

FAQs – Innovative Technology Insertion of Advanced HPC Systems

Last updated October 14, 2026

Are software-based solutions within the scope of this effort? While the opportunity appears focused on infrastructure and computing platform capabilities, would solutions that improve system performance, operational efficiency, resource utilization, reliability, automation, or the overall user experience also be considered? Specifically, are technologies that help optimize the use of existing resources, reduce administrative burden, minimize downtime, or enhance consistency across enterprise environments within scope?

No. Software-based solutions are not within scope of this effort which is based on the acquisition of HPC systems and follow-on system administration and maintenance support options. Software-based solutions that improve existing system performance, operational efficiency, resource utilization, reliability, automation, or the overall user experience will not be considered. Technologies that help optimize the use of existing resources, reduce administrative burden, minimize downtime, or enhance consistency across enterprise environments are not considered to be within scope since this effort is based on the acquisition of new HPC systems.

Could any of the acceptance testing be completed at the MFG Factory prior to shipment?

Yes. Portions of the acceptance testing may be completed in the factory prior to shipment.

Is there a general criteria on acceptance testing that can be shared?

It is expected that the government and offerors would collaborate on the final acceptance testing criteria.

For most TI acceptance processes, there are typically two phases – Capability Test (CT) and Effectiveness Level Testing (ELT).

During CT, each production system would be subject to a set of functional and capability tests. The functional set usually includes basic system hardware and software assessments to determine whether the system meets the proposed configuration. The capability set is generally a suite of application and synthetic (I/O, networking) benchmarks intended to represent the workloads anticipated on the system. CT is typically expected to be completed within 5 business days of starting.

Upon completing CT, the system enters ELT, which is generally a 60-day window in which 30 consecutive days of system metrics must be met. These metrics are usually system utilization ($\geq 75\%$), system availability ($\geq 97\%$), system interrupts ($\leq$ minimum # allowed), and node interrupts ($\leq$ minimum # allowed). If a system is unable to satisfy all of the metrics over the 30 day period, the window to complete the metrics will slide until 30 consecutive days are met. For

example, if a system exceeds its system interrupt count on day 29 of ELT, and the first system interrupt occurred on day two, ELT would need to continue for two additional days to surpass the initial system interrupt, and then one more day to reach 30 consecutive days.

What quantitative performance objectives and acceptance benchmarks apply to each DSRC system and scenario?

The government intends to collaborate with the selected offeror on performance objectives and acceptance benchmarks.

Please identify the relevant measures for CPU HPC, AI/ML, storage, and fabric performance, as well as any pass/fail thresholds for the required multi-week burn-in period.

Generally speaking, during Effectiveness Level Testing (ELT, i.e. multi-week burn-in period), the system must meet the metrics outlined in the previous response. Offerors may utilize the prescribed benchmark suite - again, developed in collaboration with the selected offeror – to both meet the required metrics while observing performance levels for issues.

What workload mix and job profiles are expected at AFRL, Navy, and MHPCC?

In general, HPCMP systems will support a wide variety of modeling and simulation workloads. For more information on these workloads, please see: <https://hpc.mil/solution-areas/technology-areas>. Consideration should be taken that some modeling and simulation applications are able to leverage accelerators configured to support double precision calculations.

Regarding AI/ML, a sample listing of the domains would include autonomous systems and robotics; predictive maintenance; situational awareness and decision support; cybersecurity and threat detection; training and simulation; intelligence, surveillance, and reconnaissance (ISR); logistics and supply chain optimization; target recognition and precision engagement; human-machine teaming; and weather modeling and forecasting. Within these domains, specific AI/ML workloads might include synthetic data generation, sparse inverse problems, open-ended design optimization, geometric deep learning, science-constrained deep learning, digital twins, uncertainty estimation, agentic systems, GPU-accelerated multi-physics codes, ethical AI and governance.

Please identify the expected balance of CPU simulation, AI/ML training, AI inference, visualization, data analytics, and interactive use, including typical job scale, GPU count/job, memory/job, concurrency, and user demand.

The offeror should consider the wide variety of workloads referenced in the previous question's response as well as the intended architectural capabilities for each system as stated in APPENDIX A: Potential Solution Scenarios.

What GPU form factor is intended for the GPU-forward use cases?

For the GPU-forward use cases, the form factor would be expected to be the most performant version, typically the OAM or SXM versus a PCIe version, and to have sufficient networking capabilities for multi-GPU data sharing.

Are there mandatory processor, accelerator, or architecture constraints?

Only that they be TAA-compliant and able to support current and future high-performance computing and AI/ML workloads.

Is the desired commonality primarily across DSRC sites, across workload types, or both?

The desired commonality is to achieve the specified “Desired End State” providing a “cost-effective, efficient, and reliable enterprise approach covering multiple HPC systems operating at multiple sites within various classification environments”. The offeror should consider the resulting enterprise acquisition, administration, and maintenance as specified “Requirements” section.

Should CPU-heavy, mixed CPU/GPU, and GPU-forward systems share common host architecture, management tools, fabric operations, and spare parts where feasible, or may systems be optimized independently by site?

The Government desires as much commonality as feasible. Systems may be optimized independently by site, but it is up to the offeror to determine the best approach in optimizing a proposal that supports one or more DSRCs.

Would the Government favor a modular rack-scale approach in which validated CPU, GPU, storage, fabric, utility, and cooling building blocks can be added independently as site demand grows?

It is up to the offeror to determine and propose an approach to support the requirements.

For the MHPCC solution, the Government favors scalable, modular architectures that allow independent expansion to align with future workload growth and phased funding profiles.

Is a common enterprise management and observability model desired across DSRCs, while preserving fully separate operational and security boundaries for each classification domain?

Yes, this is a desirable goal, but it is up to the offeror to determine and propose an approach to support the requirements.

Storage, Data, and Fabric

Is the high-performance parallel scratch file system in offeror scope, or will each DSRC provide existing storage/file-system services?

High-performance scratch storage and associated file-system services should be considered within the offeror's scope. Solutions may use parallel file systems or other high-performance scale-out or distributed storage architectures appropriate to the proposed system and workload.

If in scope, please provide target usable capacity, aggregate throughput, IOPS/metadata needs, expected file profile, checkpoint requirements, and concurrent client scale.

Details of the file system implementations will be determined during scope collaboration with the selected offeror. Offerors should provide optimized solutions to match the expected workload based on the architectures being considered as listed in Appendix A.

Please note: This CSO announcement initiates a multi-step process. The first step is the solicitation of solution briefs. The purpose of the solution brief submission is to identify potential partners that may have promising solutions. Subsequent steps consist of opportunities for selected offerors to deliver solution pitches, demonstrations, and collaborative exchanges with the Government. As such, offerors do not need to submit initial solutions with detailed technical specifications trying to satisfy technical requirements that will be identified in later steps.

Is cross-site data federation, replication, or a shared namespace desired, permitted, or prohibited?

The HPCMP is in the process of implementing an enterprise data fabric based on Hammerspace. The data fabric will be a shared namespace across HPCMP, connected to multiple HPCs at a given classification level. The initial Hammerspace installs will be in the collateral secret and unclassified realms. Any proposed system(s) in those classifications should be able to integrate the Hammerspace solution, making it available to the entire compute pool.

If any cross-site data service is allowed, please identify the approved DSRC pairs and classification boundaries.

See answer above.

What production parity is required for each Test and Development System?

See answer below.

Must the TDS replicate CPU/GPU families, fabric generation, file-system stack, scheduler/container environment, security configuration, and classification level, or is a more limited functional test environment acceptable?

In order to properly support testing of production system updates, it's generally expected that the TDS will match in all of the above-mentioned areas. Exceptions can be made in some instances. Final determinations should be made during scope collaboration with the selected

offeror.

What security deliverables must the system provider provide for RMF/STIG, encryption/key management, vulnerability management, and classified-media lifecycle handling?

Leading up to final system acceptance, the offerors must support the required Assess Only process for all contracted systems. HPCMP Security approval for the Assess Only package (including STIGs, hardware/software checklists, etc) must be received prior to the system being connected to a DSRC's network.

Requirements to support encryption/key management will need to be determined during scope collaboration.

Assuming the selected offeror provides administration and maintenance support, vulnerability management and classified-media lifecycle handling would be detailed in the administration & maintenance requirements and formalized during scope collaboration.

Can the response use a teaming model in which an offeror provides systems, factory integration, Tier 2/3 support, and training while a cleared services partner provides enterprise administration and on-site services?

Yes, a support model comprised of multiple contractor teams is acceptable. All personnel directly involved in the factory integration of the solution must be U.S. citizens.

Depending on the type of Tier 2/3 support, those personnel may also need to be cleared. If Tier 2/3 support involves user data from the system, then a clearance would be required.

Details of the relationships between the offeror teams and the government, including project management, points of contact for various types of support and escalation procedures, should be provided during scope collaboration with the selected offeror.

How should the Navy two-hour response requirement be interpreted?

The Navy two-hour response requirement is intended as a two-hour acknowledgment of an issue.

Does it mean acknowledgment, remote diagnosis, on-site dispatch, service restoration, or full resolution? Are local spare parts required?

Contracted TI offerors typically keep a spare parts cache for each system at a given DSRC.

Section A describes stove-piped support across the DSRCs and states that the Government may obtain system administration and application support from parties other than the system provider. Would the Government consider a solution brief that offers only an enterprise support coordination and governance layer (cross-site service management, SLA/KPI framework, acceptance and transition planning, and vendor coordination among the OEM,

administrator, and third parties), without proposing hardware or cleared on-site administration?

Yes, the Government will consider solutions that offer only enterprise support without proposing hardware or cleared on-site administration. Such solutions should describe how they would support hardware and on-site administration acquired through other sources.

If such a brief is acceptable, should it identify the specific OEM and administrator it would coordinate with, or may teaming with those parties be arranged after selection?

The brief should indicate how teaming would be accomplished.

May a solution brief be submitted by a team, with a prime and subcontractors or partners? If so, should each team member submit separately or should one brief cover the team?

Yes, a brief may be submitted by a team with a prime and subcontractors or partners identified. One brief will be sufficient.

Will small business or socioeconomic status (for example, WOSB, WBE, MBE, or DBE certification) be considered in selection or in later teaming and subcontracting, given the full and open basis of the CSO?

Small business participation will be considered after potential solutions are identified for selection and further actions.

Does the September 21 FAQ's exclusion of software-based solutions permit third-party orchestration software as a component of a complete new HPC-system offering with administration and maintenance support?

The exclusion was intended to remove from consideration any software-based solutions proposed for existing HPC systems. Any software-based solutions that are proposed with new HPC systems will be considered as part of the software stack delivered and/or as part of follow-on administration and maintenance support.

May an administration and maintenance provider submit independently of a new-system acquisition offer, or must its services be included in a system provider's submission?

Yes, an administration and maintenance provider may submit a solution independent of a new system acquisition. The solution should describe how it would support a new system(s) acquired from another source.

Please confirm the controlling submission deadline and time zone. The announcement lists October 23, 2026, at 4:00 pm CT, but labels that date Thursday, October 23 is Friday.

The day has been corrected. The submission deadline is Friday, October 23, at 4 pm CST/5 pm EST.

Section B, item 3 requests a ROM. Can the Government please provide representative CPU, GPU, memory, storage-capacity, I/O, network-throughput, and AI/ML workload profiles or target performance measures so offerors can size systems consistently and demonstrate applicability across traditional HPC and AI/ML workloads? Without this information it's difficult to provide a meaningful ROM estimate.

ROMs are to be based on proposed solutions that consider the Estimated Government Funding Profile in Section A and the Potential Solution Scenarios in Appendix A, as well as the infrastructure capabilities and constraints identified in Appendix B, C, and D. Proposed solutions may indicate a range of representative features and anticipated workloads.

The RFI instructs offerors not to include classified or proprietary information. Could the Government clarify whether non classified, non public information (such as elements of an offeror's product or service roadmap) may be submitted if relevant to the Government's requirements?

Yes, offeror-controlled information may be submitted if it is relevant to the solicitation requirements.

If such confidential roadmap information is ordinarily included in proposal responses when handled with appropriate protections, would the Government permit its submission for this RFI?

See answer above.

Alternatively, should offerors limit their solution briefs strictly to currently commercialized and publicly available capabilities?

See answer above.

If proprietary roadmap details are not permitted in the solution brief, may offerors identify relevant future capabilities at a high level and provide additional details during a later authorized phase?

See answer above. Offerors may also choose to submit information at a high level during Step 1 and provide additional details in subsequent steps.

For the stated annual system-purchase budget and the separate administration and maintenance budget, should the ROM include or exclude test and development systems, racks, liquid-cooling equipment, installation, integration, acceptance testing, software licensing, training, and initial support?

The ROM should include all the items stated in the question.

Please clarify the preferred breakdown for acquisition, site-specific costs, and recurring support.

See answer above. There is no preferred breakdown of costs, other than recurring support will not be part of the initial acquisition cost.

The announcement references a five-year support term after acceptance and also describes a one-year period with options for up to six additional years. Which base and option-period structure should offerors use for the ROM and solution brief, and should support pricing be provided for five years, seven years, or both?

The system administration and maintenance support is expected to have a one-year base period with six options of one-year each for a total of seven years if all options are exercised. Pricing shall be proposed by the offeror.

For the unclassified, Secret, and Above-Secret scenarios, which security requirements and authorization artifacts must be addressed in the solution brief versus deferred until the RFP or post-selection phase?

Solution briefs should demonstrate an understanding of security requirements and authorization artifacts. Specific details will be addressed in subsequent phases.

Common Technical Architecture: The CSO identifies potential HPC system acquisitions at AFRL, MHPCC, and Navy DSRCs and describes an enterprise approach to system administration across multiple HPC systems and DSRCs. Is the Government's objective to maintain a common technical architecture, or common architectural baseline, across the DSRCs, or may offerors propose substantially different architectures optimized for the mission requirements of each individual DSRC?

Offerors may propose solutions that are common across the DSRCs or substantially different. Offerors are encouraged to provide rationale for their solutions being common or different.

What types of CPU workloads will the new systems run at AFRL, Navy, and MHPCC? Will you provide representative applications or benchmarks so vendors can choose and size the right processors?

CPU workloads and applications are anticipated to be similar to the ones running on current HPC systems. In general, the new systems will support a wide variety of modeling and simulation workloads. For more information on current workloads, please see: <https://hpc.mil/solution-areas/technology-areas>. Consideration should be taken that some modeling and simulation applications are able to leverage accelerators configured to support double precision calculations.

Please note: This CSO announcement initiates a multi-step process. The first step is the solicitation of solution briefs. The purpose of the solution brief submission is to identify potential partners that may have promising solutions. Subsequent steps consist of opportunities for selected offerors to deliver solution pitches, demonstrations, and collaborative exchanges with the Government. As such, offerors do not need to submit initial solutions with detailed technical specifications trying to satisfy technical requirements that will be identified in later steps.

MILCON / Site Preparation Funding: The CSO identifies an estimated Government funding profile of approximately \$57–62M per year for TI system purchases and provides facility constraints and requirements for the individual DSRC locations. Should offerors assume that any required MILCON, facility modifications, or site-preparation activities are funded separately from the TI-27 system acquisition budget, or should costs associated with required facility modifications be included within the proposed solution and ROM?

MILCON, facility modifications, and site preparations are not included in the TI-27 system acquisition budget.

Appendix A – Potential Solution Scenarios, AFRL (p. 11): AFRL CPU/GPU Workload Growth: Appendix A identifies three potential AFRL architectures ranging from a largely CPU-based system with limited GPU capability, to a moderately CPU/GPU-based system, to a primarily GPU-based AI/ML system. Over the anticipated 5–7 year system lifecycle, does AFRL anticipate that demand for CPU-based traditional HPC workloads or GPU-based AI/ML workloads will grow at a materially greater rate than the other, and should offerors account for that anticipated workload growth when proposing the initial architecture and future scalability of the AFRL system?

AFRL is not providing a definitive forecast that either CPU-based traditional HPC workloads or GPU-based AI/ML workloads will grow at a materially greater rate than the other over the expected life of the system. Offerors should therefore base their proposed solutions on the solicitation requirements and the potential solution scenarios in Appendix A, rather than assume the Government is signaling a preferred future workload mix.

At the same time, offerors should recognize that this acquisition is generally expected to function as a one-time system buy, not a system that will likely be substantially expanded after installation. As a result, AFRL expects offerors to consider likely workload evolution over the

system's useful life when proposing the initial architecture, including the balance of CPU, GPU, memory, storage, interconnect, and scalability characteristics.

In practical terms, AFRL is looking for offerors to make a thoughtful upfront design choice that can support the expected mission space over time, rather than assume the system can later be materially rebalanced through follow-on expansion.

Can vendors propose different CPU configurations for different workloads or sites, rather than using the same processor and node design throughout?

Yes, offerors may propose different configurations in their overall solutions.

For systems built primarily for AI and machine learning, what work do you expect the CPUs to handle alongside the accelerators?

CPUs should be configured to match requirements for the AI/ML accelerators for orchestration, preprocessing, I/O or other tasks.

How closely must each small test system match its production system? Must it use the same processors, memory, and software, or can it be a smaller version with some differences?

Each small test system's components shall match those in its production system as closely as possible for the purpose of comparative testing.

After an award, when does the Government expect each system to be delivered to its site? What are the expected dates for installation, testing, and final acceptance? Does the requirement to deliver within one calendar year refer to arrival at the site or final acceptance?

Expected dates will be determined during the award process. Final acceptance within one year or less from award is a goal.

Will vendors have an opportunity to test and tune applications before submitting a final proposal? If so, how will you evaluate those test results?

The Government has not determined whether applications or benchmarks will be made available and part of any subsequent evaluations. Benchmark applications are generally available for acceptance testing after award. Offerors may test and tune any available applications prior to submitting a final proposal.

Does the government desire a combined solution brief cover all DSRCs/scenarios within the single 10-page limit, or separate solution briefs per DSRC?

The offeror shall determine the best approach to their solution brief.

Do the acronym page, cover/title page, and diagrams count toward the 10-page limit?

The cover/title page does not count toward the 10-page limit. Other pages are to be included within the 10-page limit.

Would you prefer the ROM be broken out by DSRC? Any further requirements on ROM detail?

Yes, please separate ROMs per DSRC. The offeror shall determine ROM details to include.

The FAQ notes that some modeling and simulation applications use double-precision accelerators. Will FP64 performance be an evaluation factor for accelerator selection?

FP64 performance will be considered for accelerator selection. A mixture of accelerators to match the workloads is an option that would be considered. Please note any tradeoffs/benefits in performance between memory bandwidth vs FP64 performance.

Are there any per-rack Power Density Limits? (each facility)

AFRL confirms that the distribution panels at both AFRL buildings support a maximum breaker size of 800A. Offerors may propose rack configurations up to that limit, provided the proposed rack electrical configuration is compatible with the stated site electrical constraints and all other solicitation requirements.

MHPCC does not have a per-rack limit, but power A and B-side power feeds are limited by the 400A RPP circuits.

The Navy DSRC does not have a per-rack kW limit, however, each individual circuit to the proposed solution cannot exceed 250A at 480V based on the size of the available RPP (Remote Power Panel) type available.

Will there be any required connections to public cloud / remote resources not described in the documentation? What networks are planned to provide these connections?

At this time there are no additional connections required other than described in the announcement.

Please describe anticipated use-cases and anticipated workflows for each system described in Appendix A. The FAQ was helpful in determining broad use cases, but understand if there is a subset of these use-cases anticipated by machine in Appendix A.

AFRL response: As stated previously, CPU workloads and applications are anticipated to be similar to the ones running on current HPC systems. In general, the new systems will support a wide variety of modeling and simulation workloads. For more information on current workloads,

please see: <https://hpc.mil/solution-areas/technology-areas>. Consideration should be taken that some modeling and simulation applications are able to leverage accelerators configured to support double precision calculations.

MHPCC's AI/ML focused use cases will vary by user, but current priority is to support both low-latency, high-throughput Inference and deep generative and spatial AI training.

Workloads at the Navy DSRC follow mission requirements and evolving user community HPC practices. The Navy supports RDT&E mission requirements and operational Navy meteorological and oceanographic prediction needs. The initial workloads will likely be similar to existing production workloads. (c.f. <https://hpc.mil/solution-areas/technology-areas>) Workloads will evolve with the user community over the system life. Current modelling and simulation workflows (often at FP64) are anticipated as well as continuation of the current high and growing number of AI/ML workflows.

Per the FAQ, is there any prioritization of the described use cases? Please break down anticipated for each site and anticipated machine.

There is generally no overall prioritization of use cases.

The Navy DSRC responds to existing and evolving mission needs. The goal is to deploy hardware resources and support staff to provide excellent user capability for floating point intensive modelling and simulation, machine learning intensive workflow, data intensive computing, highly scalable computing, and high throughput computing as needed by the allocated communities. In addition, a portion of the available systems are dedicated to timely simulations and data analyses to support production of meteorological and oceanographic products useful for operational Navy needs.

For AI workloads, what mix of high precision and low precision workloads are expected?

The mixture of high precision and low precision will vary. This could include surrogate models for modeling and simulation as well as LLM or other low resolution workloads.

Does the Government require partitioning of GPU resources?

GPU partitioning or sharing is an option that would be considered to facilitate running multiple workloads on one node.

The Appendix A table can be interpreted that HPCMP would like 3 systems (one of the three possible AFRL systems) or 5 systems, plus their associated TDS systems. Are you seeking 3 or 5 production systems total?

The Government plans to select a total of 3 systems – a single system for each of the AFRL, MHPCC, and Navy DSRCs. However, at Step 4 the Government may decide to acquire additional systems for one or more DSRCs.

For the TDS requirement, you state the larger systems require a TDS, does that mean the Hawaii system does NOT require a TDS?

• No TDS space is described in Appendix C.

The MHPCC DSRC will require a TDS.

What are the number of files expected to be (file count) for each systems primary filesystem?

For AFRL between 2,500,000,000 and 3,000,000,000.

For MHPCC between 10,000,000 and 100,000,000

For Navy, system inode counts are driven by user needs. Current systems have on the order of 2 billion files per parallel file system. Growth over time is anticipated. There is a trend of increasing file count and decreasing average file sizes for user groups continue moving toward more data-intensive AI/ML workloads.

What classification level will the AFRL and Navy TDS operate at, given the FAQ's expectation of production parity and that the AFRL TDS space is unclassified?

The TDS typically operates at the unclassified level. This may change at Step 4.

What is the largest file size anticipated to be used by the primary filesystem?

Offerors should assume the primary file system must support the range of file sizes reasonably associated with HPC and AI/ML workloads over the system's useful life, including large science, engineering, training, checkpoint, and output files, and should propose a solution that effectively supports that workload range. In keeping with the solicitation's intent to encourage effective and innovative solutions, the Government will evaluate the proposed storage approach based on compliance with solicitation requirements, demonstrated performance, and overall solution effectiveness, rather than on any unstated single file-size assumption.

What is the average file size anticipated to be used by the primary filesystem?

See answer above.

What is the maximum size of inference model to be used? Concurrent users?

The size of the inference and training models can vary. The models could span more than one typically deployed dense accelerator node and the interconnect between the nodes should reflect this scenario. Concurrent users could vary, but like the model size would require the ability to scale beyond one typical dense accelerator node.

What is the maximum size of training model to be used? Concurrent users?

See answer above.

Will solutions built on standards-based Ethernet fabrics be evaluated equivalently to those using proprietary or vendor-specific interconnects?

The Government is open to innovative technical approaches, including standards-based Ethernet fabric solutions, provided the proposed solution satisfies all applicable solicitation requirements and demonstrates the required performance. The Government will evaluate solutions based on compliance with requirements, demonstrated performance, and effectiveness, rather than on whether the interconnect is standards-based Ethernet or proprietary/vendor-specific.

Will the government evaluate a single standardized architecture proposed across all three DSRCs as a package, or evaluate each site independently?

The Government will evaluate the solutions offered. The offeror may submit solutions that contain all three DSRCs or separate solutions with one for each DSRC or any combination. Each submittal will be evaluated on its own merit. The Government may also select only a portion of a solution.

How is system availability measured? Are scheduled maintenance windows and government-caused outages excluded?

System availability generally excludes approved scheduled maintenance and Government-caused events.

Government states must provide scratch space storage. How much capacity should be allocated for this per site?

The Government does not prescribe a separate minimum scratch-space capacity allocation per site beyond the solicitation requirements. Because the proposed system is expected to serve the Government's workload needs primarily through the initial delivered configuration, offerors should propose and justify a scratch-space allocation at each site that appropriately supports the expected workload mix over the useful life of the system.

Please note: This CSO announcement initiates a multi-step process. The first step is the solicitation of solution briefs. The purpose of the solution brief submission is to identify potential partners that may have promising solutions. Subsequent steps consist of opportunities for selected offerors to deliver solution pitches, demonstrations, and collaborative exchanges with the Government. As such, offerors do not need to submit initial solutions with detailed technical specifications trying to satisfy technical requirements that will be identified in later steps.

Should proposals include additional high-performance storage and if so at what usable capacity?

The Government does not establish a separate requirement for additional high-performance storage or a separate usable-capacity threshold beyond the solicitation requirements. Offerors shall include whatever storage is necessary for their proposed solution to satisfy the stated requirements and required performance. Solutions will be evaluated based on compliance with the solicitation and the effectiveness of the proposed solution.

Please note: This CSO announcement initiates a multi-step process. The first step is the solicitation of solution briefs. The purpose of the solution brief submission is to identify potential partners that may have promising solutions. Subsequent steps consist of opportunities for selected offerors to deliver solution pitches, demonstrations, and collaborative exchanges with the Government. As such, offerors do not need to submit initial solutions with detailed technical specifications trying to satisfy technical requirements that will be identified in later steps.

For the distinct sites, does the gov prefer one logical fabric of resources and a common control plane with jobs placed automatically or: Each site stays operationally distinct with unified observability?

The Government does not state a preference for either a single logical fabric with a common control plane or for operationally distinct sites with unified observability, except as otherwise provided in the solicitation. In keeping with the solicitation's intent to encourage effective and innovative solutions, offerors may propose either approach, provided the proposed architecture satisfies all applicable requirements for functionality, performance, security, administration, resiliency, and usability. Offerors should clearly describe how the proposed solution will operate across the distinct sites, including job placement, management, observability, fault isolation, any operational tradeoffs, and ability to interface across separate security enclaves.

Is there an expectation of any cloud resources that may need to be managed hybrid with the on prem HPC cluster?

There is no expectation of any cloud resources that may need to be managed in a hybrid manner.

Preference to CUDA and/or ROCm for AI/ML workloads?

The Government does not establish a formal preference for CUDA- or ROCm-based solutions except as otherwise stated in the solicitation. Offerors should be aware, however, that the Government's current AI/ML workload environment is predominantly CUDA-based. In keeping with the solicitation's intent to encourage effective and innovative solutions, offerors may propose any compliant approach but should clearly address how the proposed solution will support the existing workload environments, migration considerations, software compatibility, and overall mission effectiveness.

Is there a required or preferred minimum GPU generation/interconnect (e.g., NVLink/InfiniBand vs. RoCE Ethernet)?

The Government does not specify a required or preferred minimum GPU generation or a required or preferred interconnect technology (e.g., NVLink, InfiniBand, or RoCE Ethernet) beyond the requirements expressly stated in the solicitation. Consistent with the solicitation's intent to encourage effective and innovative solutions, offerors may propose any compliant GPU and interconnect approach that satisfies the applicable functional, performance, scalability, reliability, and benchmark requirements.

Are there GPU memory bandwidth constraints or minimum desired capacity/speed?

The Government does not specify any additional required or preferred GPU memory bandwidth, memory capacity, or memory speed beyond the requirements expressly stated in the solicitation. Consistent with the solicitation's intent to encourage effective and innovative solutions, offerors may propose any compliant approach that satisfies the applicable functional, performance, scalability, reliability, and benchmark requirements.

For all sites: who is responsible for CDU/liquid-cooling loop qualification if the offeror's rack-level cooling hardware interfaces with government-owned facility CDUs – is there a liability/warranty boundary the Government has already defined?

AFRL response: There are no Government-owned facility CDUs. CDUs are expected as part of the proposed solution. The Government's liability/warranty boundary is at the Government-provided chilled-water interface.

MHPCC response: The Government will have liability/warranty of the CDU secondary loop. Vendor will have liability for rack level and interconnects to the secondary loop.

Navy response: There are no government-owned facility CDUs. CDUs are expected as part of the proposed solution. Liability/warranty boundary is at the Government-provided chilled-water interface.

For each AFRL scenario in Appendix A, which building will house the system: the first building (2 MW, unclassified) or the second building (550 kW, classified)? Should offerors propose a single AFRL scenario or provide pricing for each AFRL alternative?

The AFRL DSRC's first building is generally used to house unclassified systems, and the second building is used to house classified systems. Offerors should provide pricing for each AFRL alternative.

The government states:

"The Government budgets approximately \$57-62M per year for all TI system purchases. The estimate for 2027 will include equal amounts for the AFRL and Navy DSRCs with

approximately 8% of the total budget for the system at the MHPCC DSRC The HPCMP budgets 6% of the systems' original purchase price per year for up to 5 years of system administration and maintenance of purchased hardware and software. Beyond 5 years, annual costs usually increase and budgets are adjusted as necessary."

The MHPCC share is described as less than 10% of total investment in Requirements and approximately 8% in the Funding Profile. Which figure should the ROM target?

The ROM should target approximately 8% of the total budget for the MHPCC solution.

Is 6% of systems original purchase price budgeted for maintenance is a separate budget from that traditionally allocated \$57m-\$62m or if that larger number is inclusive of the 6%?

The annual budget for systems administration and maintenance support is separate from the total acquisition budget.

The floor-space figures differ between the summary tables and the section text: AFRL unclassified (1,400 vs. 1,440 ft²), AFRL classified (320 vs. 432 ft²), and MHPCC (224 vs. 244 ft²). Which figures should offerors use for layout?

For AFRL, 1400 ft² is available in the unclassified area and 320 ft² is available in the classified area.

For MHPCC, the correct floor space availability is 224 ft².

For the AFRL TDS space, the summary table indicates no UPS/generator backup, while §3.3 references a 3 MW UPS/generator system. The table also lists a 45°F chilled-water loop, while §3.2 describes air cooling only. Please confirm whether TDS power is UPS-backed and whether facility water is available at the TDS location. Given the 600 lb/tile and 150 lb/ft² floor ratings, can the Government accommodate a TDS rack matching the weight and cooling method of the production system, as the FAQ indicates TDS parity is generally expected?

For AFRL, the area hosting the TDS is generator backed up. Facility water is available at the TDS location. The TDS must stay within the 600 lb/tile and 150 lb/ft² floor ratings. Offerors should not assume that the TDS can necessarily accommodate a rack identical in weight, utility demand, or cooling method to the production system unless that capability is expressly confirmed by the solicitation. For purposes of this acquisition, TDS parity is intended to support representative test and development functionality and does not, by itself, guarantee full physical rack parity with the production system.

For each site, please provide:

- **maximum allowable return temperature and allowable ΔT**
- **connection sizes and types**
- **water quality specifications**

- **seasonal variation in supply temperature**

The site facility specifics will be provided at Step 4.

Are rear-door heat exchangers acceptable for achieving room-temperature neutrality? Does the requirement apply to storage and network racks as well as compute?

For AFRL rear-door heat exchangers are acceptable. This requirement applies to all racks.

For MHPCC RDHx solutions are acceptable. The Government preference is for liquid cooling of the whole solution, but air cooling is available for a limited thermal load.

For Navy rear door heat exchangers are acceptable. This requirement applies to all racks.

Is the full available power UPS-backed, or should offerors assume UPS coverage for critical infrastructure only? Is there a maximum per-rack power density? Each site.

For all sites, full available power is backed up by both UPS and Generator. The Navy DSRC does not have a per-rack kW limit, however, each individual circuit to the proposed solution cannot exceed 250A at 480V based on the size of the available RPP (Remote Power Panel) type.

How will you compare the performance of proposed systems? For example, will you measure how quickly they complete a workload, how many jobs they can run, power use, cost, or some combination?

Performance metrics generally include all those factors.

Is the offeror expected to provide the complete network architecture supporting the proposed HPC and AI/ML systems, including the high-speed compute/AI fabric, storage networking, management and out-of-band networking, and production/user-facing connectivity, or only the fabric internal to the system, with the remaining networks treated as Government-furnished equipment?

Proposed solutions should include fabric internal to the system and should identify other network architecture components needed for integration with external systems such as, but not limited to, storage systems. Government-furnished equipment will be identified during scope collaboration with the selected offeror(s).

Please note: This CSO announcement initiates a multi-step process. The first step is the solicitation of solution briefs. The purpose of the solution brief submission is to identify potential partners that may have promising solutions. Subsequent steps consist of opportunities for selected offerors to deliver solution pitches, demonstrations, and collaborative exchanges with the Government. As such, offerors do not need to submit initial solutions with detailed technical specifications trying to satisfy technical requirements that will be identified in later steps.

Appendix B describes the AFRL DSRC network as a 100/40/10/1 Gbps LAN with native IPv6/dual-stack support, a 100 Gbps DREN connection, an eight-switch core spanning two buildings, and a separate 10 Gbps private storage network requiring 10GBASE-SR on any node mounting that storage, while Appendix D describes Navy DSRC as a 40 Gbps Juniper/Arista backbone with a minimum of 10 Gbps user-facing and 1 Gbps administrative interfaces. Should offerors assume these center-wide backbones remain unchanged for the 5–7 year lifecycle of the TI-27 systems, and what specific interface speeds, media types, and port counts does the Government expect the new systems' login, data-transfer, and management nodes to present to each network?

AFRL response: Offerors can assume that the center-wide backbone will remain unchanged for this period.

Navy response: This level of detail can be explored in later detailed discussions. Current DSRC networking configuration can be assumed. Current systems have 100 Gbps and 40 Gbps connections to the login nodes, transfer nodes, and visualization nodes. High density BiDi transceivers are commonly used.

Should the proposed architecture provide network and fabric headroom, spare spine/leaf capacity, ports, and optics for expansion of compute and GPU resources across the anticipated 5–7 year lifecycle, and if so, should offerors state a target expansion percentage or design to a specific Government-provided figure?

AFRL response: The Government does not establish a required expansion percentage or a specific Government-provided future-growth figure for network or fabric headroom beyond the requirements expressly stated in the solicitation. Offerors should recognize that the systems are generally expected to function as one-time buys and should therefore propose an initial architecture that appropriately supports the expected workload needs over the useful life of the system.

Offerors may include network or fabric headroom, spare spine/leaf capacity, ports, optics, or other design features that they determine are necessary to provide an effective and innovative solution, but proposals should clearly explain the rationale, assumptions, and tradeoffs associated with that approach.

MHPCC response: Yes, the government prefers support for future expansion of GPU resources during the solution lifecycle. Specific expansion targets can be explored following solution selection.

Navy response: Pre-planned expansion paths are not required. To the extent that such provisioning reduces maximum system performance it is not preferred.

Please note: This CSO announcement initiates a multi-step process. The first step is the solicitation of solution briefs. The purpose of the solution brief submission is to identify potential partners that may have promising solutions. Subsequent steps consist of opportunities for selected offerors to deliver solution pitches, demonstrations, and collaborative exchanges with the Government. As such, offerors do not need to submit initial solutions with detailed technical specifications trying to satisfy technical requirements that will be identified in later steps.

Are there specific IPv6 or dual-stack requirements for the compute/AI fabric, storage network, management network, and production-facing interfaces of the new systems, consistent with the DoD IPv6 transition direction and the native-IPv6 LAN that Appendix B describes for AFRL?

The Government requires IPv6 dual-stack support only for production-facing nodes/interfaces, consistent with the solicitation's stated site network environment. The Government does not otherwise require IPv6 dual-stack support for the internal compute/AI fabric, storage network, or management network unless expressly stated elsewhere in the solicitation. Offerors shall clearly describe the proposed IPv4/IPv6/dual-stack posture for all production-facing interfaces and any assumptions regarding integration with Government-furnished network environments

Does the Government have a preference or constraint regarding the high-speed interconnect technology for the CPU-compute and AI/ML partitions, for example InfiniBand, HPE Slingshot, or lossless Ethernet with RoCE, or should offerors select the fabric on the basis of measured MPI latency, collective performance, and tail-latency behavior for the anticipated workload mix and present that evidence?

The Government does not specify a required or preferred high-speed interconnect technology (e.g., InfiniBand, Slingshot, or RoCE Ethernet) beyond the requirements expressly stated in the solicitation. Consistent with the solicitation's intent to encourage effective and innovative solutions, offerors may propose any compliant interconnect approach that satisfies the applicable functional, performance, scalability, reliability, and benchmark requirements.

Does the Government expect physically separate fabrics for compute/AI, storage, management, and production/user traffic, or may these functions be logically segmented over shared infrastructure (for example, traffic classes and congestion control on a converged fabric) where offerors demonstrate that performance isolation and security separation are maintained under concurrent MPI and distributed-training load?

Offerors may propose solutions consisting of shared infrastructure that will satisfy applicable functional, performance, scalability, reliability, and benchmark requirements.

Are there specific requirements for fabric and network redundancy, including redundant switches, planes, links, paths, supervisors, and power supplies, and does the Government have an availability or fault-tolerance target the fabric design should meet, such as ensuring no single switch failure reduces the schedulable partition by more than a stated fraction?

The Government considers non-blocking and/or rail-optimized fabric designs to be desirable, particularly where they improve workload performance, resiliency, and operational effectiveness. However, offerors may propose other compliant approaches, provided they clearly demonstrate how the proposed design satisfies the solicitation's requirements and achieves comparable solution effectiveness.

Please note: This CSO announcement initiates a multi-step process. The first step is the solicitation of solution briefs. The purpose of the solution brief submission is to identify potential partners that may have promising solutions. Subsequent steps consist of opportunities for selected offerors to deliver solution pitches, demonstrations, and collaborative exchanges with the Government. As such, offerors do not need to submit initial solutions with detailed technical specifications trying to satisfy technical requirements that will be identified in later steps.

Are there requirements for fabric and network telemetry, automation, configuration management, and performance monitoring across the TI-27 systems, such as per-port congestion and error counters exposed to the scheduler and operations tooling, and should these integrate with existing DSRC network-management platforms and interfaces?

The Government does not establish, by this response, a separate mandatory requirement that proposed telemetry or monitoring capabilities integrate with any specific existing DSRC platform or interface beyond the solicitation requirements. However, offerors should clearly describe how the proposed capabilities will support Government operations, observability, troubleshooting, and performance management in the target environment.

Please note: This CSO announcement initiates a multi-step process. The first step is the solicitation of solution briefs. The purpose of the solution brief submission is to identify potential partners that may have promising solutions. Subsequent steps consist of opportunities for selected offerors to deliver solution pitches, demonstrations, and collaborative exchanges with the Government. As such, offerors do not need to submit initial solutions with detailed technical specifications trying to satisfy technical requirements that will be identified in later steps.

For the data-transfer nodes the solicitation asks offerors to include, what external-facing bandwidth and protocols does the Government expect toward DREN and toward HPCMP enterprise storage, such as 100 Gbps per node with parallel transfer tooling, and are there existing DSRC data-transfer conventions or tools the new nodes must interoperate with?

The current solicitation does not establish additional mandatory external-facing bandwidth, protocol, or transfer-tool requirements for Data Transfer Nodes beyond the requirements and site environment information expressly provided. Offerors should recognize that the Government environment includes production-facing network connectivity and, at AFRL for example, a separate 10 Gbps private storage network for nodes that mount enterprise storage. Offerors should describe the data-transfer-node bandwidth, interface approach, protocols, tooling, and interoperability assumptions necessary to support an effective and operable solution in the Government environment.

The solicitation identifies “high-performance file systems needed for temporary workspace.” Is this intended to be HPC scratch only, or should the proposed solution also include persistent project storage, home directories, or a capacity tier, and if so, with what retention expectations?

Offerors should base their proposed storage architecture on the requirements expressly stated in the solicitation. The Government does not, by this response, prescribe a specific storage-tier model beyond those requirements, including whether a proposed solution should include only high-performance temporary workspace or also separate persistent project, home-directory, or capacity-tier storage.

Offerors may propose the storage architecture they determine best satisfies the solicitation requirements and the Government's operational needs, provided they clearly describe the proposed storage tiers, intended use, retention assumptions, performance characteristics, and associated operational tradeoffs.

Will the Government provide target figures for usable capacity, aggregate bandwidth, metadata operation rates, or small-file/IOPS performance for the file system, or should offerors size storage in proportion to the compute partition, such as a ratio of aggregate bandwidth to system memory, and state the basis?

Offerors should provide optimized solutions to match the expected workload based on the architectures being considered as listed in Appendix A. Details of the file system implementations will be determined during scope collaboration with the selected offeror.

Please note: This CSO announcement initiates a multi-step process. The first step is the solicitation of solution briefs. The purpose of the solution brief submission is to identify potential partners that may have promising solutions. Subsequent steps consist of opportunities for selected offerors to deliver solution pitches, demonstrations, and collaborative exchanges with the Government. As such, offerors do not need to submit initial solutions with detailed technical specifications trying to satisfy technical requirements that will be identified in later steps.

Is the Government open to any parallel or distributed file system architecture, for example Lustre-based, IBM Storage Scale, or NVMe-native disaggregated designs such as Weka or VAST, provided the offeror demonstrates POSIX-parallel semantics, metadata scalability, and encryption at rest to the solicitation's requirements, or are there file-system technologies that must be maintained at a given DSRC for interoperability with existing systems?

The Government does not require offerors to maintain any specific incumbent file-system technology at a given DSRC unless expressly stated elsewhere in the solicitation. Offerors should not assume that interoperability with existing systems requires retention of a specific incumbent file-system product unless such a requirement is expressly identified in the solicitation. Offerors may propose any compliant parallel or distributed file-system architecture, including non-Lustre approaches, provided the proposed solution satisfies the applicable solicitation requirements for functionality, performance, scalability, security, interoperability, and supportability.

Is the new system's file system expected to integrate with existing DSRC center-wide storage, such as the 10 Gbps private storage network at AFRL, and if so, what protocols (NFS, parallel-

file-system client, S3), mount conventions, and performance expectations should offerors assume?

The Government does not require the proposed system's primary file system to replace, extend, or become part of the existing DSRC center-wide storage architecture. Rather, where identified in the solicitation, the Government expects the proposed system to support connectivity to Government-furnished center-wide storage for data access and transfer purposes. Offerors should not assume that a common or unified file-system implementation with existing DSRC storage is required unless expressly stated elsewhere in the solicitation. Offerors shall clearly describe the proposed method for enabling connectivity, access, and data transfer between the proposed system and Government-furnished storage environments, including any assumptions regarding protocols, mounts, clients, and supporting software.

Is there an expected data lifecycle between HPC scratch, persistent project storage, the HPCMP Enterprise Data Fabric, and archival storage, such as automated purge policies on scratch, tiering or staging tools, and transfer paths to enterprise storage, that the proposed system should support?

The Government does not, in this phase, prescribe a single required data-lifecycle architecture or a mandatory set of automated purge, tiering, staging, or archival mechanisms beyond the information expressly provided in the solicitation. However, offerors should recognize that the proposed system will need to operate effectively within a broader HPC data environment that may include temporary high-performance workspace, persistent storage environments, enterprise data-management services, and archival workflows. Offerors should not assume that the proposed primary file system itself must become part of each Government storage environment, but should explain how the proposed solution would support required access, connectivity, and data transfer across those environments.

Offerors should describe any assumptions regarding temporary workspace behavior, retention, purge or cleanup mechanisms, staging/tiering approaches, and transfer paths to Government-furnished storage or enterprise data environments.

Which cybersecurity capabilities are expected to be included in the proposed system versus provided as Government-furnished enterprise or DSRC security services, such as boundary firewalls, IDS/IPS, network access control, log forwarding to a Government SIEM, vulnerability scanning infrastructure (ACAS), and host-based security (HBSS/ESS), and can the Government identify those capabilities where integration is required?

For the unclassified option, AFRL will use HPCMP Enterprise tooling. For the classified options, AFRL will provide the security stack for integration.

Please note: This CSO announcement initiates a multi-step process. The first step is the solicitation of solution briefs. The purpose of the solution brief submission is to identify potential partners that may have promising solutions. Subsequent steps consist of opportunities for selected offerors to deliver solution pitches, demonstrations, and collaborative exchanges with the Government. As such, offerors do not need to submit initial solutions with

detailed technical specifications trying to satisfy technical requirements that will be identified in later steps.

Is the offeror expected to provide encryption for data in transit within and between the new systems' networks, should offerors assume Government-furnished cryptographic devices such as HAIPE/Type 1 encryptors for classified WAN paths or commercial encryption within the proposed solution, and what approval requirements apply to commercial encryption in transit, such as FIPS 140-3 validation, CNSA algorithm compliance, or NSA CSfC?

Yes, assume Government-furnished cryptographic devices for classified WAN paths.

The solicitation requires full-disk encryption of all provided non-volatile storage using a cipher approved for national security information on classified systems and on unclassified systems requiring additional data protection. Does this requirement extend to every non-volatile device in the system, including node boot devices, local NVMe, parallel-file-system media, and management/BMC storage, and does the Government accept self-encrypting drives operating under a FIPS-mode host operating system whose kernel cryptographic module holds an active CMVP validation as a compliant implementation?

AFRL requires that all non-volatile storage provided as part of the proposed system be encrypted in accordance with the solicitation requirements. The Government also requires the use of FIPS 140-3 validated cryptographic modules where required by the solicitation.

The Government does not, by this response, pre-approve any specific implementation model—such as self-encrypting drives, host-based encryption, or a combined hardware/software approach—solely based on the architecture description. Offerors proposing any such approach shall clearly demonstrate how the proposed implementation satisfies the solicitation's encryption requirements across all applicable non-volatile storage, including how encryption is implemented, managed, enforced, and validated.

Who is expected to provide and manage encryption keys for data at rest, the system provider via an OEM-internal or file-system key manager, the DSRC, HPCMP enterprise services, or a Government-furnished key-management/HSM capability, and for classified systems, is an OEM-internal key manager acceptable to the ISSM or is an external, Government-controlled key manager required?

The Government does not prescribe a single required key-management architecture or assign a specific provider for encryption-key generation, custody, or management beyond the requirements expressly stated in the solicitation. Offerors should recognize that key management for data-at-rest protection is a significant operational and security consideration and should clearly describe the proposed key-management approach, including key generation, storage, protection, access control, recovery, rotation, revocation, auditability, and any reliance on OEM-internal, file-system-native, enterprise, or Government-furnished capabilities.

For classified-system use cases, offerors should not assume that any specific OEM-internal key-management approach is automatically acceptable. Proposals should clearly identify all assumptions and dependencies associated with the proposed key-management model.

What level of RMF support is expected from the system provider, such as STIG implementation, SCAP/ACAS scan results in eMASS-ready format, NIST SP 800-53 control-implementation narratives, POA&M development for HPC components that lack a distribution-specific STIG (parallel file system, fabric manager, GPU driver stack, BMC firmware), and support through authorization, and are TI-27 systems expected to meet CNSA 2.0 quantum-resistant cryptography requirements at initial acceptance or on a schedule within the 5–7 year lifecycle?

Offerors should assume that the system provider is expected to support the Government's RMF and cybersecurity accreditation activities to the extent necessary for the proposed system to achieve initial authorization and to operate in accordance with the solicitation requirements. This support should include, as applicable, implementation of required security configurations, remediation of scan results and supporting artifacts, development of corrective-action information, contribution to control implementation documentation, and support for Government authorization activities.

If administration/system support options are exercised, offerors should assume that this support is expected to continue through the end of the exercised option periods, consistent with the solicitation requirements.

With respect to specific artifacts, offerors should clearly describe the proposed approach for supporting STIG implementation, scan and assessment outputs, control documentation, POA&M-related support, and treatment of components for which no directly applicable STIG exists.

The solicitation requires a small Test and Development System aligned with the architecture of each of the largest production systems for testing updates before production. Is the intent that the TDS reproduce the complete production stack at reduced scale, including compute and GPU node types, high-speed fabric, a representative slice of the parallel file system, management plane, and security configuration, so that firmware, fabric, file-system, scheduler, and STIG changes can be validated end to end before production rollout?

In order to properly support testing of production system updates, it's generally expected that the TDS will match in all of the above-mentioned areas. Exceptions can be made in some instances. Final determinations should be made during scope collaboration with the selected offeror(s).

How closely must the TDS match production in terms of identical CPU/GPU generations, fabric ASICs/firmware, storage controllers, and software versions versus functional equivalence at the processor, accelerator, and interconnect level, and where TDS space provides air cooling

only (Appendix B, Section 3), may the TDS cooling topology differ from a liquid-cooled production system?

See answer above. The TDS cooling options may be different than the liquid-cooled production system and TDS cooling specifications will be determined during scope collaboration with the selected offeror(s).

Is the TDS expected to reproduce production network segmentation and connectivity to DSRC enterprise services (authentication, logging, scanning, Enterprise Data Fabric) so that integration changes can be tested, is virtualization or digital-twin modeling acceptable for management/scheduler testing, and is a TDS required for MHPCC given its omission from Appendix C?

AFRL response: AFRL does not prescribe a single mandatory method for achieving TDS parity across network connectivity, segmentation, enterprise-service interaction, management, scheduler testing, or integration validation beyond the information expressly provided in the solicitation.

Offerors should propose the TDS architecture and supporting approach they determine are necessary to provide an effective and supportable test and development capability consistent with the required system parity and operational objectives of the solicitation. This includes clearly describing any assumptions regarding connectivity to Government-furnished enterprise services, the use of representative versus full-fidelity integration environments, and whether any virtualization, emulation, modeling, or digital-twin techniques are used as part of the proposed solution.

Navy response: Yes. The TDS provides a platform to test proposed production system configuration changes. To be effective, it should faithfully replicate the production system in all ways except scale. Virtualization approaches can be considered and will be evaluated as to their ability to actually exercise required functionality and provide extrapolatable performance predictions.

MHPCC response: In addition to the above responses from the other DSRC's, MHPCC requires TDS capability to support the solution. This requirement is dependent on the offeror's innovative solution across the three DSRCs. For example, if the offeror proposes a common architecture and support model between MHPCC and another DSRC, then MHPCC would support a shared TDS supporting both centers.

Appendix A describes the desired systems by CPU/GPU balance but gives no performance targets, node counts, or memory and storage capacities. Should offerors size the proposed systems to the funding profile in Section A, and will the Government provide representative HPC and AI/ML workload characteristics, benchmark suites, model classes, precision requirements (FP64 vs. FP8/FP16/BF16), or reference applications used in prior Technology Insertion cycles?

Offerors should base their proposed solutions on the requirements, potential solution scenarios, and other information expressly provided in the solicitation. The Appendix A scenarios are

intended to communicate the range of architectural approaches the Government is willing to consider, including different CPU/GPU balances, rather than to establish a single prescribed system size or configuration in this phase.

Offerors may consider the available funding context in developing a practical and supportable solution; however, the Government is not, by this response, directing offerors to size their proposed systems to any unstated performance, node-count, memory-capacity, or storage-capacity target beyond the solicitation requirements. Solutions meeting requirements and committing to higher performance within the funding envelope will be evaluated more positively.

Please note: This CSO announcement initiates a multi-step process. The first step is the solicitation of solution briefs. The purpose of the solution brief submission is to identify potential partners that may have promising solutions. Subsequent steps consist of opportunities for selected offerors to deliver solution pitches, demonstrations, and collaborative exchanges with the Government. As such, offerors do not need to submit initial solutions with detailed technical specifications trying to satisfy technical requirements that will be identified in later steps.

Are there minimum requirements or preferences for CPU architecture, core count, memory capacity and bandwidth per node, GPU architecture and memory, or accelerator interconnect domain size (such as NVLink or Infinity Fabric), and for mixed systems, does the Government have a desired CPU-to-GPU resource ratio or preferred partitioning model between tightly coupled MPI work and loosely coupled AI training/inference?

AFRL response: Offerors should base their proposed solutions on the requirements, potential solution scenarios, and other information expressly provided in the solicitation. AFRL does not, in this phase, establish additional minimums or preferences for CPU architecture, core count, memory capacity, memory bandwidth per node, GPU architecture, GPU memory, accelerator-interconnect domain size, CPU-to-GPU ratio, or partitioning model beyond what is stated in the solicitation.

The Appendix A scenarios are intended to communicate the range of architectural approaches that AFRL is willing to consider, including differing CPU/GPU balances and workload emphases, rather than to prescribe a single Government-directed architecture.

Navy response: Offerors' solutions should seek to maximize system utility for allocated users. It is left to offerors to optimize architectures and provide innovation.

Please note: This CSO announcement initiates a multi-step process. The first step is the solicitation of solution briefs. The purpose of the solution brief submission is to identify potential partners that may have promising solutions. Subsequent steps consist of opportunities for selected offerors to deliver solution pitches, demonstrations, and collaborative exchanges with the Government. As such, offerors do not need to submit initial solutions with detailed technical specifications trying to satisfy technical requirements that will be identified in later steps.

Does the Government expect offerors to present effective-performance estimates (such as projected HPL/HPCG results, STREAM bandwidth, and MPI micro-benchmarks) in the solution brief, what minimum system-utilization or availability targets must be demonstrated during acceptance testing under concurrent load, and would the Government consider Government-witnessed pre-shipment source inspection under FAR 46.402 to shorten the on-site acceptance period?

The Government expects offerors to present effective-performance estimates and commit to achieving them. The Government will consider Government-witnessed pre-shipment source inspection.

Should proposed systems comply with HPCMP Baseline Configuration policies defining the common user environment across DSRCs, what is the preferred container runtime and orchestration model for HPC and AI work (unprivileged containers under batch vs. Kubernetes for persistent inference), and is the system provider expected to build and maintain the application stack via tools like Spack or provide platform-level support only?

Yes. Offerors should assume that the proposed system is expected to support compliance with applicable HPCMP Baseline Configuration Team (BCT) policies governing the common user environment across DSRCs. The applicable BCT policies also define the container solution to be supported. (<https://centers.hpc.mil/consolidated/bc/projects.html>)

Beyond that, the Government does not prescribe a single required division of responsibility between application-stack support and platform-level support beyond the requirements expressly stated in the solicitation. Offerors should describe the support model they propose and explain how that approach provides an effective and innovative solution for delivering, sustaining, and operating the required environment.

Should offerors assume integration with HPCMP Kerberos-based authentication and allocation-accounting systems, is the Government seeking a common management, monitoring, and orchestration architecture across AFRL, Navy, and MHPCC despite differing classification levels, and can the Government provide protocol specifications for the Enterprise Data Fabric referenced in Section A?

Offerors should assume that the proposed solution must be capable of integrating with the Government-furnished HPCMP Kerberos-based authentication environment and applicable allocation-accounting systems, as required for operation within the Government enterprise environment.

Beyond those identified enterprise integration points and any other requirements expressly stated in the solicitation, the Government does not prescribe a single mandatory common management, monitoring, or orchestration architecture across AFRL, MHPCC, and Navy DSRCs. Offerors should recognize that the systems will operate in differing site and classification environments and should clearly describe how the proposed approach addresses interoperability, administration, observability, and operational supportability across those environments.

Does the Government expect centralized monitoring and automation across compute, GPU, fabric, storage, power, and cooling telemetry with a unified operational view integrating with existing DSRC/HPCMP tools, and is predictive hardware-failure detection with proactive component replacement a desired capability?

The Government is interested in proposed solutions that are innovative and improve effectiveness, performance, and operability.

For the 24x7 support requirement and response SLAs (4-hour nominal / 2-hour Navy), does "response" mean initial acknowledgement, remote engagement, or on-site arrival, which administration and troubleshooting activities may be performed remotely by cleared personnel versus requiring physical presence, and which specific functions fall under contractor "system administration" ROM scope versus Government scope?

The response requirement is intended as acknowledgment of an identified issue. Specific response activities and responsibilities will be addressed during scope collaboration with the selected offeror(s).

Does the Government expect on-site spare hardware for compute, GPU, fabric, and storage components, are there target spare levels or replacement time objectives, and would the Government consider a program-level spares pool shared between the AFRL and Navy systems and managed by the enterprise maintainer under applicable property and classification rules?

The Government is interested in proposed solutions that are innovative and improve effectiveness, performance, affordability, and operability.

What portions of physical data-center integration fall within contractor scope (racks, PDUs, cabling, CDUs, secondary manifolds/hoses, structured cabling, rigging, physical installation), where is the demarcation to Government infrastructure, and should offerors design directly to Government-provided secondary loop conditions or include offeror-supplied CDUs where compute architectures require different supply parameters, specifically regarding the Navy DSRC dual-loop setup and MHPCC liquid-cooling limits relative to available power?

AFRL response: Racks, CDUs, and secondary manifolds/hoses, are provided and physically installed by the vendor/contractor. PDUs are provided by the government as well as the power whips/receptacles. The plugs are provided by the vendor/contractor. Vendor should include CDUs as part of their solution. Government provides primary chilled water supply/return to CDUs or individual racks as required. Please provide definition of structured cabling in this context.

Navy response: The demarcation for Government infrastructure is as follows - Electrical power whips with identified receptacle configurations including voltage and amperage will be provided under the raised floor for the offeror's provided rack level power solution to plug into. Primary chilled water will be provided to the system's cooling solution at a mutually agreed upon

location and size. It is the responsibility of the offeror to provide the final primary cooling connection hoses, and the secondary cooling system. It is also expected that the offeror design this system to accept Navy's chilled water. The preferred chilled water loop at the Navy DSRC operates at 50 – 53 deg F supply temperature.

MHPCC response: The offeror will be responsible for all power and cooling in the racks and lines connecting to the MHPCC RPP's and CDU secondary loop. If the offeror proposal includes an alternate CDU, the offeror responsibility will then extend to include that CDU and the secondary loop and associated manifolds.

Please confirm authoritative values where appendix figures conflict, including AFRL unclassified space (1,400 vs 1,440 sq ft), subfloor clearance (2 ft vs 52 in), AFRL classified space (320 vs 432 sq ft), MHPCC space (224 vs 244 sq ft), AFRL TDS power/loading specs, Navy 4,000-pound rack limits versus load-spreading plates, available distribution voltages/circuits per site, DCIM telemetry protocols, and how offerors should handle the <1 year delivery objective for MHPCC if award occurs prior to November 2027 facility occupancy.

AFRL response: Unclassified space 1,400 ft². The raised floor depth for this space is 24" and there is a 42" deep trench around the perimeter of the room for primary chilled water supply/return lines. AFRL classified space is 320ft². TDS specs are correct.

MHPCC response: The correct floor space availability is 224 ft². It is expected that final contract award will be less than one year from the MHPCC site availability. However, after selection, MHPCC and offeror can negotiate if temporary storage is required following factory integration and testing and prior to site readiness.

Navy response: 1300 sq ft is available for the large system, 100 sq ft is available for the test system. Load spreading plates and load distribution rack designs are allowed so that the 4,000-pound rack limit can be exceeded. The Navy DSRC utilizes ConCor 2000 floor tile, solutions cannot exceed the ratings of this raised floor type. Available electrical distribution voltages are 480/277V and 208/120V. There are ample available circuits at each available voltage type. BaCnet or ModBus are the preferred facility telemetry protocols.

For the TDS requirement, you state the larger systems require a TDS, does that mean the Hawaii system does NOT require a TDS?

No TDS space is described in Appendix C.

MHPCC requires TDS capability to support the solution. This requirement is dependent on the offeror's innovative solution across the three DSRCs. For example, if the offeror proposes a common architecture and support model between MHPCC and another DSRC, then MHPCC would support a shared TDS supporting both centers.

What classification level will the AFRL and Navy TDS operate at, given the FAQ's expectation of production parity and that the AFRL TDS space is unclassified?

The TDS typically operates at the unclassified level. This may change at Step 4.

The Navy DSRC has an 8.5' ceiling and prohibits interconnect cabling in the subfloor. What is the maximum rack height after sprinkler and cable-tray clearance? Is overhead cable tray existing or offeror-provided?

The proposed solution cannot exceed 96". This includes all racks, cable trays, and any other support infrastructure needed for the racks. Overhead cable tray is offeror-provided.

Please note these questions are in response to the Innovative Technology Insertion of Advanced HPC Systems CSO. There appears to be a system issue where this is no longer an option on the question topic dropdown.

Workload & Architecture: For the AI/ML portions of the AFRL, Navy, and MHPCC systems, what is the anticipated training-vs-inference workload mix, expected model/dataset scale, and required numerical precision (e.g., FP32/FP16/BF16/FP8/FP4)? (Ref: Sec. A, Requirements: Enterprise System Acquisition (describes CPU vs. GPU mix only qualitatively))

Appendix A, Potential Solution Scenarios (architecture rows have no precision/workload detail))

The accelerators that are part of the solution will likely have to be a mixture to meet the nature of various workloads. The numerical precision could be as high as FP64 down to quantized AI/ML models. Please see note below.

Please note: This CSO announcement initiates a multi-step process. The first step is the solicitation of solution briefs. The purpose of the solution brief submission is to identify potential partners that may have promising solutions. Subsequent steps consist of opportunities for selected offerors to deliver solution pitches, demonstrations, and collaborative exchanges with the Government. As such, offerors do not need to submit initial solutions with detailed technical specifications trying to satisfy technical requirements that will be identified in later steps.

Workload & Mission Fit: What is the GPU capacity primarily intended for — e.g., engineering/CFD simulation, weather/ocean modeling, ISR/sensor data processing, autonomy, or generative AI/LLM workloads — at each of the AFRL, Navy, and MHPCC systems? (Ref: Sec. A, Introduction/Problem Statement (names "AI/ML platforms" generically))

Sec. C, Proposal Evaluation Applicability criterion (references "traditional and AI/ML workloads in an HPCMP context" without naming specific mission applications))

The anticipated workload will be a mixture of AI/ML (e.g. LLM's, surrogate models, reinforcement learning) as well as traditional modeling and simulation that has historically run on HPCMP's systems.

Please note: This CSO announcement initiates a multi-step process. The first step is the solicitation of solution briefs. The purpose of the solution brief submission is to identify potential partners that may have promising solutions. Subsequent steps consist of opportunities for selected offerors to deliver solution pitches, demonstrations, and collaborative

exchanges with the Government. As such, offerors do not need to submit initial solutions with detailed technical specifications trying to satisfy technical requirements that will be identified in later steps.

Facilities & Power: Of the stated available power (2 MW at AFRL and Navy

350 kW at MHPCC), what percentage is expected to serve IT load (GPU+CPU) vs. cooling/facility overhead, and is there a target rack power density (kW/rack) or PUE the Government is designing against? (Ref: Appendix B Sec. 1.3 / 2.3 (AFRL Power/Electrical), Appendix C Sec. 2.3 (MHPCC Power/Electrical), Appendix D (Navy Power/Electrical) all give total kW/MW available but no IT-load split or PUE target)

AFRL response: 2MW is available to the total system residing in the computer room, including any CDUs and other ancillary equipment the OEM requires to cool and operate the system. There is no target rack/power density. The OEM is responsible for providing the solution that fits within the facility specifications given, such as floor loading limitations and 45 deg-F chilled water supply. There is no PUE target.

MHPCC response: There is no specific split or PUE target; however, innovative solutions that optimize PUE are encouraged.

Navy response: The 2MW identified for the TI-27 solution is expected to serve 100% of the IT solution proposed. This shall include any CDUs or rear door heat exchangers that need to be included to meet the expectations of the TI-27 procurement. The PUE target is managed by Navy facility staff and has already been considered.

Sizing & Performance: Is there a target aggregate compute capacity (e.g., PFLOPS) or reference benchmark (HPL, HPL-AI, MLPerf) used to derive the funding allocation across AFRL, Navy, and MHPCC, or is sizing intended to be budget-constrained? (Ref: Sec. A, Estimated Government Funding Profile (gives \$ split only, no capacity/benchmark target)

Appendix A (architecture scenarios have no numeric sizing))

Sizing is intended to be budget constrained.

Please note: This CSO announcement initiates a multi-step process. The first step is the solicitation of solution briefs. The purpose of the solution brief submission is to identify potential partners that may have promising solutions. Subsequent steps consist of opportunities for selected offerors to deliver solution pitches, demonstrations, and collaborative exchanges with the Government. As such, offerors do not need to submit initial solutions with detailed technical specifications trying to satisfy technical requirements that will be identified in later steps.

Security & Classification: For AFRL's above-secret option and the MHPCC TS-SCI facility, what accreditation baseline applies, is a cross-domain solution required between enclaves, and what is the expected workload split between secret and above-secret tiers? (Ref: Sec. A,

Constraints and Considerations (security/RMF/STIG bullets name NIST SP 800-53, CNSSI 1253, STIGs generally, but no IL level, no CDS requirement, no workload split))

AFRL response: For classified AFRL environments, offerors should assume the security and accreditation approach will align with the applicable Government-approved baseline for the target operating environment. For TS//SCI environments, this includes NIST SP 800-53 Rev. 5, CNSSI 1253, the applicable classified overlay, the applicable Integrity level (A, B, or C), and applicable Accessibility requirements depending on the environment.

For SAR environments (S/TS), offerors should assume NIST SP 800-53 Rev. 5, CNSSI 1253, JSIG (2016) (transitioning to SOWIG, as applicable), the applicable classified overlay, applicable Accessibility requirements, and the applicable Integrity level (A, B, or C) depending on the environment, including TS//SAR//SCI contexts where applicable.

AFRL does not have a cross-domain solution (CDS) requirement associated with this question. AFRL's workload is probably going to be 100% above secret.

MHPCC response: A Cross-Domain solution is not required.

Software & Integration: Are there required or preferred job schedulers, container orchestration tools, AI/ML frameworks, or integration protocols for the HPCMP Enterprise Data Fabric and centralized User Authentication systems referenced in the CSO? (Ref: Sec. A, Constraints and Considerations: "Integration with DoW HPCMP enterprise capabilities (e.g. Enterprise Data Fabric, User Authentication)" bullet names the systems but gives no technical/API detail)

Offerors should assume that the proposed solution must be capable of integrating with the Government-furnished enterprise capabilities identified in the solicitation, including centralized user-authentication services, job reporting and accounting, and the HPCMP Enterprise Data Fabric, as applicable. For user authentication, offerors should assume interoperability with the Government's existing Kerberos-based authentication environment.

Beyond that, the Government does not prescribe a single required job scheduler, container orchestration tool, AI/ML framework, or specific enterprise-integration protocol or API beyond the requirements expressly stated in the solicitation. Offerors should clearly describe the software, orchestration, and integration approach they propose, including any assumptions regarding authentication, data access, interoperability, and operational supportability.

Facilities Monitoring / Digital Twin: Do the DSRCs currently operate a Data Center Infrastructure Management (DCIM) or Building Management System (BMS) at AFRL, Navy, and MHPCC, and would the Government value a facility/system digital twin for real-time power, cooling, and utilization telemetry as part of the Administration or Maintenance scope? (Ref: Not Applicable)

AFRL response: AFRL uses an HPCMP funded tool called the Energy Monitoring and Management System (EMMS) to monitor electrical/mechanical plant equipment status, power consumption, chilled water supply/return temperature and flow, and room temperature/humidity. AFRL's implementation of EMMS does NOT control mechanical/electrical equipment, it only monitors. Facility/system digital twin is not required but the vendor may propose one if they believe it provides significant value to the Government.

MHPCC response: Sunbird's DCTrack and PowerIQ are under development. The Government would welcome if an offeror's solution easily integrates into that DCIM solution.

Navy response: Navy utilizes a BMS system. It is not desired to have a digital twin for the facility as part of the TI-27 solution.

Data Rights & IP: What data rights or access does the Government expect to retain for telemetry, performance, and configuration data generated by monitoring/analytics tooling (e.g., a digital twin platform) delivered under this effort? (Ref: Sec. D, Award notes "additional terms and conditions may be required...examples...would be data rights" but gives no specifics on scope or ownership expectations)

The Government anticipates requiring sufficient rights and access to telemetry, performance, configuration, and other operational data generated under this effort to support Government operations, cybersecurity, sustainment, analysis, transition, and follow-on support. The Government's intent is to ensure that operational data and associated deliverable documentation are accessible and usable by the Government for Government purposes, while specific data-rights treatment for vendor pre-existing proprietary software, tools, and other intellectual property will be addressed in the final solicitation and contract terms, as appropriate.

Facilities — Cooling Headroom: Beyond the CDU/tonnage figures stated in the facility appendices (e.g., ~31 tons redundant CDU capacity at MHPCC), is there headroom for offeror-supplied supplemental CDUs or liquid-cooling infrastructure if a GPU-dense architecture is proposed? (Ref: Appendix C Sec. 2.2, Cooling (MHPCC) gives current CDU tonnage and notes offeror may propose an "alternate CDU solution" but doesn't state available physical/electrical headroom for one)

If offeror has alternative innovative solutions beyond the specified facility CDU capacity, the Government is open to consider allocating additional physical/electrical headroom to the solution. Specifics would be discussed during negotiations.

Constraints and Considerations / Estimated Period of Performance: Please confirm the intended administration and maintenance period of performance. Should offerors price a five-year term, or a one-year base period with up to six one-year options? (Ref: Proposed solutions shall include administration and maintenance services for a five-year term after acceptance administration and maintenance period is one year with options for up to six additional one-year increments)

The system administration and maintenance support is expected to have a one-year base period with six options of one-year each for a total of seven years if all options are exercised. Pricing for the base period and each option shall be proposed by the offeror.

Estimated Government Funding Profile: How should offerors develop ROM assumptions for administration and maintenance beyond year five, when the stated 6% annual budget assumption applies only through year five? (Ref: HPCMP budgets 6% of the system's original purchase price per year for up to five years of administration and maintenance)

Beyond year five, offerors may develop ROMs that do not adhere to the 6% assumption; however the Government is interested in innovative approaches that either sustain or reduce costs across all administration and maintenance years.

Solution Brief Preparation and Submission: Does the ten-page limit include the acronym page, ROM estimate, figures, and diagrams and does the 12 pt font also apply to graphics and diagrams? (Ref: General Formatting requirements: Type Font 12 point TNR. Solution briefs shall be no more than ten pages

one page of the proposal body is allocated to acronyms, abbreviations, and symbols.)

The cover/title page does not count toward the 10-page limit: other pages are to be included within the 10-page limit. The 12 pt font applies to graphics and diagrams.

Constraints and Considerations: How will system availability be calculated for purposes of prorating maintenance fees, including the treatment of planned maintenance, DREN or other Government-caused outages, and the monitoring source of record? (Ref: Maintenance fees will be prorated on a monthly basis based on system availability to customers)

System availability generally excludes approved scheduled maintenance and Government-caused events.

Appendix B - AFRL First Building: Please confirm the available floor space for the AFRL deployment location, as Appendix B references both 1,400 square feet and 1,440 square feet. (Ref: The summary table identifies 1,400 square feet available for the new system, while Section 1.1 states that 1,440 square feet will be set aside)

For AFRL, 1400 ft² is available in the unclassified area and 320 ft² is available in the classified area.

Appendix C - MHPCC Classified Floor Space: Please confirm the available raised floor space at MHPCC, as Appendix C references both 224 square feet and 244 square feet. (Ref: The summary table identifies 224 square feet available for the new system, while Section 2.1 states that 244 square feet has been set aside)

For MHPCC, the correct floor space availability is 224 ft².

Constraints and Considerations / Appendix C: Given the objective to deliver systems within one calendar year and the stated November 2027 availability of the MHPCC facility, how

should offerors account for deployment scheduling and system acceptance assumptions at that location? (Ref: Systems should be deliverable within less than one calendar year the MHPCC facility is scheduled to be available for occupancy in November 2027)

After solution selection, MHPCC and offeror can negotiate if temporary storage is required following factory integration and testing and prior to site readiness.

Constraints and Considerations: Is the offeror responsible for producing all SSP / RMF packages and what is the existing average timeline to achieve ATO? (Ref: All acquired systems must meet all physical and information security requirements set by DoW, the Government worksite, and the Information Systems Security Manager, primarily detailed by the DISA Risk Management Framework (RMF) NIST SP 800-53r4.)

Requirements and responsibilities for producing RMF artifacts and achieving ATOs will be determined during scope collaboration with the selected offeror. Timelines to achieve ATOs vary depending on classification level and different Authorizing Official processes.

Constraints and Considerations: The CSO requires full disk encryption "approved for national security information," can the government provide additional detail on the level of encryption required (e.g., NSA CSfC/Type 1, FIPS 140)? (Ref: For classified systems and unclassified systems identified as requiring additional data protection, all provided non-volatile storage must support full disk encryption utilizing a cipher approved for national security information.)

Requirements and responsibilities for ensuring data protection and associated encryption methods will be determined during scope collaboration with the selected offeror.

Please note: This CSO announcement initiates a multi-step process. The first step is the solicitation of solution briefs. The purpose of the solution brief submission is to identify potential partners that may have promising solutions. Subsequent steps consist of opportunities for selected offerors to deliver solution pitches, demonstrations, and collaborative exchanges with the Government. As such, offerors do not need to submit initial solutions with detailed technical specifications trying to satisfy technical requirements that will be identified in later steps.

For AFRL, should an offeror select one Appendix A architecture/classification scenario and design to the corresponding facility, or may multiple scenario and facility alternatives be proposed for Government consideration during collaboration?

Offerors may propose multiple scenario and facility alternatives to be considered.

For classified TI-27 requirements, if an otherwise qualified offeror does not hold the required Facility Clearance Level, will the Government sponsor the offeror for the required FCL and

issue a DD Form 254 during the RFP phase, or must the required FCL already be in place prior to award?

The required FCL must already be in place prior to any award.

Beyond the identified Hammerspace enterprise data fabric, does HPCMP have mandated or incumbent standards for workload scheduling, container orchestration, cluster provisioning, monitoring/observability, or configuration management, or may the offeror propose those elements of the management software stack? The FAQ confirms Hammerspace integration is expected and that software delivered with a new HPC system is in scope.

Offerors should assume that the proposed solution must be capable of integrating with the Government-furnished enterprise capabilities identified in the solicitation, including the HPCMP Enterprise Data Fabric and centralized user-authentication services, as applicable. Beyond those identified integration points and any other requirements expressly stated in the solicitation, the Government does not prescribe a single mandated or incumbent software-stack standard for workload scheduling, container orchestration, cluster provisioning, monitoring/observability, or configuration management.

Offerors may propose those elements of the management software stack, provided the proposed approach satisfies the solicitation requirements and supports effective integration, operability, supportability, and cybersecurity within the Government environment. Offerors shall clearly describe the proposed software-stack approach and all assumptions regarding interoperability.

Please note: This CSO announcement initiates a multi-step process. The first step is the solicitation of solution briefs. The purpose of the solution brief submission is to identify potential partners that may have promising solutions. Subsequent steps consist of opportunities for selected offerors to deliver solution pitches, demonstrations, and collaborative exchanges with the Government. As such, offerors do not need to submit initial solutions with detailed technical specifications trying to satisfy technical requirements that will be identified in later steps.

Are there any incumbent application codes or benchmarks offerors should use for pre-award sizing and ROM development before the final Capability Test suite is established? The FAQ says final CT benchmarks will be developed collaboratively after selection, but it does not identify pre-award sizing benchmarks.

The final benchmarks will be developed collaboratively after selection. They are not predetermined. Offerors wishing to prepare pre-award materials may wish to consider the HPCMP user community's common workflows and application mix. Commonly used HPC community benchmarks as well as representative specific full codes have been used for past TI's. Determining final benchmarks later provides flexibility to accommodate innovation and maximize system cost effectiveness.

Please note: This CSO announcement initiates a multi-step process. The first step is the solicitation of solution briefs. The purpose of the solution brief submission is to identify

potential partners that may have promising solutions. Subsequent steps consist of opportunities for selected offerors to deliver solution pitches, demonstrations, and collaborative exchanges with the Government. As such, offerors do not need to submit initial solutions with detailed technical specifications trying to satisfy technical requirements that will be identified in later steps.

May an offeror submit separate 10-page solution briefs for different DSRCs, and may the same offeror submit more than one brief for a single DSRC when the proposed technical approaches are materially different? The FAQ only says one brief is sufficient for a prime/subcontractor team. It does not establish a one-brief-per-offeror rule.

Yes, offerors may submit separate 10-page solution briefs for different DSRCs and more than one brief for a single DSRC if the proposed technical approaches are materially different.

Will AFRL, Navy, and MHPCC solutions be evaluated and selected independently by DSRC, or can an integrated multi-site offering affect the selection decision for an individual DSRC? The FAQ says commonality is desirable but allows site-specific optimization.

Proposed solutions will be evaluated for both individual DSRC as well as multi-site potential.

For the initial solution-brief downselect, are innovation/feasibility, mission effectiveness, and funding availability weighted equally, or are any of these factors given greater consideration?

The first two factors are more important than funding availability.

Please clarify the AFRL TDS power-backup requirement: the facility summary states that UPS/generator backup is not available, while Section 3.3 states that the building is backed by a 3 MW UPS/generator system.

For AFRL, the area hosting the TDS is generator backed up. Facility water is available at the TDS location. The TDS must stay within the 600 lb/tile and 150 lb/ft² floor ratings. Offerors should not assume that the TDS can necessarily accommodate a rack identical in weight, utility demand, or cooling method to the production system unless that capability is expressly confirmed by the solicitation. For purposes of this acquisition, TDS parity is intended to support representative test and development functionality and does not, by itself, guarantee full physical rack parity with the production system.

The CSO states that approximately 31 tons of liquid cooling will be available through the MHPCC facility CDUs and that additional chilled-air capacity is available. For ROM and system-sizing purposes, how much supplemental chilled-air capacity should offerors assume is available to the TI-27 system?

While there is sizable amount of chilled-air capacity initially available, the Government prefers that only a residual amount of cooling requirement be air cooled. Specifics can be negotiated after solution selection.

Given that the MHPCC MILCON facility is scheduled for occupancy in November 2027, how should offerors apply the requirement to deliver systems within less than one calendar year, and from what event will the MHPCC delivery period be measured?

It is expected that final contract award will be less than one year from the MHPCC site availability. However, after selection, MHPCC and offeror can negotiate if temporary storage is required following factory integration and testing and prior to site readiness.