Future Generation High Performance Computing Center (FG-HPCC)

Attachment (1) RFI Technical Considerations

Table of Contents

Section 1. Introduction

This RFI Technical Considerations document provides interested respondents with information about LLNS' FG-HPCC vision and background details.

Section 2. Summary

Section 2.1. LLNS

Lawrence Livermore National Security, LLC (LLNS) is interested in receiving information about technologies that could be available in the 2029-2030 timeframe that may serve to enable the vision for a Future Generation High Performance Computing (HPC) Center (FG-HPCC) described in this document. The future HPC Center vision has been conceived to meet the future mission needs of the Advanced Simulation and Computing (ASC) Program within the National Nuclear Security Administration (NNSA). LLNS envisions a center composed not of many independent clusters, but of heterogeneous elements accessible to users as a single system. The capabilities will be integrated to create a scalable, flexible, yet tightly coupled computing center capable of integrated HPC, AI, and cloud-like workloads.

Section 2.2. NNSA

NNSA, a semi-autonomous agency within the Department of Energy (DOE), is responsible for the management, security, and modernization of the nation's nuclear weapons, nuclear nonproliferation, and naval reactor programs. The NNSA Stockpile Stewardship Program, which underpins confidence in the U.S. nuclear deterrent, has been successful since its inception in 1995, largely as a result of HPC-based modeling and simulation (ModSim) tools. HPC tools have increasing roles in understanding evolving nuclear threats posed by adversaries, both state and non-state, and in developing national policies to mitigate these threats.

The NNSA's Advanced Simulation and Computing (ASC) Program provides the computational resources that are essential to enable nuclear weapon scientists to fulfill stockpile stewardship and modernization requirements through simulation without underground testing. Modern simulations on powerful computing systems are key to ensuring no return to testing, development and deployment of cost-effective and high-quality solutions, and that the stockpile can address an evolving threat landscape.

Section 2.3. Evolving Mission Needs

The stockpile continues to move further from the nuclear test base, through aging of stockpile components and modifications involving system refurbishment, reuse, or replacement. The realism and accuracy of ASC simulations must continue to increase over time to track the aging

stockpile through development and use of improved physics models and solution methods, which require orders of magnitude greater computational resources than are currently available. In the coming decade, weapon modernization efforts are expected to become a much larger fraction of the NNSA workload, and simulation teams at design agency (DA) sites not only must achieve higher fidelity but also must closely collaborate with production agency (PA) sites to understand their processes, capabilities, and manufacturing constraints.

NNSA simulations will increasingly use data and models from across the NNSA complex to ensure that designs are optimized for real-world production facilities. HPC use cases will no longer be confined to studies or workflows conducted at a single site at a time; they will span a web of connected sites across the complex, and sites will iterate on design and production processes. As the need for HPC expands across the complex, multi-physics codes will need to integrate an increasing number of capabilities and make these capabilities available to more users across the complex. Simulation results will need to reflect the latest data from PA sites, maintaining consistency with manufacturing processes and also guiding them.

Section 2.4. Revolutionizing Productivity to Accelerate the Mission

Traditionally, HPC procurements in and outside of NNSA have sought to accelerate delivery on mission priorities by accelerating modeling and simulation jobs. High performance accelerators, particularly GPUs, have become an essential tool for NNSA codes, enabling floating-point intensive simulations to complete in hours when previously they required weeks. GPUs will continue to be relevant for physical simulations and for AI, well into the foreseeable future. However, NNSA's new mission also needs to accelerate complex workflows and cross-complex collaboration. Simulation runtimes are not the main bottlenecks for these problems, at least not yet. Instead, human factors dominate the runtime for large workflows. To accelerate this work, the NNSA tri-labs must pursue other directions that increase the productivity of developers and designers, while improving upon the modeling and simulation performance that we have already achieved. The FG-HPCC vision includes three ways to enhance productivity: artificial intelligence (AI), increased automation, and persistent services.

Section 2.4.1. Artificial Intelligence (AI)

AI is emerging as a tool to guide and to accelerate large simulation workflows, to understand and to model complex scientific phenomena, and to automate critical non-simulation tasks (e.g., image or document analysis) in mission workflows. Fundamentally, AI will increase the ability to synthesize knowledge from our data. Increasingly, NNSA workloads will be coupled with AI, leading to much more complex data-centric workloads in FG-HPCCs.

Future workloads will integrate large-scale simulations and simulation ensembles alongside AI training, AI inference, complex data lakes, live data streams, services, storage, and extensive

automation. For the vast majority of AI use cases, simulations will generate the training data. In many NNSA-relevant domains, no other means exists to obtain enough data to use modern AI techniques. As such, simulation ensembles will be tightly coupled with AI training, and data generated by simulations must be available to train models rapidly and iteratively.

Section 2.4.2. Increased Automation

AI, continuous deployment of codes and Machine Learning (ML) models, and workflow scheduling all depend on the ability of users and staff at FG-HPCC sites to leverage automation. Users must automatically and securely set up and run large scientific workflows. Code teams and facility developers must automatically build, test, and deploy code. Facility staff and users must automatically deploy infrastructure and services. The scientific and AI library landscape is constantly changing, and users must also be able to test rapidly with the latest versions of internal and external scientific and AI packages. Moreover, developer workflows will need to be augmented with ML-ops, so that developers can version and manage ML models as easily as they currently manage software packages. Regular updates and rapid development will be critical to exploit the full power of next-generation NNSA systems without sacrificing correctness or end-user productivity.

Section 2.4.3. Persistent Services

To enhance productivity in a connected NNSA complex, NNSA HPC applications will be made available as remotely accessible services. An expectation is that production sites will leverage DA-developed simulations, models, and databases. Conversely, DA-hosted simulations will leverage models and data from production facilities, enabling more accurate simulation scenarios through rapid iteration.

Designers and analysts will use hosted data sets, either as inputs for codes, as training data for AI, or simply for large-scale data analysis. Rather than hosting these in traditional filesystems and moving them from site to site when needed, data sets may be hosted on demand in S3-like object storage, or they may be searchable via continuously updated indexes and databases. The FG-HPCC must enable NNSA developers to make codes, data, and AI models available to other sites through hosted web portals and other services.

Section 2.5. Future Infrastructure Vision

With over 3 double precision exaflops in peak compute capability, an abundance of utilities available to our computer room floors, the combined NNSA data centers are world-class. While NNSA has a long history of delivering highly capable HPC resources that meet program needs, NNSA data centers must evolve to meet the changing landscape of the NNSA mission.

Section 2.5.1. Gaps

The future generation HPC center must solve these three issues:

1. Most HPC procurements and deployments, including those of LLNS, currently focus on entire integrated systems, and centers operate systems largely independently. Incremental upgrades of center resources are difficult and often entail changes to unrelated resources to accommodate a desired improved capability.
2. The security model of current HPC centers does not allow for strong isolation or multi-tenancy. Ensuring the security of HPC center resources and jobs running on them often requires choices, like rigidly defined network zones, that limit the specific workload for which a given resource can be used.
3. Users must explicitly specify the resources on which their jobs run, using per-machine batch schedulers. Decomposed workflows with multiple interacting components are, at best, poorly supported. Execution of workflows that require disparate or disaggregated resources cannot currently be coordinated automatically through any common system or scheduling layer. Ultimately, overall center efficiency suffers because the center cannot be optimized as a unified whole.

Section 2.5.2. Anticipated Solution

The Future Generation HPC Center (FG-HPCC) will facilitate greater efficiency in using center resources and, most importantly, greater productivity for users. To support NNSA's anticipated mission needs, LLNS will transform LC into a converged data center for HPC-style workloads, workflows, and persistent services. HPC, AI training, AI inference, web services, analytics, and continuous integration (CI) jobs will all be well supported. Users will be able to employ heterogeneous compute and storage resources seamlessly to support complex workloads and orchestrated workflows. The FG-HPCC will enable composable but tightly integrated technologies to serve multiple purposes.

Effectively, the center will become the system and the system should not be optimized only for any one workload (e.g., ModSim or AI) although large fractions of it will be. For example, the center may have a GPU partition that works well for both AI and HPC, and another CPU-only partition that is optimized for service workloads. FG-HPCC operations will support incremental updates that target specific improvements in capabilities and homogeneous, tightly integrated resources will be procured when they provide advantages over incremental updates. Diverse systems resources will only be co-scheduled and orchestrated using common interfaces. Ultimately, users will not need to care what resources are used to execute portions of their workflows, only that their workflows are executed quickly and efficiently.

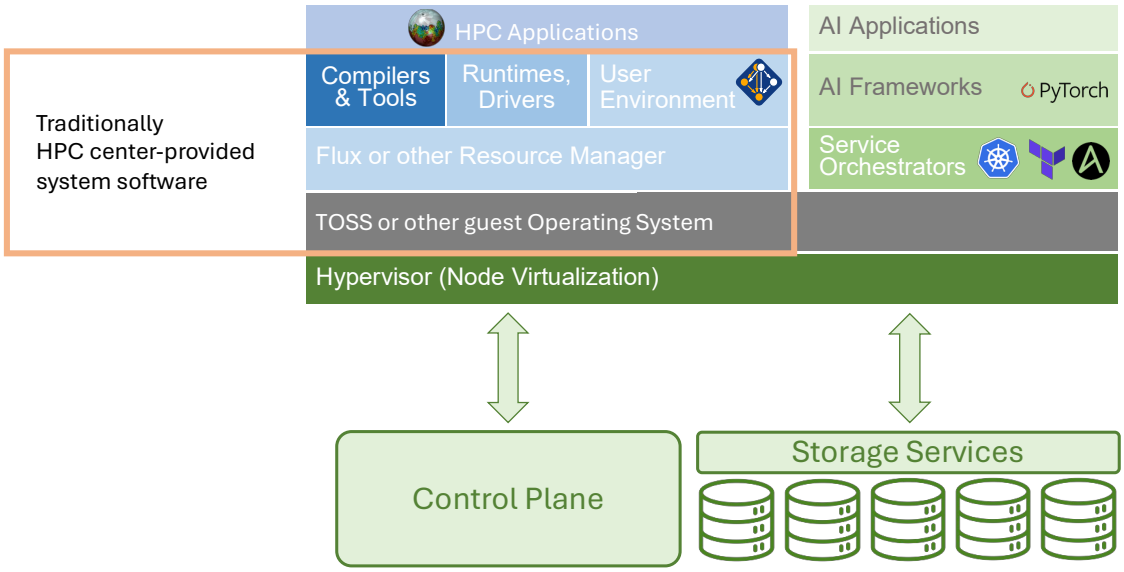

Figure 1: Open source system software stack diagram

To support efficient execution of the wide range of workloads that meet NNSA's mission needs, the FG-HPCC must support automated infrastructure provisioning so that diverse compute, storage, and networking resources can be assembled and tailored to particular jobs, applications, and workflows. Some workload scenarios may have jobs that require the computing power of an entire platform, taking days or weeks to complete, but the workload must also simultaneously store, consume, analyze, transform, and train on data that it produces. Other scenarios may require millions of small jobs that use the same computational resources. Users will provision and manipulate the system automatically through well-defined APIs. The FG-HPCC will support rapid development of new codes, services, *and* workflow tools, and web interfaces must support the availability of secure persistent services to users across the NNSA complex.

The required FG-HPCC transformation is a convergence with technologies used in multi-tenant hyperscale data centers. In addition to a heterogeneous, center-as-a-system deployment model, LLNS will enhance the center with cloud-like capabilities. This transformation requires two major innovations in center management. First, LLNS will integrate and operate the data center, using a common, open-source software stack to provision hardware resources flexibly and securely. Second, by integrating and operating the data center, LLNS system procurements will not require a single vendor to provide all system elements, and will expand its procurement model to include HPC/cloud hardware from a wider range of vendors.

Section 2.5.3. Open Source System Software Stack

LLNS has a history under the ASC program of developing and integrating the system software used in LC, and it will leverage this strength to evolve into an FG-HPCC. ASC funded the

development of many tools that have become widely used for HPC system management. SLURM, originally developed under the ASC program, is now the de-facto standard for HPC batch scheduling and is used at many HPC sites. Flux, the next-generation resource manager used on El Capitan has enabled LLNL scientists to run heterogeneous, massively parallel workflows. The Trilab Operating System Software (TOSS), the NNSA tri-lab's operating system stack derived from Red Hat Enterprise Linux (RHEL), is used at the tri-labs, NASA, and other NNSA sites. ZFS on Linux is an ASC project and now powers Lustre file systems across the tri-labs. Spack, started under ASC in 2023, is the de-facto standard package manager for user-level HPC software. Additional ASC-funded tools and environments, such as OpenCHAMI, and others are emerging to modernize and simplify system management.

To enable the FG-HPCC, LLNL and the NNSA tri-labs must work with each other, other data centers and industry on a consolidated software environment for modern resource provisioning at the data center level and, potentially, at the NNSA-complex-wide level. Figure 1 shows the envisioned stack. Alongside the traditional HPC stack are frameworks for services and AI, and below both are primitives for strong isolation, control-plane APIs, and storage APIs. A primary aim of this RFI is to assess potential designs and components of such a stack and to prioritize collaborations that can build and harden pieces of it. Aspirational requirements for the FG-HPCC software and provisioning environment are detailed in Section 3. Section 4 and Section 6 outline the collaborations LLNS envisions, along with supporting information for RFI respondents' consideration.

Section 2.5.4. Flexible FG-HPCC Procurement

Even though LLNS develops and manages much of its own software, LLNS cannot implement this FG-HPCC vision alone; it must procure *elements* of a hybrid HPC/cloud data center from hardware vendors. However, these elements must work together as an integrated whole, *even if they come from separate providers*. The transformation will necessitate a new, more flexible procurement strategy. Past large NNSA acquisitions have been awarded to a single offeror to deliver a single, integrated computer with associated infrastructure. LLNS's contemplated new acquisition approach will take procurements in a different direction, supporting potentially multiple awards that target specific resource types to evolve the LC hybrid HPC/cloud center to meet NNSA ongoing needs.

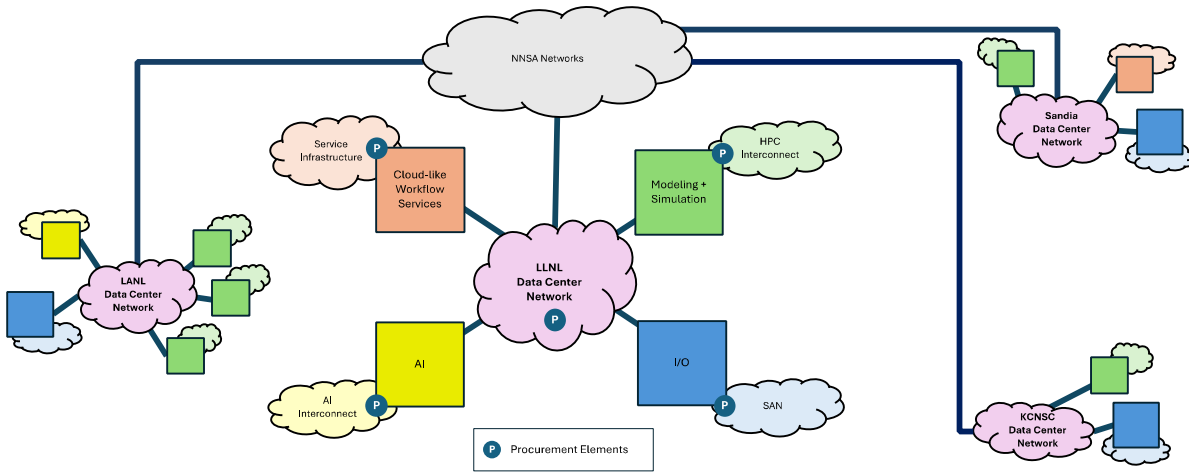

Figure 2: LLNL's vision of FG-HPCC procurement elements across the NNSA complex.

Figure 2 is a conceptual diagram of the evolved HPC data center and its relationship to future system procurements at LLNL and other NNSA sites. A large procurement might include 5 major elements provided by one or more entities, labeled with gray circled "P" symbols in the figure:

i. A next-generation, integrated ModSim capability (green);
ii. An element to train AI models efficiently and to use those models for inference in conjunction with ModSim activities (yellow);
iii. Compute resources highly optimized to support persistent, data-intensive services (orange);
iv. Center-wide storage resources and/or services (blue); and
v. A high-speed data-center-wide network to support composition of resources to execute complex workflows, and jobs within those workflows, efficiently (pink).

LC will manage the software stack and will integrate elements into its center on an ongoing basis, and acceptance will focus on the functionality and performance of APIs and interfaces specific to each element in addition to the functionality and performance of the element. LLNS may issue RFPs that include opportunities to bid on many elements, or it may issue RFPs that target only certain elements, depending on mission need at any time. LLNS anticipates making this transition gradually over the next 4-5 years and is particularly interested in technology developments that are deliverable between now and 2030[1].

This model is not limited to LLNL—other NNSA laboratories can take part and can similarly tailor their deployments to meet their own mission needs and those of the NNSA complex. Figure 2 shows FG-HPCC deployments at Sandia, Los Alamos (LANL), and the Kansas City National Security Complex (KCNSC), each with a different mix of resources tailored to the host site. Like the LLNL FG-HPCC, other sites can compose HPC center components and achieve strong isolation across elements through a central data center network. Ultimately, the envisioned

[1] This FG-HPCC RFI (the subject of this document) will inform LLNS's vision described herein and will potentially help LLNS to better establish the basis for future hardware acquisitions.

approach will enable more coordinated deployment of computational resources across the NNSA complex, allowing users at one site to coordinate resources in a unified manner across the entire complex.

Section 3. Cross-Cutting Requirements

To address the gaps described in Section 2.5.1, the FG-HPCC transformation will require a number of technological innovations. This section outlines those that LLNS views as most critical.

Section 3.1. Security

Security is paramount in this transformation. HPC centers traditionally implement security using OS-level controls, trusted filesystems, and disparate network zones, but this solution is too coarse-grained for future NNSA needs. Future data center elements must incorporate robust security measures at every level, from hardware to software, ensuring data integrity, confidentiality, and availability. This incorporation will include (but is not limited to):

- **Multi-Tenancy**: Efficient use of center resources requires that they can be used flexibly and securely by multiple users at different security levels. Users should be able to compose arbitrary subsets of center resources to run jobs, workflows, and services as needed for different workloads, and no portion of a system should be dedicated to any one security level. While LLNS still expects to use an air gap to separate classified from unclassified workloads, on either side of the air gap, the FG-HPCC should be able to use strong *logical* rather than physical separation to ensure secure separation between jobs across the center and within a node.
- **Strong on- and off-node Isolation**: Serving a diverse set of users and missions requires that users be isolated from each other—that one user cannot observe the data or actions of another user without permission. This model must support fine-grained partitioning of system components (e.g., CPUs , GPUs, and elements such as networks and storage) while maintaining this heightened level of security for each partition.
- **Flexible access control**: Modern workflows across diverse teams require that users have greater flexibility and control of access to their data while still ensuring that the data is secure. Advanced authorization and authentication mechanisms must protect data, must allow the owner of the data to determine who has access and must ensure that only those authorized users can access it. In addition to traditional user-based authentication, role-based authentication is also a critical capability to enable data-centric workflows. Filesystems and other system-level data services should support authorization mechanisms so that a compromise of any node does not result in a breach of an entire filesystem.

In short, *resources in the FG-HPCC must be easily usable for any workload without compromising security*. The type of job should not matter—a given piece of hardware should be allocatable to a batch job, on-demand and orchestrated jobs, a persistent service or an untrusted CI job without changing the architecture or security model of the center. The size of a job should also not matter, and users should be able to securely share a node using different *fractions* of CPUs, GPUs, or other node components. The *user* running a job should not matter—LLNS should have confidence that users on the same node cannot access each other's data without authorization. Storage resources should be usable for large-scale parallel filesystems (e.g., Lustre), for small block volume claims, or for object storage. Multiple tenants, potentially from different programs or laboratories, should be able to run separate jobs on the same network or within a single node. Users should be able to build secure, persistent services without assistance from facility administrators.

NNSA developers work across a wide range of environments, from completely open GitHub repositories to internal unclassified machines to air-gapped systems. NNSA software necessarily leverages many open-source libraries in addition to internally developed libraries and tools, and often developers need to run tests with on-prem hardware in response to new releases or changes in external software. Such integration testing ensures that on-prem applications continue to function reliably. LLNS Security restrictions currently prevent running untrusted tests on-prem, but the expectation is that the FG-HPCC will provide sufficient isolation (e.g., network and compute virtualization or containerization) to allow it.

Public cloud systems provide this degree of isolation for their users, and overheads of isolation technologies like virtualization, software defined networking, trusted execution environments, and encryption have become low enough to be considered in HPC environments. Figure 1 is a notional diagram of the envisioned stack, with a hypervisor at the lowest level, below a guest operating system. That traditional HPC environments lack the guarantees provided by these technologies is a major barrier to unifying on-prem infrastructure with technologies used by hyperscale data centers. LLNS welcomes research directions that lower the performance overheads and costs of using these technologies because they will be critical for future mission workloads.

Section 3.2. Interoperability

The FG-HPCC will require interoperability between different instance of elements, whether they are of different types or the same type. Resources from each element must be usable together with strong isolation that spans the elements involved. For example, a user might request ModSim resources, AI resources, and I/O resources in order to run an AI-augmented ModSim workflow. The resource manager (e.g., Flux or Kubernetes) must rapidly set up an isolated network that spans the ModSim, AI, and I/O elements, connected through the Data

Center Network Core. While the FG-HPCC would ideally use the same isolation technology for all elements, a viable, all-encompassing solution may not be available in the 2030 timeframe. Therefore, depending on the networking technology used by the different elements, center-wide interoperability may require bridging different types of network isolation mechanisms.

Once an isolated network is constructed, the resource manager must also set up virtual machines on the nodes (if they are not bare metal nodes), and ensure that they are associated with the newly created network. The node-local hypervisor must be able to associate networks with jobs running on subsets of CPU cores and with subsets of GPU resources.

LLNS seeks information describing methods for building virtual networks that span different network technologies, information on the interoperability of isolation mechanisms (VLANs, VXLAN/VNIs, encryption-based mechanisms, pkeys, and others) across HPC and ethernet networking technologies. LLNS also seeks information on low-overhead on-node virtualization technologies that could be available by 2030, and how these technologies inter-operate with isolated networks. Open source, low overhead, performant, and cost-effective mechanisms with strong security guarantees are most preferred.

Section 3.3. Data Center Control Plane

The FG-HPCC will provide the baseline services required to run an HPC data center like an on-premises cloud. The on-premises cloud does not need to include *all* services that hyperscalers offer, as the support burden and required investment would be too high. However, the FG-HPCC will provide *at least* baseline storage, compute, and network allocation services from which higher-level services and/or Platform-as-a-Service (PaaS) systems can be constructed, either by users or facility staff. Ideally facility staff will maintain resources for large, center-wide filesystem services using the same underlying interfaces that a user would leverage to construct their own smaller, job-specific filesystem. These capabilities will enable FG-HPCC users to construct their own persistent services, frameworks, ensemble runs, and workflows rapidly. FG-HPCC users will be able to find TerraForm scripts, Helm charts, and other Infrastructure-as-Code (IaC) solutions online and to adapt them quickly to provision their own services and HPC workflows securely within the FG-HPCC.

To manage and to control resources in the FG-HPCC, LLNS will use a common open-source software stack, common system administration practices, and industry-standard APIs where available. No widely used open standard for an on-premises control plane like those used by cloud providers currently exists; respondents to this RFI should describe any projects that have the goal of providing open, on-premises cloud-like control interfaces, particularly that support integration of heterogeneous hardware and software components from multiple vendors. Most importantly, the control plane must enableusers to compose FG-HPCC elements into

private, isolated enclaves within the FG-HPCC. It must also seamlessly allow elements to be added, upgraded or removed from the center over time, without disturbing user workflows.

Section 3.3.1. APIs and other Control Interfaces

The FG-HPCC will manage on-premises hardware through low-level IaaS APIs similar to those used to provision cloud resources. Open, industry standard APIs, e.g., Sunfish, Redfish, or de-facto standard APIs, like S3 and TerraForm providers, will provision and manage FG-HPCC resources. If an industry standard does not exist for a given API, the API must be open and well documented. APIs will support management of nodes, VMs, network, storage, AI processors, volumes, and other system hardware. The FG-HPCC requires both in-band and out-of-band management support for finer grain access control with strong isolation.

Section 3.3.2. Resource Management and Orchestration

Resource management in the FG-HPCC will require support for isolated allocations that span system elements and thus support for heterogenous resources models. LLNS seeks information on engagements that can enable these types of models in HPC resource management systems (e.g., Flux, SLURM, or others).

LLNS will use Flux (flux-framework.org) to allocate HPC resources and to schedule HPC jobs on future systems, and other NNSA laboratories may use SLURM and/or other systems. Flux is open source, and LLNS encourages research and collaborations that extend its capabilities to support the FG-HPCC. As Flux already supports a heterogeneous graph model for resource allocation and the Flux development team has explored integration with Kubernetes, close engagement with the Flux development team on such efforts is strongly encouraged.

LLNS and the NNSA tri-labs also use Kubernetes extensively, to provision facility services and to allow users to provision their own containerized services. LLNS encourages R&D that would enable Kubernetes to use the capabilities of the FG-HPCC, either independently or in conjunction with a traditional HPC resource management system. LLNS has also developed integrations between Kubernetes and Flux to enable converged HPC/cloud workflows, and R&D that leverage the two together is also encouraged. LLNS anticipates that Flux, Kubernetes, and other higher-level resource management tools will leverage lower-level IaaS APIs to manage networks, VMs, and other low-level system resources.

Section 3.3.3. Guest Operating System

The guest operating system in the FG-HPCC is a critical component, particularly for HPC jobs. While LLNS anticipates that the FG-HPCC will use a virtualized environment, hardware support for GPUs, network fabrics, and other accelerator hardware will still need to be maintained for

guest OS kernels, and the guest operating system will need to have a Security Technical Implementation Guide (STIG) that allows it to be used in classified environments.

LLNS uses TOSS in production. TOSS augments RHEL with kernel patches and packages for large-scale system management, configuration support, and hardware support for key GPUs and networks. It has a STIG that allows it to be used in classified environments. For collaborations with LLNS, LLNS expects the guest OS for new systems to be based on TOSS version N or N-1, where N is the latest TOSS release. Major TOSS releases are based on RHEL releases and follow their cadence.

NNSA laboratories may use other operating systems, such as Rocky Linux, Alma, or SLES. The OS will depend on the use case, and any such OS will require hardware enablement if it is to be used in an HPC environment. Software required to enable hardware should be released in a form that allows the NNSA laboratories to compile it from source against the kernel of their guest OS of choice. LLNS expects that such software will be distributed outside of LLNL, to other sites. Early and frequent engagement with kernel.org to upstream any needed hardware support patches is required.

Section 3.3.4. Virtualization and Hypervisors

The FG-HPCC will use virtualization software extensively. Respondents to this RFI are encouraged to provide information related to virtualization software that can integrate seamlessly in an HPC environment. Information regarding overheads of open-source virtualization software, both on the CPU and on GPUs and accelerators is of particular interest. VMs in an HPC setting must spin up rapidly, provide little to no overhead over bare metal, and must support efficient leverage of node hardware by the guest OS. LLNS is also interested in open-source software that could bridge performance and security gaps in CPU and GPU virtualization, as well as software that would allow multiple VM tenants on a node to share GPUs.

Section 3.3.5. Network Isolation

Many ways to isolate VMs on a node and over a network exist; LLNS seeks information related to high performance, low-latency implementations of isolated virtual networks. Mechanisms for traffic encryption, data-center-wide key management, memory encryption, other isolation techniques, and any relevant performance characteristics are of particular interest.

Section 3.3.6. Storage Isolation

LLNS and other NNSA laboratories envision modernizing their storage infrastructure to support the sharing of storage resources across security domains. Technologies that can enable storage for one security domain to be isolated from storage for other domains on the same device will

be needed for the FG-HPCC vision. Encrypted volumes and how they might work with the data-center wide key management mentioned in Section 3.3.5 are of interest.

Section 3.3.7. Complex-wide Federation

While LLNS envisions that the initial transformation of NNSA sites to FG-HPCCs will take place locally, sites will eventually want to federate resources and to make FG-HPCC capabilities possible across data centers. Much as clouds can coordinate resources across different regions, LLNS envisions that FG-HPCC instances can coordinate across sites and organizations. LLNS seeks information about any capabilities that could enable federation of data center networks across sites, potentially spanning the continental US. Permissions models that allow sites to federate local authentication through technologies such as DOE OneID, NNSA's ESNHub, or other technologies are of interest.

Section 4. Collaborations

LLNS seeks information on possible collaborations on software and hardware to enable the FG-HPCC.

Section 4.1. Open-Source Software Projects

The FG-HPCC will rely critically on an open-source software stack to manage and to control the data center, as well as on-node and network virtualization technologies. LLNS and other NNSA laboratories use open-source extensively, including software developed under the ASC program (e.g., Flux, Spack, Lustre, TOSS) as well as external open-software projects that LLNS builds as part of its OS distribution and to which it contributes. LLNS seeks information regarding potential collaborations that could improve existing open source or open standards, as well as collaborations that can develop new, well documented APIs or control plane infrastructure for the FG-HPCC.

Section 4.2. Hardware Testbeds

To enable the FG-HPCC and to begin to bring cloud-like technologies into the HPC center, LLNS seeks information on potential collaborations around hardware and software testing. LLNS can provide a testbed cluster within LC to conduct experiments and tests with novel hardware elements, and LLNS is willing to consider committing effort to investigate and to ensure that potentially high-value hardware and enabling software are production-ready.

Section 4.3. Software Development Practices

LLNS strongly prefers that software is open-source, and that it is provided through a third-party hosting site such as GitHub, where LLNS can work directly with offerors and collaborators on

issues. Software should be clearly version-tagged and should have a well-structured release process. As LLNS and other laboratories build their own OS releases, source availability is critical for enabling rebuilds with potentially ABI-altering extensions and modifications to the kernel.

Section 5. Future HPC Center Use Cases

This section describes several emerging workflows that the FG-HPCC will support. These scenarios are not intended to be all encompassing. Instead, they are examples intended to give readers an idea of how the envisioned technologies will be used in practice.

Section 5.1. Management and Orchestration of Services

Section 5.1.1. AI-Augmented Simulation

Traditional numeric ModSim codes will be enhanced, both with embedded AI surrogate models tightly incorporated into multi-physics, multi-scale simulations, and with AI orchestration models driving simulation campaigns. Emerging workflows will require a blend of traditional double-precision floating point operations for ModSim codes, tightly interwoven with AI-centric, low-precision floating point for “small” model inference, likely on the same node. The orchestrating AI model that launches ModSim jobs will also need to run on AI-capable hardware, at modest to large scale.

Section 5.1.2. Digital Twins

In addition to AI integrated directly into the simulation workload, the FG-HPCC will support increased use of digital twins to model components produced at PA sites (e.g., Y-12, Kansas City National Security Complex (KCNSC), Pantex). These sites manufacture components according to designs produced by DA sites. A digital twin of a 3-D printed part can be used as an input to a traditional simulation, enabling parts to be “born certified” via increased simulation accuracy. AI models will serve two roles for digital twins. First, AI models will be used to monitor physical systems and to generate streams of measurements based on the monitoring data. Second, AI models will be used to monitor and to integrate multiple streams of measurements to annotate existing digital models. In both of these use cases, the computational needs of the AI models are more focused on real-time execution and on-demand scheduling to be consistent with live experiments at the production facility.

Section 5.1.3. Inverse Design

Generative AI models will be used for inverse design capability, specifically to explore a design space rapidly and to identify key regions that should be examined with further ModSim runs. In this type of workflow, thousands or more traditional ModSim runs will be orchestrated with scientifically varied inputs. The outputs of the simulations will be collected and the full input

and output data will inform a human designer or train a surrogate model of the simulation. These types of generative models can be substantially more expensive to execute than a typical surrogate model, but lighter weight than the orchestration models discussed in Section 5.1.1. These generative models may benefit from AI accelerators, possibly in distinct AI elements.

Section 5.2. AI Workloads and their Computational Motifs

FG-HPCC users will begin to exploit AI in even more scenarios than the three mentioned in Section 5.1. A more complete list with computational characteristics is:

1. AI-augmented simulation (Section 5.1.1): HPC ModSim integrated with AI such that many small inference requests short circuit calculations, tightly coupling AI and 64-bit interaction;
2. AI-augmented simulation campaigns (Section 5.1.1): Using an AI model (possibly a large reasoning model) to orchestrate modeling and simulation, which will include LLM inference, traditional ModSim with coupled surrogate models, surrogate model re-training, and orchestration and fine-tuning of LLMs;
3. Inverse-design (Section 5.1.3): LLM or other model to orchestrate batched jobs, modeling and simulation jobs along with AI surrogates;
4. Specialized foundation model development: Large-scale, compute-intensive training of transformer models for a specialized scientific domain, possibly coupled with modeling and simulation to generate data;
5. Creation of data-surrogate models: Train a DSM (domain specific model) to represent and to compress a multi-modal data set (including rare events); and
6. HPC Code assistant: An LLM assists with porting HPC code, likely run as a persistent service on AI-capable nodes.

For each of the models above, the training requirements vary substantially, from hundreds of compute hours to potentially exaflop days for the largest models. The data I/O requirements for hybrid AI workflows differ substantially from a traditional ModSim checkpoint paradigm:

1. Data sets will range from Terabytes (TB) to Petabytes (PB) and contain up to billions or trillions of samples / tokens;
2. Storage systems will be required to serve complex sets of this data in a read-mostly, near random-access pattern, or to allow for in-situ ingestion of data streams for online training;
3. Provenance of the data will become a first-class property that is critical for understanding the fidelity and veracity of trained models;
4. Data may be sourced from real-time edge experimental facilities, such as the National Ignition Facility, the Advance Manufacturing Lab, the Vera Rubin telescope array, or the

Scorpius laser facility. This live data is crucial for digital twins, but also the calibration of ModSim codes; and

5. Finally, as models become part of AI-augmented simulation workflows, reproducibility and auditability require the ability to recreate exact workflows, including specific variants of the models.

Models that are used for inference and training will range from millions of parameters that run efficiently on portions of a modern GPU or AI accelerator, up to trillions of parameters that require hundreds to thousands of accelerators just to load, let alone train efficiently. Training costs of the most sophisticated models will stretch into many exaflop days, with an exaflop month(s) of ModSim runs and large surrogate model inferences required to generate the supporting training data set.

Section 5.3. Services & Orchestration Capabilities

A majority of these workloads require significantly more services to orchestrate large sets of runs, training, inference, and model updates than HPC centers have previously encountered. Unlike prior workflows, these scenarios require the HPC user to manage the dynamic scheduling of their own ensemble runs. Moreover, the workflows require an intelligent system that can place different types of jobs on the most appropriate resources, at different scales, in conjunction with needed data. These types of workflows require thinking about more than MPI batch jobs—services communicate through many more network layers and libraries. In many cases, the granularity of computational work is much finer than what we have seen in the past. In large ensembles, the training jobs may be small and we may need to run several tens, hundreds, or thousands of small simulations to amass sufficient training data.

Anticipated FG-HPCC workflows clearly need co-scheduled services, deep resource awareness, and strong user isolation. The FG-HPCC must ensure that services do not unintentionally expose or leak data to other users, and the FG-HPCC will need orchestrators to ensure that analysis, training, and inference jobs from the same workflow are efficiently co-scheduled.

Section 5.4. Data-Centric Computing in a Connected Complex

For the digital twin use case, the FG-HPCC must deploy persistent services, not only to orchestrate jobs *within* the center but to connect with production agencies *outside* it. The FG-HPCC will need frequent updates of data from the PA sites to be sent to persistent services to update and to redeploy models.

Nearly all of the new ML components require careful placement near training or analysis data. Large data sets of current HPC centers reside in filesystems and tape archives across sites, but FG-HPCC workloads will need to manage data sets differently. The FG-HPCC will keep persistent

data warehouses, data lakes, and large data sets that will need to be used across multiple sites. Regular cross-site replication, versioning, and updates of these data sets will be critical, as will staging of compute jobs near data and data near appropriate compute resources.

Many NNSA data sets are subject to need-to-know restrictions, so flexible access controls as well as flexible data movement will be essential. The FG-HPCC cannot allow unauthorized access to data sets, but FG-HPCC users who should have access should not have to struggle to gain it. Sharing should be simple, fine-grained, and should be possible without excessive copying or wait times. The connected complex requires that the FG-HPCC enable users to work together efficiently within and across sites.

Section 5.5. Section 5.5 – Developer and Operational Workloads

ML models and simulations must be versioned and managed like code to support these workflows, and the deployment tools for developer and operational workflows are typically distributed services. Developers need Continuous Integration (CI) to trigger easily both for code and model changes that occur inside the FG-HPCC, at trusted sites outside the FG-HPCC, and for support libraries on external sites like GitHub. ML workloads require frequent model updates and redeployments, and often require a human in the loop. Engineers use tools like Jupyter, Colab, and SageMaker to create, to tune, and to deploy AI models, and ensuring that this type of productive iteration is possible at the FG-HPCC is critical.

Section 6. RFI Topics and Questions

This section provides a list of specific topics and questions for respondents to this RFI. Topics of interest are based on the FG-HPCC vision described in prior sections. Respondents to this RFI may address any technologies that they believe are relevant to the FG-HPCC vision. LLNS does not expect respondents to address all topics—they may respond to portions of the RFI, or even a single part of the RFI, without addressing every aspect.

Section 6.1. Response Format

No specific format of responses is required although it is recommended that respondents provide a detailed high-level overview of each topic that they address, describing the capabilities that the solution will provide and how these capabilities will enable the FG-HPCC objectives. LLNS is particularly interested in impact on the HPC market and the potential for broad adoption (e.g., in cloud or other industry data centers). Solutions that have the potential for broader adoption beyond HPC are highly desired, as they do not rely solely on NNSA or its laboratories for sustainment.

In addition to topic responses, we are requesting responses that outline potential collaborations around:

1. Open-source software development;
2. Hardware and software testing;
3. Software hardening; and
4. Standardization efforts.

As part of the review process for this RFI, LLNS and partner laboratories intend to identify the most promising potential collaborations and to identify NNSA staff to work with their respective RFI respondents to effect long-term FG-HPCC goals. Activities that may lead to broader community adoption and contribution are of particular interest, but RFI responses addressing NDA or otherwise private activities among LLNS, the tri-labs, and specific industry partners are also welcome for consideration.

LLNS does not mandate a particular software architecture, and respondents are encouraged to detail their own ideas of how the FG-HPCC control plane could work. Figure 1 in particular is meant as a guideline.

Comments that include information that is not widely published should include source data or citations.

Section 6.1.1. Roadmap Description

LLNS is interested in hardware and software technologies in existing roadmaps that can enable the FG-HPCC. Suggested categories:

1. Low overhead security and isolation capabilities;
2. Processors, GPUs, and accelerators for modeling and simulation;
3. Processors, GPUs, and accelerators for AI;
4. Low-overhead virtualization for CPUs and GPUs;
5. I/O, storage and composable storage services;
6. Secure data-center networks with support for isolation;
7. Long-range connectivity between data centers (e.g., between NNSA laboratories), while preserving isolation and other capabilities of an on-prem data center network;
8. Network configuration, isolation mechanisms, and bridging;
9. Open-source system software stack;
10. Security guarantees for any of FG-HPCC capabilities;
11. Platform -as-a-Service (PaaS) or other high-level solutions that enable users to leverage FG-HPCC benefits more easily without having to learn low-level IaaS capabilities; and
12. Any FG-HPCC enabling technologies not listed above.

Section 6.1.2. Gaps

LLNS is interested in gaps that may prevent realization of FG-HPCC objectives in the 2029-2030 timeframe. Respondents should focus particularly on:

1. Gaps that would prevent technologies of interest from being used in an HPC environment (or vice versa)—especially overheads, latencies, and other performance gaps that may make roadmap items unsuitable for an HPC environment;
2. Gaps in the HPC center software stack that prevent broad usage of cloud-like technologies in on-premises environments like HPC centers;
3. Advancements needed to satisfy the needs of both HPC and broader industry; and
4. Standardization efforts that could mitigate gaps with open-source solutions, particularly with broad support from industry partners;

Section 6.1.3. Software Collaboration Areas

LLNS is interested in software collaborations that could fill gaps mentioned in Section 6.2. and described in Section 4. Suggested (but not exhaustive) list of topics:

1. Advancements in low-overhead isolation software that could enable users to partition nodes, CPUs, GPUs, AI accelerators, memory and other relevant resources in a multi-tenant system;
2. Encryption or confidential computing technologies with sufficiently low overhead for use in an HPC or AI-oriented system;
3. Low or zero-overhead network virtualization technologies that enable isolated (cryptographically or otherwise), user-and job-specific silos within a larger multi-tenant system;
4. Adaptations of existing HPC software packages for the FG-HPCC, e.g., explorations of how MPI, network libraries, or GPU drivers could work in virtualized or containerized environments; and
5. Open APIs, control planes, and their implementations to facilitate automation throughout a large-scale HPC/cloud data center, or within a node. We require interoperability and extensibility with existing industry standards and, where standards do not exist, APIs that could establish industry standards are of particular interest.

Open source software collaborations are *strongly* preferred, but respondents are welcome to mention source-available or other solutions that could be integrated into this vision of the FG-HPCC, as well as a plan for how NNSA, other DOE laboratories, and vendors might collaborate on a proprietary solution.

Section 6.1.4. Other Collaboration Areas

LLNS is interested in other ways could collaboration help to further FG-HPCC goals. Suggested topics:

1. Standards documents—either through work on an established committees or work to propose and to build such a standards body if none exists;
2. Collaborations on hardware/software testing in the open LC HPC Center, especially to enable co-design;
3. Hardware, software, and network Integration of AI accelerators into the LC HPC center; and
4. Any other collaboration deemed by respondents to further FG-HPCC goals.

For hardware collaborations: LLNS and other NNSA laboratories are particularly interested in working with respondents to test, to prove out, and to harden FG-HPCC related hardware. The NNSA laboratories have testbed clusters where this type of experimentation is feasible, under NDA if needed.

Section 6.2. RFI Review Process and Collaboration Selection

RFI responses will be reviewed by experts from the NNSA tri-labs (LLNL, LANL, and Sandia) and potentially by external experts from DOE Office of Science laboratories as well as DOE employees and other federal employees.

Responses will be used to for two purposes:

1. To make DOE laboratory and federal staff aware of upcoming technology developments, *and*
2. To prioritize a set of industry/NNSA collaborations to pursue at LLNL and other tri-labs.

Reviewers will rate potential collaborations according to their laboratories' priorities, and will recommend the most promising project descriptions to NNSA HQ for consideration and prioritization.

The anticipated timeline is as follows:

- 2025 August 08 — RFI responses due
- 2025 Late August / September — Laboratories review RFI responses
- 2025 Fall/Winter — Feedback provided to RFI respondents if requested

-end-