprocessing and simulation.

The following updates were provided by CMS HI:

- CMS HI would not expect to take live data until November/December 2022 (prior to the seasonal shutdown).
- Research and development work is focusing on the adoption of the new data format (mini-AOD). This new format is one-tenth the size of AOD and will be in use by the next run.
- Use of the new mini-AOD format will not replace any other formats: it will co-exist with AOD (and thus be stored and exchanged, forcing a slight increase in storage space/network needs). Overall, it is expected that mini-AOD will reduce streaming network requirements over time as it begins to be adopted for distributed analysis.
- CMS HI has no other significant changes to the case study process of science, data projections, or operational aspects to report.

6.7 ALICE and ALICE-USA Computing

ALICE and ALICE-USA Computing are also tied to the overall schedule of the LHC. Currently, the LHC project is in “long shutdown” (since 2018), with an expected start of 2021. Due to the COVID-19 pandemic, startup has been pushed out to 2022. This means no live experimentation, but LHC-affiliated work continues in the form of reprocessing and simulation.

The following updates were provided by ALICE:

- ALICE has shifted most requirements by a full year due to the COVID-19 pandemic.
- ALICE is still working to implement more perfSONAR testing within the experiment (and through OSG), but this is not complete.
- ALICE is still formalizing the concept of “analysis facilities” instead of doing analysis at every participant site. If this comes to fruition, there would be ~10 of these facilities worldwide, with the U.S. getting one (possibly at LBNL). The overall network requirements for such a facility are not expected to be extraordinary (e.g., 1PB a week).
- ALICE continues to explore using NERSC HPC resources and integrating these into the Worldwide LHC Computing Grid. Dynamic access to data is one area of improvement. Currently, this is accomplished by also using the LBNL Information Technology Division cluster to stage data. ALICE is working jointly with STAR on this aspect, as the work will benefit both use cases.
- ALICE rates the peering between LBNL, NERSC, and ESnet as a critical component going forward as ESnet6 is implemented.
- ALICE has no other significant changes to the case study process of science, data projections, or operational aspects to report.

7. Synergistic Discussion

The following additional topics were discussed by the Requirements Review participants during the virtual event:

- In recent years NP experiments have increased its use of HPC resources for both simulations and experimentation . An area of improvement to this adoption is related to data access. Methods to stream data work well in an HTC context (e.g., OSG software), but do not always work on integrated HPC systems like those at NERSC, which expect data to be pre-staged (e.g., via bulk data movement). Workarounds to stage the data “near” these resources are promising, but still experimental. Future collaboration to understand the requirements and explore possible mechanisms to support the streaming data use case on HPC systems (for example, using the Cray Slingshot interconnect that will be present on future machines to provide external data access to jobs running on compute nodes) is recommended. These efforts could potentially be combined with ESnet6 caching efforts.
- The design and implementation of EIC in the coming years is a priority for NP, BNL, and JLab. ESnet is committed to working with all parties to ensure that requirements are captured to support experimental operation. ESnet requests to be a part of future planning discussions that relate to federated computing and use of networks.
- ESnet6 caching has the potential to benefit several NP use cases. Interested parties are requested to help define requirements and use cases, to provide input to ESnet as it plans this service.

8. Appendix

List of Abbreviations

ALICE	A Large Ion Collider Experiment
ANL	Argonne National Laboratory
AOD	analysis object data
API	application programming interface
ASCR	Advanced Scientific Computing Research
ATLAS	Argonne Tandem Linear Accelerator System
BNL	Brookhaven National Laboratory
CEBAF	Continuous Electron Beam Accelerator Facility
CMS	Compact Muon Solenoid
CMS HI	Compact Muon Solenoid Heavy Ion research program
DOE	U.S. Department of Energy
DTN	data transfer node
EIC	Electron-Ion Collider
ESnet	Energy Sciences Network
FRIB	Facility for Rare Isotope Beams
GlueX	Gluonic Excitations Experiment
GRETA	Gamma-Ray Energy Tracking Array
HL	high luminosity
HPC	high-performance computing
HTC	high-throughput computing
JLab	Thomas Jefferson National Accelerator Facility
LBNL	Lawrence Berkeley National Laboratory
LCF	Leadership-class Computing Facilities
LHC	Large Hadron Collider
MOLLER	Measurement of a Lepton-Lepton Electroweak Reaction
NERSC	National Energy Research Scientific Computing Center
NP	DOE Office of Nuclear Physics
OSG	Open Science Grid
PB	petabyte
PHENIX	Pioneering High-Energy Nuclear Interaction eXperiment
RACF	RHIC and ATLAS Computing Facility
RHIC	Relativistic Heavy Ion Collider
SC	DOE Office of Science
SDCC	Scientific Data and Computing Center
SoLID	Solenoidal Large Intensity Device
STAR	Solenoidal Tracker at RHIC

