PERFORMANCE IS A SECURITY REQUIREMENT

SECURING FEDERAL HPC WITHOUT BREAKING THE SCIENCE



By Ian Lee
Director of Advanced Computing Systems
ShorePoint, LLC

NOVEMBER 2026

SHOREPOINT

CONTENTS

Executive Summary....3

Key Insights.... 4

1. The Challenge: What Changed....5

1.1 The standards gap closed....5

1.2 The objections are a property of the system class....5

1.3 Logging requirements were rewritten, and the deadline is close....6

1.4 SI policy compressed the timelines....6

2. Why This Keeps Happening.... 8

3. What Is at Stake: The Arithmetic.... 8

4. Three Decisions.... 10

4.1 Decision 1: Sort the objection.... 10

4.2 Decision 2: Measure the cost, in both performance and schedule....11

4.3 Decision 3: Govern the users these systems are acquiring.... 12

5. Applying the Framework: The Agency Logging Plan.... 13

6. Conclusion.... 15

About ShorePoint.... 15

References.... 16

EXECUTIVE SUMMARY

Federal agencies standing up or operating high-performance computing (HPC) and Super Intelligence (SI) [formerly referred to as artificial intelligence (AI)] compute capabilities encounter a recurring objection from their research computing organizations: not all security controls work or should apply on this system. Some of those objections are well founded and architecturally correct. Most, however, describe a poor fit for the system, not a rejection of the security outcome itself, and can be satisfied through different or compensating controls. A smaller number are true resourcing or architectural constraints.

The underlying differences are now documented in federal guidance. NIST SP 800-223 defines a four-zone HPC reference architecture and identifies the characteristics that separate HPC from enterprise IT. NIST SP 800-234, finalized in May 2026, tailors 60 controls from the NIST SP 800-53B moderate baseline with HPC-specific supplemental guidance. Of the 288 controls in the overlay, NIST identified only two that it describes as generally not applicable to HPC. A blanket claim that security controls do not apply to HPC, or can be set aside on these systems, is therefore not a defensible position. The rest of the baseline still applies, and the risk to these systems' data and mission requires it.

What federal guidance does not provide is a method for telling a valid objection from an invalid one. That decision belongs to the agency, and in most agencies, the technical and compliance staff end up at odds with one another. The pressure to resolve that is rising. The challenge: the compute that trains and serves federal SI models is HPC, and recent SI and cybersecurity directives are reaching these systems on compressed deadlines, often before tailoring work can be completed.

This paper describes three decisions that make the problem tractable:

Decision 1: Sort each objection into one of three objection categories.

Decision 2: Measure what the control costs in mission terms, in both performance and schedule.

Decision 3: Govern the user population these systems are now acquiring: user-built software, multi-tenant access, an expanding external collaborator population, and models and data as assets.

The paper also identifies a cost asymmetry that most agencies overlook. The performance overhead they argue about is real and measurable. The unmeasured authorization delay is typically several times larger, and at the longer delays practitioners describe, more than an order of magnitude larger. In mission terms, that delay is time national security and scientific research programs – including work aimed at medical breakthroughs, nuclear fusion energy advancements, and geopolitical insights – spend waiting on a system the government has already purchased.



KEY INSIGHTS

1. Federal HPC security requires tailored implementation.
NIST SP 800-223 and NIST SP 800-234 recognize the architectural differences between HPC and enterprise IT. Agencies should use this guidance to determine how controls can be implemented effectively while preserving the security outcomes required to protect their systems and missions.

2. Security decisions must account for both performance and authorization delays.
The cost of security encompasses more than the computational overhead. Delays in operational authorization can leave expensive computing resources underutilized while mission work waits. Agencies should quantify performance impacts and schedule impacts together to make informed, defensible decisions.

3. Every security objection should have a category, owner, and documented resolution.
ShorePoint's three-category framework distinguishes architecturally valid objections, objections that can be resolved through compensating controls, and constraints that are actually related to staffing, funding, or unfamiliarity. This distinction helps agencies route issues to the right decision-makers and avoid unnecessary security exceptions.

4. SI and federated science are expanding the HPC security boundary.
User-developed software, multi-tenant access, external collaborators, third-party datasets, and trained SI model weights create security challenges that traditional enterprise control assumptions may not adequately address. Agencies need repeatable governance processes that account for how people, software, models, and data enter, operate within, and leave their environments.

5. Effective HPC logging starts with security outcomes, not maximum data collection.
Agencies should prioritize telemetry according to its value for continuous monitoring, threat hunting, investigation, response, and forensics. Zone-aware logging, representative workload measurements, and documented collection decisions can help agencies manage storage and ingestion costs without losing sight of mission-critical security needs.

1. THE CHALLENGE: WHAT CHANGED

Federal high-performance computing (HPC) security has been contested for as long as agencies have operated supercomputers. What changed between 2024 and 2026 is that the argument now takes place against published federal guidance, on a compressed schedule, in front of a rapidly expanding user population.

1.1 THE STANDARDS GAP CLOSED

NIST SP 800-223 "High-Performance Computing Security: Architecture, Threat Analysis, and Security Posture"[1], published in February 2024, established the first federal reference architecture for HPC security. It divides an HPC system into four function zones: the High-Performance Computing Zone, the Data Storage Zone, the Access Zone, and the Management Zone. Each zone carries a distinct security posture, reflecting how these systems are actually built and operated.

NIST SP 800-234 "High-Performance Computing (HPC) Security Overlay"[2], finalized in May 2026, builds on that architecture with a security control overlay. It tailors 60 controls from the NIST SP 800-53B "Control Baselines for Information Systems and Organizations"[3] moderate baseline, adding HPC-specific supplemental guidance intended for direct adoption by federal HPC operators.

The practical effect for an agency leader is that two things are now true at the same time:

A. An HPC organization's claim that its systems differ architecturally from enterprise IT is a documented federal position rather than a local assertion.

B. The limits of that difference are documented as well. The overlay tailors 60 controls. It does not exempt an HPC system from the rest of the baseline, and it does not authorize a blanket exception from all security controls.

Of the 288 controls in the overlay, 60 required HPC-specific guidance, and of those, only two were identified as generally not applicable to HPC. Roughly one third of those 60 were tailored because of performance or capacity impacts. The other 228 apply as written.

1.2 THE OBJECTIONS ARE A PROPERTY OF THE SYSTEM CLASS

An agency leader hearing an unfamiliar technical objection has to judge whether it reflects the system or the organization. In HPC, the evidence points to the system.

The standards are cross-institutional, not one agency's opinion. SP 800-223 was authored by contributors from NIST, the Department of War (DoW) High-Performance Computing Modernization Program, the Laboratory for Physical Sciences, Los Alamos National Laboratory, MIT Lincoln Laboratory, Oak Ridge National Laboratory, Sandia National Laboratories, industry, and three universities. Both SP 800-223 and the companion SP 800-234 were developed through the NIST HPC Security Working Group, which brings together government, industry, and academic practitioners to consider security challenges from multiple vantage points.

The same objections surface outside the standards process. The HPC Security Technical Exchange (STX), a practitioner forum spanning more than 20 federal agencies and national laboratories, sees them raised consistently across mission types, classification levels, and agency cultures. An objection that appears

across multiple agencies operating similar systems is an engineering property, and it should be adjudicated on engineering terms rather than escalated as a compliance dispute. The converse is also true; a blanket objection that security controls do not apply, or can be set aside, is no longer credible. The security posture of the system and the risk to its data and mission require them.

1.3 LOGGING REQUIREMENTS WERE REWRITTEN, AND THE DEADLINE IS CLOSE

OMB Memorandum M-26-14 "Ensuring Effective and Efficient Agency Logging and Network Visibility to Defend Against Evolving Cyber Threats"[4], issued May 22, 2026, rescinded M-21-31 "Improving the Federal Government's Investigative and Remediation Capabilities Related to Cybersecurity Incidents"[5] and replaced prescriptive collection and retention tiering with two outcome objectives:

- Continuous Event Monitoring, and
- Threat Hunting, Investigation, Response, and Forensics.

Agencies are measured against a five-level maturity model, with baseline expectations for data that is searchable for six months and retrievable for a minimum of 12 months.

CISA published the Logging Reference Architecture[6] on August 20, 2026, developed with OMB and the CISO Council. It instructs agencies to "determine what telemetry the agency needs" from the security outcomes they must support, rather than from what a platform already offers or what is easiest to integrate. The stated goal is to answer the operational questions that matter, with collection volume following from that. With the publication of the Logging Reference Architecture, agencies must create their Agency Logging Plans and achieve specific minimum maturity level milestones and reach Advanced (Level 3) status within 320 days (July 6, 2027).

The updated logging requirements in M-26-14 apply to Federal Civilian Executive Branch systems. They do not apply directly to national security systems (NSS) or the DoW and Intelligence Community systems. In practice, the boundary is thinner than the scoping language suggests. Equivalent expectations are frequently applied to those systems through NSS-specific policy or by local authorizing officials, and agencies with mixed estates should plan for the substance of these requirements to reach systems the memorandum does not name.

1.4 SI POLICY COMPRESSED THE TIMELINES

Executive Order 14363 "Launching the Genesis Mission"[7] signed November 24, 2025, directed the Department of Energy to build an integrated SI platform that pairs federal scientific datasets with national laboratory computing resources. OMB Memorandum M-26-16[8], issued July 21, 2026, then directed agencies to treat support for the Genesis Mission as a central research and development priority in their FY2028 budget submissions. Together, the two documents make SI workloads on federal HPC systems a standing budget priority.

Executive Order 14409 "Promoting Advanced Artificial Intelligence Innovation and Security,"[9] issued June 2, 2026, set 30-day deadlines for the Committee on National Security Systems to prioritize the cyber defense of national security systems, for the Secretary of War to do the same for DoW systems, and for CISA to issue Binding Operational Directives for civilian federal systems. Together, these orders put federal SI workloads and their supporting HPC systems squarely under directives built for enterprise IT.

HPC TIMELINE: KEY DATES

AUGUST 27, 2021
OMB M-21-31 issued, setting the prescriptive logging tiers

DECEMBER 2022
TOSS 4 STIG publicly released

FEBRUARY 2024
NIST SP 800-223 published

NOVEMBER 24, 2025
Genesis Mission launched by Executive Order 14363

MAY 4, 2026
NIST SP 800-234 published

MAY 22, 2026
OMB M-26-14 issued, rescinding and replacing M-21-31

JUNE 2, 2026
EO 14409 issued, with upgrading federal cyber defenses among its three focus areas

JULY 21, 2026
OMB M-26-16 issued, making Genesis a central FY2028 R&D priority

JULY 22, 2026
First Genesis Mission awards: 278 projects selected from more than 5,000 applications

AUGUST 20, 2026
CISA Logging Reference Architecture published

NOVEMBER 18, 2026
Agency Logging Plans due to OMB and CISA

JULY 6, 2027
Advanced (Level 3) logging maturity deadline

2. WHY THIS KEEPS HAPPENING

The standoff between security organizations and research computing organizations is structural. Neither side is behaving unreasonably. Both act with the mission in mind. Research computing organizations want speed to results. Security organizations want the integrity and protection of those results. Agreement starts with each side recognizing that the other's priority serves the same mission.

- **Different measures of success.** Research computing organizations are evaluated on utilization, throughput, and time to result. Security organizations are evaluated on control compliance and audit findings. Both are doing what their mandate rewards.
- **Performance primacy is a documented design consequence.** SP 800-223 lists, among general HPC threats, the vulnerabilities introduced by prioritizing performance in system design and operation. Designers, for example, make deliberate choices to build less redundant systems in order to achieve performance. A spare or backup HPC system is often too costly to justify, and most HPC missions lack a service-level agreement that would require one.
- **The user population is changing faster than the access model.** The Department of Energy's Genesis Mission[10] drew more than 5,000 applications and selected 278 projects for award negotiation in July 2026, with additional awards expected. New SI systems are being built through public-private partnerships, such as the Lux SI supercomputer at Oak Ridge National Laboratory (ORNL), co-developed with AMD, HPE, and Oracle[11]. Each expansion brings collaborators, models, and data that the original access model did not anticipate.

Both security and research computing organizations are being asked to reach agreement faster than they ever have, using a vocabulary they do not share.

3. WHAT IS AT STAKE: THE ARITHMETIC

The security and performance debate in federal HPC is usually conducted in percentages, which is easiest to judge against a real system. Lawrence Livermore National Laboratory's El Capitan[12], the fastest U.S. system on the June 2026 TOP500 list[13], has a peak performance of 2.75 exaFLOPS and was verified at 1.809 exaFLOPS on High-Performance Linpack (HPL) across more than 11,000 compute nodes. The system was announced in 2019 at a contract value of $600 million[14]. A one percent performance cost on a system of that class is roughly 18 to 28 petaFLOPS depending on the basis used, or approximately $6 million measured against acquisition cost. The low end of that range, 18 petaFLOPS, is roughly the full performance of Sequoia, the #1 TOP500 system in 2012[15].

Now consider the other cost. Amortized across a typical seven-year lifecycle, a system in that class represents roughly $86 million per year of delivered capability. In non-attributed discussions within the HPC Security Technical Exchange, practitioners from multiple agencies have consistently described delays of 6–24 months between system acceptance and operational authorization for mission-oriented HPC systems, with the longer end more common at higher classification levels.

At 18 months, the capability purchased, but never realized, is roughly $130 million. Even at six months, the unrealized capability is roughly $43 million, about seven times the one percent performance cost.

The one percent argument is real. Across the range practitioners describe, it is also between seven and nearly 30 times smaller than the cost of the delay that the argument itself produces. Agencies measure performance impacts of HPC systems carefully and seldom quantify the cost of delayed authorization.

Four consequences follow from leaving this unadjudicated.

A. **The standoff resolves itself poorly**. Two outcomes are common when no one owns the decision. The first is a broad exception that may leave the system with unnecessary risk. The second is a baseline enforced without tailoring, which makes the system unusable for the work it was purchased to do.

B. **Mission objectives are missed.** Delivered capability sits idle while the mission that acquired it waits to start working. On a system supporting nuclear stockpile stewardship, that is deferred national security work that can place our country in an unnecessary risk posture in a time of dynamic geopolitical shifts. On one supporting biomedical research, it is deferred progress toward new cures, delays that could cost human lives.

C. **Genuine risk goes unexamined.** While the familiar controls are debated, the exposures specific to these systems receive less attention: access controls and authorization on shared parallel file systems at scale, job and scheduler isolation, the supply chain for user-built scientific and SI software, node integrity, and trained model weights (the learned parameters that constitute a trained SI model) as an asset class in their own right.

D. **Collection costs grow without detection value.** The failure pattern under the previous logging memorandum was unsustainable storage spend and compliance activity that mistook collection for detection. HPC telemetry volumes make that pattern more expensive, not less.

None of these outcomes is inevitable. Agencies that adjudicate deliberately get both the security posture and the science they paid for, and the three decisions that follow describe how.



4. THREE DECISIONS

The following framework is intended for the agency leader who must adjudicate between the mission organization, typically research computing, and the security organization, both of which have a defensible case. It assumes no HPC engineering background and produces a documented record.

Before adjudicating anything, establish a shared description of the system. Use the four zones from SP 800-223 as the common map. State explicitly what is being protected: data, trained models, allocation, availability, or the integrity of results. Separate the system's security categorization from the sensitivity of individual workloads.

4.1 DECISION 1: SORT THE OBJECTION

Every objection belongs in one of three categories. The category determines the response, and the response determines who owns the next step.

THREE-CATEGORY FRAMEWORK FOR SECURITY OBJECTIONS

CATEGORY	THE TEST	THE RESPONSE
Architecturally valid	The control cannot function as written given how the system is built.	Tailor against the overlay, document the tailoring, cite it.
Valid with a compensating path	The control is a poor fit as written, but the security outcome is achievable elsewhere in the architecture.	Implement the compensating measure and document the evidence it produces.
Not a technical objection	The constraint is staffing, funding, or unfamiliarity rather than architecture.	Record it as a resourcing decision, not a security exception.

ARCHITECTURALLY VALID.

Ask whether the control assumes per-host agents, persistent local storage, or inspection in the data path. Ask whether its cost is bounded or scales with node count or I/O volume. Ask whether the impact has been measured or only asserted. SP 800-223 supports this category directly, noting that enterprise security tools are generally designed for standalone devices, that installing forensic tooling across all compute and storage nodes is unreasonable, that remotely mounted storage can disable some tools outright, and that available Security Technical Implementation Guides (STIGs) are typically written for servers and desktops rather than HPC systems. The notable exceptions are the DISA STIGs for Tri-Lab Operating System Stack (TOSS)[16] 4[17] and 5[18] STIGs developed with Lawrence Livermore National Laboratory as the TOSS vendor.

VALID WITH A COMPENSATING PATH.

Ask what outcome the control is actually trying to produce, where else in the architecture that outcome can be produced, and what evidence the alternative would generate for an assessor. Placement is the usual answer. Identity and access enforcement belongs at the access boundary. Configuration and integrity management belong in the Management Zone. Compute nodes should carry as little in-path tooling as the outcome allows. Encryption decisions belong at the storage layer, where their cost can be measured.

Malware scanning illustrates the category. SP 800-223 observes that HPC storage easily reaches a petabyte or more, that existing scanning tools are built for single machines and are not efficient at that scale, and that scanning can degrade the performance of running jobs. It then supplies the compensating approach: scan at write and perform subsequent batch scanning at read.

NOT A TECHNICAL OBJECTION.

Ask whether the control has been attempted or only assumed infeasible, what it would cost to implement properly, and whether the objection concerns feasibility or priority. In practice, this category is usually a legitimate resourcing gap rather than obstruction, and treating it as obstruction guarantees the wrong outcome. What matters is that it is recorded as a resourcing decision, because resourcing gaps and architectural constraints are remediated by different people on different timelines.

Audit logging on a parallel file system shows how the sorting works in practice. During a large simulation, file creation, modification, and access events on a parallel file system generate volumes that no enterprise collection design anticipated. The volume is genuine, so this is more than a resourcing gap. The desired security outcomes remain achievable through the application of risk-based, HPC zone-aware, prioritized logging, so it is not an architectural impossibility. It belongs in the second category, and Section 5 walks through this example in more detail.

WHAT GOOD LOOKS LIKE: every open objection has a category, a date, and a named owner.

4.2 DECISION 2: MEASURE THE COST, IN BOTH PERFORMANCE AND SCHEDULE

An unmeasured cost is argued about indefinitely. This is the decision that ends the standoff, and federal guidance has already recommended it.

SP 800-223 recommends that organizations test and benchmark the performance penalty of security tools to determine whether it is acceptable and observes that doing so would also encourage more performance-aware tool design. It further recommends building security into initial system design rather than adding it afterward, since add-on tools tend to cost more performance than native measures, and avoiding one-size-fits-all approaches by differentiating node types and applying appropriate controls to each.

More than two years after that recommendation, systematic measurement remains uncommon across the federal HPC community. As of this writing, a cross-agency practitioner effort is working to quantify control impacts in a comparable way, and research investment is being directed at the same question. Until that work matures, agencies can begin to close most of the gaps themselves:

- **Express security cost in mission units.** Specify the percent of job throughput, wall-clock delta on a representative workload, or allocation hours consumed. The phrase “some overhead” cannot be evaluated and therefore cannot be resolved.
- **Benchmark with representative workloads rather than synthetic ones.** A control that is invisible on a benchmark may be significant on a production simulation, and the reverse is equally common.

- **Measure the schedule cost with the same discipline.** Track elapsed time to authorization as a program cost. Agencies that regularly report performance overhead and speed-ups to two decimal places routinely carry authorization delay as an unquantified footnote.
- **Agree on the decision criteria before the measurement, and name who decides.** The criteria can be a risk-based range rather than a single number, but they should be set in advance. Deciding afterward converts a technical result into a negotiation.
- **Make the measurement an authorization and assessment artifact.** A documented tradeoff survives staff turnover and does not have to be relitigated at each assessment.

WHAT GOOD LOOKS LIKE: a control-by-control record showing measured performance cost, measured or estimated schedule cost, and a named decision owner.

4.3 DECISION 3: GOVERN THE USERS THESE SYSTEMS ARE ACQUIRING

Two properties of HPC systems break assumptions embedded in enterprise baselines. Two new changes are arriving with SI and federated science programs. All four concern people and software rather than hardware, and all four are getting less attention than they warrant.

A. **User-developed software.** HPC users routinely modify and recompile their own code to tune it to the hardware. Publication of common vulnerabilities and exposures for scientific software is not routine, and within that community the identification and remediation of performance regressions is generally the higher priority. Enterprise software supply chain expectations assume vendor-supplied binaries and a patch cadence. Neither assumption holds here, and extending supply chain governance to user-built scientific and SI software stacks is where most of the real exposure sits.

B. **Multi-tenancy.** SP 800-223 identifies exploitation of multi-tenant environments as a major threat to HPC systems, and notes that because of the implied delegation of authority across distributed HPC systems and file systems, root on a compute node may be equivalent to root across all zones depending on configuration. This is one exposure where the enterprise security concern is well founded, and agencies should confirm it has been addressed rather than tailored away.

C. **An expanding external population.** Federated science and SI initiatives introduce collaborators whose vetting, scoping, and offboarding were not part of the original access model. Federate identity rather than provisioning locally, implement role-based access control tied to specific projects rather than the individual, and set expiration by default, with periodic re-evaluation informed by continuous monitoring.

D. **Models and data as assets.** Determine how third-party models and datasets enter the environment, how outputs leave it, and what happens to model weights when a project ends. Weights are frequently the most valuable artifact on the system and the least likely to appear in its system security plan. A copy of the weights is, in practice, a copy of the model and the knowledge embedded within and therefore contains the investment that produced it.

WHAT GOOD LOOKS LIKE: external collaborator onboarding is documented, repeatable, and expires without intervention.

5. APPLYING THE FRAMEWORK: THE AGENCY LOGGING PLAN

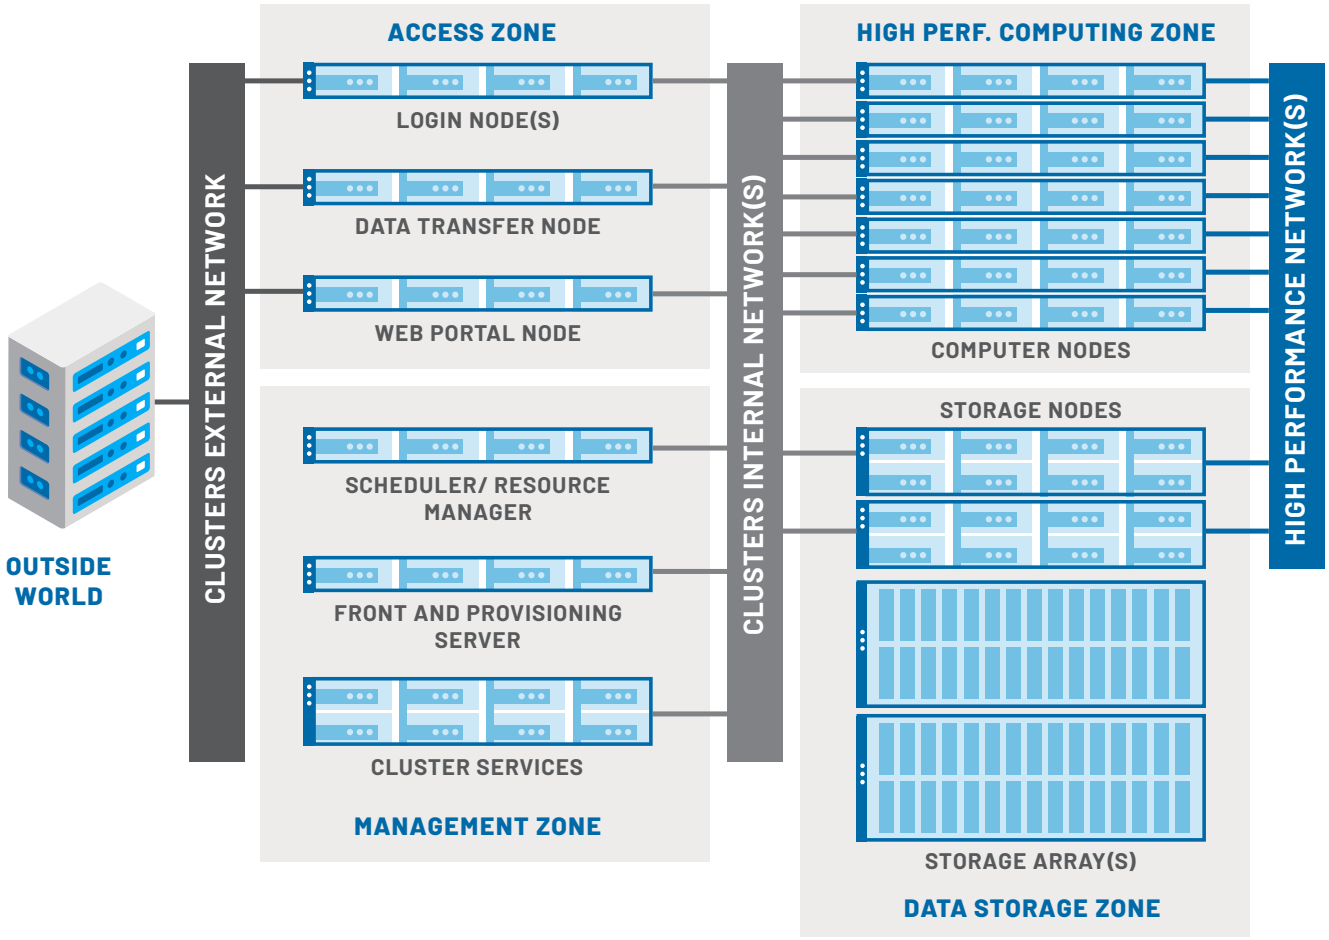
The framework is easiest to see applied to a requirement agencies are working through now. In accordance with M-26-14, Federal Civilian Executive Branch agencies must submit an Agency Logging Plan to OMB and CISA by November 18, 2026, using a template CISA has provided, against the Logging Reference Architecture published August 20, 2026.

- **The objection.** Audit logging on a parallel file system during a large simulation produces event volumes that exceed what the agency's collection design was sized for. Scheduler, interconnect, and node health telemetry add to it. The HPC organization states that comprehensive collection is not feasible.
- **Sort it.** The volume claim is measurable and correct, so it is more than a resourcing issue. The security outcome remains achievable through prioritized selection of logging sources, so it is not an architectural impossibility. This is a compensating-path decision.
- **Measure it.** Translate the proposed audit configuration into gigabytes per day under representative job load, then into annual cost at the agency's per-gigabyte ingest rate for its analytics tier. The resulting figure focuses the discussion on the real cost of the logging source selections being made.
- **Resolve it.** Separate telemetry that must support continuous event monitoring, and therefore must be searchable in near real time, from telemetry that supports investigation and forensics and needs only to be retrievable. Sample or aggregate where per-event fidelity carries no investigative value, and document why.
  This is what the Logging Reference Architecture asks for when it directs agencies to work from required outcomes rather than from what is easiest to collect. SP 800-234 gives the zone ordering directly, directing that logging detail be prioritized in the Management Zone over the Access Zone, with the Data Storage Zone and High-Performance Computing Zone lowest, and that storage and performance resources be allocated according to the risks being managed.
- **Scope it.** Identify which systems are in scope. Systems designated as national security, DoW, or Intelligence Community systems fall outside of M-26-14 but should expect equivalent expectations through their own policy.

The outcome is a documented, prioritized logging decision the agency can defend at assessment, produced in weeks rather than quarters, and reusable the next time a collection requirement arrives.

One caution when working from the overlay: Its logging guidance references M-21-31 and CISA's M-21-31 implementation guide, both of which have since been superseded by M-26-14 and the Logging Reference Architecture. The zone-based prioritization logic still holds, and agencies should map the overlay's guidance onto the current outcome objectives rather than the retired event categories.

HPC SYSTEM REFERENCE ARCHITECTURE, FROM SP 800-223





6. CONCLUSION

The tradeoff between security and performance in federal HPC is real in specific, identifiable controls, and absent in most of the rest. Knowing which is a decision an agency makes deliberately or by default, and the default outcomes are poor: an exception that cannot be defended, a baseline that drives users off the system, or a system that sits unauthorized while its capability and value depreciate.

Three actions are available to any agency leader whose portfolio includes HPC or SI compute, independent of any particular deadline:

- Name someone who owns the adjudication.
- Run the first sorting exercise against the exceptions already open.
- Begin measuring what controls cost in mission terms, including the cost of time spent waiting for authorization.

There is also a near-term forcing function. The Agency Logging Plan due to OMB and CISA on November 18, 2026, is the first logging requirement written around security outcomes rather than prescribed event categories, and it applies to systems for which enterprise control baselines were never written. Agencies that have adjudicated their objections will write that plan in a matter of weeks. Agencies that have not will discover the argument while writing it.

The standards now exist. The tailoring work does not do itself. Systems arriving under current SI and advanced computing initiatives will be authorized against whatever adjudication practice an agency establishes this year, and the leaders best positioned are the ones who know which risks they are accepting on these systems and who decided to accept them.

ABOUT SHOREPOINT

ShorePoint is a cybersecurity services firm dedicated to strengthening the cyber resilience of federal agencies and their missions. Its Advanced Computing Solutions practice supports federal agencies, national laboratories, and research institutions in securing HPC, SI/ML, and data-intensive environments, spanning security architecture and engineering, risk management and authorization support, security operations and data integration, and independent assessment. The practice includes a co-author of NIST SP 800-234 and the founder of the HPC Security Technical Exchange.
Learn more at ShorePointInc.com

REFERENCES

[1] NIST SP 800-223, High-Performance Computing Security: Architecture, Threat Analysis, and Security Posture. February 2024. https://doi.org/10.6028/NIST.SP.800-223

[2] NIST SP 800-234, High-Performance Computing (HPC) Security Overlay. May 2026. https://csrc.nist.gov/pubs/sp/800/234/final

[3] NIST SP 800-53B, Control Baselines for Information Systems and Organizations. September 2020. https://csrc.nist.gov/pubs/sp/800/53/b/upd1/final

[4] OMB Memorandum M-26-14, Ensuring Effective and Efficient Agency Logging and Network Visibility to Defend Against Evolving Cyber Threats. May 22, 2026. https://www.whitehouse.gov/wp-content/uploads/2026/05/M-26-14-Ensuring-Effective-and-Efficient-Agency-Logging-and-Network-Visibility-to-Defend-Against-Evolving-Cyber-Threats.pdf

[5] OMB Memorandum M-21-31, Improving the Federal Government's Investigative and Remediation Capabilities Related to Cybersecurity Incidents. August 27, 2021. https://www.whitehouse.gov/wp-content/uploads/2021/08/M-21-31-Improving-the-Federal-Governments-Investigative-and-Remediation-Capabilities-Related-to-Cybersecurity-Incidents.pdf

[6] CISA, Logging Reference Architecture. August 20, 2026. https://www.cisa.gov/resources-tools/resources/logging-reference-architecture

[7] Executive Order 14363, Launching the Genesis Mission. November 24, 2025. https://www.whitehouse.gov/presidential-actions/2025/11/launching-the-genesis-mission/

[8] OMB Memorandum M-26-16, Ushering in a New Golden Age of American Innovation. July 21, 2026. https://www.whitehouse.gov/wp-content/uploads/2026/07/M-26-16-Ushering-in-a-New-Golden-Age-of-American-Innovation-Fiscal-Year-2028-Administration-Research-and-Development-Budget-Priorities.pdf

[9] Executive Order 14409, Promoting Advanced Artificial Intelligence Innovation and Security. June 2, 2026. https://www.whitehouse.gov/presidential-actions/2026/06/promoting-advanced-artificial-intelligence-innovation-and-security/

[10] U.S. Department of Energy, Genesis Mission. https://genesis.energy.gov/

[11] Oak Ridge National Laboratory, Lux: ORNL's AI + HPC Supercomputer. https://www.olcf.ornl.gov/olcf-resources/compute-systems/lux/

[12] Lawrence Livermore National Laboratory, El Capitan system specifications; TOP500 November 2024 List. https://top500.org/lists/top500/2024/11/

[13] Lawrence Livermore National Laboratory, June 2026 TOP500 List. https://top500.org/lists/top500/2026/06

[14] U.S. Department of Energy, DOE's NNSA signs $600 million contract to build its first exascale supercomputer. August 13, 2019. https://www.energy.gov/articles/does-nnsa-signs-600-million-contract-build-its-first-exascale-supercomputer

[15] Lawrence Livermore National Laboratory, Sequoia system specifications; TOP500 June 2012 List. https://top500.org/lists/top500/2012/06/

[16] Lawrence Livermore National Laboratory, TOSS: Tri-Lab Operating System Stack. https://hpc.llnl.gov/toss

[17] Intelligence Community News, DISA releases TOSS 4 STIG. January 3, 2023. https://intelligencecommunitynews.com/disa-releases-toss-4-stig

[18] NIST National Checklist Program, Tri-Lab Operating System Stack (TOSS) 5 STIG Ver 1, Rel 2 Checklist Details, October 6, 2026. https://ncp.nist.gov/checklist/1332



SHOREPOINT

12901 Worldgate Drive • Suite 650 • Herndon, VA 20170-6014

shorepointinc.com