A Configurable Testbed for Evaluating Novel HPC Architectures to Enhance NOAA Mission Capabilities

Gerry Kolosvary[a], Steve Heibein [a], David W Miller [b], Terri T Miller [b]

[a] FedCentric Technologies Inc, [b] Miller Engineering and Research Corporation (MERC)

Abstract

High-Performance Computing (HPC) is at the cutting edge of IT technology, where trends rapidly change through successive generations of technology, requiring that applications constantly adapt to ensure they can best exploit the computational platforms. While the NOAA High Performance Computing and Communications Branch (HPCC) has the charter to provide both operational and research and development (R&D) services to the NOAA missions, it is faced with challenges in funding which result in a less than optimal R&D budget to support the evolution of the HPC hardware and software solutions development. This white paper describes a cost-effective and targeted approach to help NOAA achieve significant goals and objectives as stated in the NOAA High-Performance Computing (HPC) Strategic Plan 2015-2020.

Three key features of this approach are

1) The provision of a configurable HPC Testbed as a Service (TaaS) (See Figure 1).
2) The provision of application tuning and HPC subject matter expertise (SME) and systems integration services that guide the development, use and exploitation of the testbed in partnership with NOAA's HPC experts to accomplish NOAA's objectives.
3) Provision of High Performance Data Analytics on existing NOAA data sets to further enrich the value of NOAA assets.

This approach will allow NOAA to more fully explore technology advances and techniques in a risk-free, cost-effective environment to better understand hardware and software approaches, that when implemented, greatly reduce operational risk and provide a path to exponential performance increase needed to support current and future weather and climate models.

Figure 1. Our configurable testbed provided by a systems integration team as a service, not as a capital purchase.

Overview

In their Strategic Plan, NOAA outlines 15 objectives mapped to 5 specific goals. There is a very cost-effective solution that can be implemented to target up to 13 of these 15 objectives. The solution is centered on several key elements:

4) Provision of a configurable testbed, dedicated to NOAA's use and *provided as a service* to NOAA, not as a capital purchase. Using this Testbed as a Service (TaaS) implies NOAA only pays for the configuration and use of the testbed for specific tasking defined jointly by NOAA HPCC and the systems integration team who owns and operates the testbed;
5) The implementation of this testbed, under the guidance of an experienced HPC systems integration team where hardware and software requirements are met with low-risk, high performance integrated solutions;
6) The addition of a systems integration team member that is intimately knowledgeable of NOAA science requirements that drive their HPC applications;
7) Execution by a cost-effective small business team that is driven by results without the overhead costs of large systems integration contractor;
8) Testbed houses flexible compute, storage, and network hardware that supports reconfiguration as scale-out, scale-up, and hybrid architectures;
9) Testbed equipment includes regular refresh to stay abreast of performance and capacity advances; and,
10) HPC vendor loaner hardware and software are encouraged to supplement NOAA HPC testbed so investigations are not limited to equipment on hand.

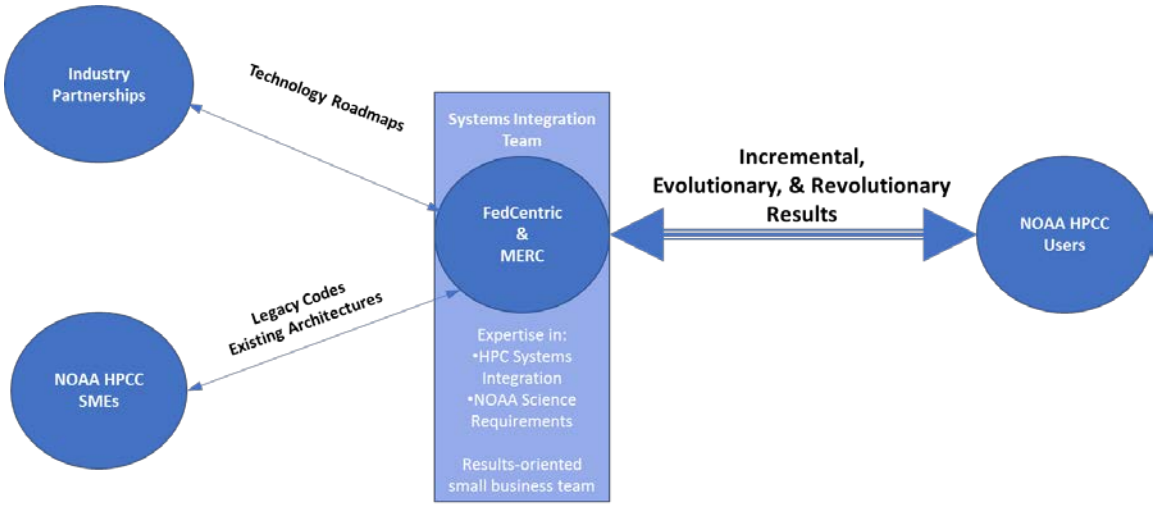

Figure 2. Systems Integration Team Competencies and Contribution

The systems integration team, as shown in Figure 2, performs the following functions:

1) Engage HPC vendors to ensure NOAA has insight into company and industry roadmaps and access to latest technologies for testing.
2) Engage NOAA HPCC and model (SMEs) to set priorities and use of the TaaS resources.
3) Work with model owners to determine performance bottlenecks.
4) Propose alternatives for testing to improve performance.

5) Configure the TaaS environment for the investigation.
6) Work with NOAA SMEs to run investigations and validate results.
7) Communicate investigation results to HPCC stakeholders and the greater NOAA community.
8) Incorporate feedback into future investigations
9) Provided recommendations on HPC strategy and tactics based on R&D performed on the TaaS platform

Table 1 provides a brief description of the solutions promoted in this white paper and the resulting benefits to NOAA HPCC's goals and objectives.

Goal & Objective		Description	Solution	Benefit
1	**Provide enterprise HPC services to enable the agency's mission**			
	1	Maintain core competency and reliable HPC enterprise capability	• A configurable, dedicated NOAA HPC Testbed as a Service (TaaS). • Identify bottlenecks in current approaches. • SMEs suggest possible alternatives. • Reconfigure TaaS to test NOAA applications and/or data and provide quantitative and qualitative results.	• A production scale environment to measure impact of software optimizations, hardware enhancements, architecture changes, and alternative approaches. • Data points for NOAA decision-makers. • Hands-on access to emerging technology
	2	Leverage other federal agencies HPC shared services to increase NOAA's access to leadership-class computing and novel architectures	Leverage industry partners that have provided significant improvements in HPC technologies for other federal agencies by employing novel architectures including HPE, SGI, Cray, Intel, NVIDIA hardware	• Lessons learned from other agencies' HPC evolutions • Access to industry's top leadership • Collaboration reduces risks associated with new technology
	3	Provision and allocation of HPC services consistent with requirements of NOAA's mission workflow	Results driven HPC investigations using NOAA application codes and analogous benchmarks that represent NOAA mission and workflow.	• Provides actionable information on code, system, and workflow optimization • Stays on mission • Eliminates research for curiosity's sake
2	**Improve Linkage Between Mission Requirements and HPC Solutions**			
	1	Establish integrated support of NOAA's applications to optimize HPC solutions and to guide model development based on changing technologies	A configurable, dedicated HPC testbed provided as a Service to the NOAA software engineering team responsible for the model developments. SI and industry partner SMEs can tune code for new architectures.	• Risk reduction for applications development, optimization and implementation as part of advancing operational improvements • Reduced implementation costs through leveraged capabilities
	2	Develop scientific and software assurance methodologies to deal with increasing complexity of HPC systems	A configurable, dedicated NOAA HPC testbed provided as a Service as to test alternative workflows, user tools, and methodologies.	• Offline testing of software automation techniques without adding risk to production operations • Provides pathway for exponential (not just incremental) improvements.

Goal & Objective		Description	Solution	Benefit
			Prototype platform to demonstrate best practice from other gov and commercial HPC sites.	• Reduces missed opportunities for improved HPC capabilities
	3	Explore innovative methodologies and solutions to improve efficiency of data storage and effectiveness of data analytics	Industry partners who are both HPC technology experts and systems integration engineers who understand most efficient and effective approaches to architect hardware/software solutions. Reconfigurable system to accommodate emerging HPC architectures and support vendor loaner equipment, extending the range of alternatives investigated.	• Most affordable solutions provided by system architects who have performed their own data analytics, developing their own toolsets (i.e. graph analytics, machine learning, memory centric computing) • Reduced time to insight • Broader range of alternatives
3		**Recognize and Plan for Emerging Uses of HPC**		
	1	Perform outreach and education to NOAA programs about HPC services	Sponsor a vibrant NOAA HPC users group, provide content for NOAA internal & external web sites, highlight investigation results, disseminate knowledge, provide training materials, technical libraries and on-line documentation.	• Training and demonstrations provided to new and potential users • Shared knowledge among the user community enhances best-practices, provides SOTA HPC knowledge that can extend the HPC user community and provide user and application efficiencies.
	2	Enhance support to accommodate new communities and technologies	Internal seminars demonstrating NOAA HPC testbed investigation results and NOAA successes. Extensible testbed encouraging vendor loaners and SME involvement.	• Demonstrations remove the fear, uncertainty, and doubt holding back new users. • Vendor loaners bring latest technology and additional points of view to extend alternatives that can be considered.
	3	Evolve funding model to account for emerging HPC users	HPC TaaS to identify the most successful approaches/solutions prior to making HPC expenditure.	• Sizing is more accurate. • Cost estimates are based on measured results. • Reduces risk
4		**Effectively Adopt Latest HPC Technology to Drive Efficiencies**		
	1	Access novel architectures	Leverage industry partners who are developing novel architectures. Supplement NOAA HPC testbed with vendor loaners and SMEs to test NOAA applications on these architectures.	• Access to and trial of novel and architectures and emerging products. • Lessons learned from other installations and vendor SMEs.

Goal & Objective		Description	Solution	Benefit
	2	Create training opportunities and workshops to enable exchange of knowledge between model developers and HPC solution providers	• Facilitate solution whiteboarding sessions with model developers and HPC SMEs. • Facilitate training seminars by testbed provider and vendor SMEs • Train the trainer • On-site and remote training • On-demand training videos	• Extends knowledge transfer • Enthuses model developers to embrace new solutions.
	3	Develop integrated enterprise-wide competency to recognize, plan, and evaluate emerging technologies for future acquisitions	Facilitate HPC outreach via multimedia and face-to-face knowledge transfer with audience appropriate information on HPC trends, best practices, and investigation results.	• Provides NOAA decisionmakers with actionable information. • Provides model developers with tradeoffs of alternative approaches. • Provides operations team with best practices • Provides regular information exchanges between HPC stakeholders
5		**Maximize Effectiveness of HPC Solutions**		
	3	Improve the user experience	Provide audience appropriate information of quantitative and qualitative results. Include impact on user workflow when testing innovative approaches.	• Results are more digestible • Increases user acceptance of new approaches

Table 1. Strategic solutions that help NOAA HPCC achieve their goals.

NOAA HPC Goals and Objectives and this White Paper

The goal of the NOAA test bed would be to investigate and make recommendations regarding new technologies and how it might greatly enhance NOAA's Mission. In the "High Performance Computing Strategic Plan 2015-2020 (Final Draft)", the Office of the Chief Information Officer of NOAA states five Goals and associated Objectives for High Performance Computing to achieve its mission. Listed below are NOAA's stated goals, objectives, and accompanying dialog (in italics). With the MERC/FC team's capabilities and experience, coupled with a NOAA-dedicated testbed facility, NOAA can specifically and measurably address the following goals and objectives.

Goal 1: Provide enterprise HPC services to enable the agency's mission.

Objective 1: Maintain Core Competency and Reliable HPC Enterprise Capability

Computing capacity, short and long-term storage, and scientific analysis are central to NOAA's mission. NOAA will continue to operate and maintain its core enterprise capability and services for computing; post-processing and analysis; and long-term storage assets. In-house HPC expertise familiar with the mission applications, is a key

asset essential to optimize, continuously modernize, and exploit the use of NOAA's HPC capabilities for the maximum return.

NOAA's stated objective is to provide in-house HPC expertise familiar with mission applications as a key asset essential to optimizing and continuously modernizing their HPC capabilities for maximum return. A configurable HPC testbed provided as a service provides NOAA HPCC with the needed resources to test and vet their unique requirements with alternative options including application optimization, hardware modernization, and innovative approaches. The systems integration team works with NOAA HPCC to architect the testbed testing priorities and approaches to ensure NOAA's HPC capabilities are maintained and continuously modernized.

Objective 2: Leverage other federal agencies HPC shared services to increase NOAA's access to leadership-class computing and novel architectures

On the path to exascale computing, NOAA will leverage leadership-class facilities that are available throughout the Government. Access to these national assets will enable NOAA to scale their work further and leverage engineering and resources that would not be otherwise attainable. Access to leadership class computing will enable breakthrough research.

Whereas the configurable HPC testbed would be dedicated to NOAA and not leveraged from other federal agency resources or technology, it would increase NOAA's access to leadership-class computing and novel architectures if it were provided by industry partners with cross-agency experience and operate at the leading edge of HPC technology through direct partnerships with the leading developers and manufacturers of the technology. This partnership provides NOAA HPCC with access to lessons learned from other federal agency HPC achievements.

While industry partners do not represent national assets, NOAA can use dedicated industry resources to scale their work further and thereby leverage engineering and resources that would not be otherwise available. Access to industry's top leadership in HPC greatly reduces risk associated with implementation of new configurations and reduces time to insight for NOAA scientists.

Objective 3: Provision and allocation of HPC services consistent with requirements of NOAA's mission workflow

Computational and storage capabilities are determined by model and workflow requirements. Working with agency partners, resources will be acquired and configured to achieve optimal results.

Since computational and storage capabilities are determined by model and workflow requirements, it is imperative that industry partners who are providing configurable HPC testbed resources understand the systems engineering aspect of this technology. Configuration of HPC hardware in and of itself can provide increased performance. However, coupling improvements in hardware configuration, even without technology upgrades, with software optimization that leverage the hardware configuration can bring

sizeable performance increases with input/output capability and reduced latency with minimal capital expenditure.

Goal 2: Improve Linkage Between Mission Requirements and HPC Solutions

Objective 1: Establish integrated support of NOAA's applications to optimize HPC solutions and to guide model development based on changing technologies

As heterogeneous computing architectures and their associated programming environments mature, the interplay between hardware and software application design has become more tightly coupled. NOAA needs to continually adapt to a more complex set of emerging programming standards and wider array of heterogeneous computing solutions. NOAA will develop an integrated software engineering team, comprised of lab and center personnel, to enhance its software engineering discipline and expertise to achieve optimal code performance and scaling (while maintaining code portability), maximize the efficiency of transitioning research to operations, and enable effective collaborative model development with partners both internal and external to NOAA. In close collaboration with the modeling community, this integrated software team will provide the necessary bridge for environmental modeling to work in tandem with the existing integrated management team, which currently manages the acquisition, provisioning, and day to day operation of HPC.

An industry-provided configurable testbed is the resource to be used under the leadership of the NOAA software engineering team. Through their partnership with the right systems integration team, this resource allows the team to accomplish its stated objectives as follows:

- Testbed configured similarly to the existing operations hardware can be used to develop new applications, optimize existing ones, and ensure risk has been reduced or retired in their transition to operations, as part of advancing operational improvements
- Testbed configured with newer technology hardware can be used to develop and trial models created by the integrated software team demonstrating and validating expected performance improvements to management teams
- Memory scale-up can be trialed, advancing operational improvements such as moving the compute to the data and reducing inherent latency

This approach can reduce implementation costs by leveraging this risk-reducing capabilities.

Objective 2: Develop scientific and software assurance methodologies to deal with increasing complexity of HPC systems

The mean time between failures for HPC will decrease dramatically as the number of processing cores and integrated components within the HPC increase by orders of magnitude. This increase in system complexity leads to more frequent job runtime interruption and potential data corruption requiring additional sophistication in the applications to achieve the current level of runtime reliability and data integrity. NOAA will need to invest in software automation techniques required for fault tolerant data movement, and error detection, correction, and handling to ensure mission applications run reliably. NOAA scientists will have to increasingly work together with software

engineers to develop workflows that have fault tolerance built in from their inception. Interdisciplinary teams will become the norm for model development within NOAA.

With a dedicated NOAA HPC testbed, software automation techniques are tested offline without adding risk to current operations. Development of successful workflows requires an understanding of the software platforms and the data sets that are to be processed. This approach results in a scale-up approach, a scale-out approach, or a hybrid approach all of which can be trialed with the testbed to ensure the optimal solution is identified.

Objective 3: Explore innovative methodologies and solutions to improve efficiency of data storage and effectiveness of data analytics

Within the next five years, NOAA's hierarchical storage management system, within its HPC program, will store hundreds of petabytes of data. While growth in computational demand continues at an exponential rate, storage performance and capacity is lagging behind creating an incrementally increasing cost for storage over time.

NOAA will refine its workflows to ensure appropriate trade-offs are made when storing data. Engagement with industry will continue to ensure that storage and data movement technologies can continue to scale to the required capability and reliability. Hardware and software innovations will be exploited to enhance scientific analysis.

To thoroughly exploit hardware and software innovations in HPC technologies, the expertise required includes not only partners who are experts in the evolution of HPC technology but also systems integration engineers who understand most efficient and effective approaches to architecting hardware/software solutions that in turn provide the most effective and efficient computational environment for NOAA's software applications and scientific needs. Expertise in HPDA techniques to exploit analysis of stored data provides a specific application example. There is an inherent cost-effectiveness to be gained by having access to system architects that understand and have developed tools such as graph databases to generate data analytics. Software platforms as well as model and workflow requirements must drive hardware requirements. True systems integrators know how to manage this risk to provide seamless transition to operations through testbed experimentation. Having a flexible, i.e. configurable, testbed system allows for trialing the most promising technologies providing an alternative platform and/or novel hardware configuration.

Goal 3: Recognize and Plan for Emerging Uses of HPC

Objective 1: Perform outreach and education to NOAA programs about HPC services.

Outreach and education will be provided to communities within NOAA which have not traditionally used HPC. The HPC program helps translate the complexities of using HPC environments involved with parallel programming, debugging, data movement and storage. Training will be provided by experts from NOAA, our partners, and industry. Enterprise HPC wikis, help desk support and NOAA application analysts will assist this NOAA community in making efficient use of HPC resources.

Results from testbed investigations could be leveraged for outreach and education. Training, demonstrations, and findings could be provided to new and potential users. The systems

integration team could sponsor a NOAA HPC users group to create training materials, technical libraries and on-line documents. Shared knowledge among the user community always enhances best-practices and provides state-of-the-art (SOTA) HPC knowledge that increases efficiencies and throughput.

Objective 2: Enhance support to accommodate new communities and technologies

> *Computational testbeds will be established to support initial "development" allocations. These testbeds will be associated with applications and technology liaisons to assist new users. The NOAA applications lead will engage new Principal Investigators to ensure that they can adapt their codes to new platforms and that they are aware of the latest programming standards.*
>
> *As NOAA's HPC program has evolved and matured, the user community has grown and diversified. To foster a strong relationship between the HPC program and its users, NOAA will establish an HPC User Committee with representation from the major user communities and some of the "pioneer" users identified by the NOAA Allocation Committee. The User Committee will bring the advice, guidance, and point of view of computer users to the attention of the HPC program; and to exchange information concerning effective utilization of the HPC resources available to NOAA.*

The configurable testbed can be used to develop NOAA's approach to supporting new end users integration into the HPC program. Unlike operational HPC assets of NOAA, this testbed can be configured to demonstrate to new users performance expectations from the current operational HPC environment or to determine the optimal environment for new users. Through internal seminars, demonstrated successes can be shared with the current user community as well.

Objective 3: Evolve funding model to account for emerging HPC users

> *Currently, NOAA's HPC assets are funded primarily by OAR and NWS. (Operational supercomputing is augmented by NOS pursuant to an MOA with NWS on hydrodynamic modeling for oceans and coasts.) NOS, NESDIS, and NMFS, each receive small allocations on NOAA's HPC. As these new users' mission requirements grow, a new funding model will need to be established to recognize the cost of providing those services. Significant new computational and storage allocations will only be provided when programming, network, and HPC support are provided by the requestor. In the long term, emerging users will need to contribute funding and some personnel support towards the overall cost of the HPC enterprise and provided services.*

By using the configurable testbed approach, the most successful approaches that bring the most value for the expenditure can be identified. This approach enhances NOAA HPCC's ability to reduce risk inherent in the capital budgeting process, while increasing the branch's capabilities and informing their capable user community. The testbed can be leveraged to demonstrate to new end users how their requirements translate into system usage expectations for the purpose of quoting a price for the HPC service to be provided by NOAA. Baseline user metrics can be developed and communicated as a basis for provisions of the HPC service.

Goal 4: Effectively Adopt Latest HPC Technology to Drive Efficiencies

Objective 1: Access novel architectures

High Performance Computing is at the cutting edge of IT technology. HPC trends rapidly change through successive generations of technology and require that applications constantly adapt to ensure they can best exploit the computational platforms. Access to novel architectures enables developers to explore the opportunities and challenges associated with running new and existing applications on the next generation of computing.

NOAA will identify technology testbeds for novel architectures. *The capabilities of other agencies, university partners, industry and our own assets will be leveraged. Early adopters will be able to evaluate the programming, I/O handling, software, and storage environments for NOAA's mission workload.*

The configurability of a testbed is a necessity in the field of HPC technologies as this hardware is costly and the success of its architecting and integration into existing HPC systems truly determines the performance enhancement. Access to novel architectures comes through leveraging industry partners who are systems integrators, have an understanding of the historical development of HPC hardware and software, have access to the latest technological developments in hardware and have used HPC for their own computation projects. Scalability is a key requirement when using a testbed to evaluate specific performance characteristics such as I/O handling, programming and storage. Systems integration experts can both determine and communicate the logic required to assess scalability of given system architectures. Proof of concepts are easily achieved by obtaining latest hardware from vendors to demonstrate capabilities.

Objective 2: Create training opportunities and workshops to enable exchange of knowledge between model developers and HPC solution providers

Understanding the programming, software, and technology limitations of next-generation architectures is beneficial when implementing new algorithms for numerical models. Advanced techniques ensure that applications can scale and exploit new platforms without having to be completely rewritten.

NOAA experiences will be communicated with industry and our partners to ensure that our concerns, "lessons learned", and successes can be leveraged to improve the next-generation technologies. NOAA will benefit by ensuring that appropriate enhancements will be made to enable environmental codes to exploit these architectures. This will serve as a benefit not just for NOAA, but the HPC community as a whole - as the technology will be more stable, functional, and better performing.

Trialing of advance techniques on a properly configured testbed provides many benefits to NOAA. This resource can be used to facilitate training seminars including participation of appropriate vendors. Training can be provided on-site, remotely, through YouTube videos, and train-the-trainer courses. This approach serves for risk reduction in deploying new applications into the operational system; it serves to trial the next-generation architecture through the selection of configuration for the testbed; it serves to develop and establish performance metrics for enhancements in environmental codes; it serves as a training and

workshop opportunity for the HPC community under NOAA's leadership; and, it creates potential for an expanded user base.

Objective 3: Develop integrated enterprise-wide competency to recognize, plan, and evaluate emerging technologies for future acquisitions

The leading NOAA applications engineers will work with domain scientists to identify agency-wide benchmarks that are both relevant and computationally challenging. These applications will be cross-compiled on various architectures and made available to the entire modeling community and HPC vendors. As appropriate, performance numbers will be made publicly available. Since new software capabilities are typically frozen during operational transitions to new technologies, these transitions should be planned and resourced to occur with maximum reliability and minimum elapsed time.

Regular interactions with HPC vendors will be ongoing to ensure that both technology challenges and successes are communicated. Through these regular interactions, NOAA's leading application engineers and community of users will remain informed of the latest software and technology trends by participating in various community forums.

The experience to be gained by NOAA applications engineers through use of a dedicated configurable testbed cannot be quantified. Not only will it lead to a substantial increase in the competency of both NOAA engineers and scientists to recognize, plan and evaluate emerging technologies for future acquisitions, but also reduces the risk inherent in planning for future hardware and systems upgrades. Leading the testbed development with a systems integration team that is "vendor" agnostic and has close working relationships with all suppliers, gives NOAA the advantage of working with any HPC vendors technology, whether it is COTS or a cutting-edge development or something in between such as a prerelease product. It is most imperative your systems integration team know and understand the current operational environment so the test bedding activities can be mapped directly to the future architecture to be evaluated and employed by NOAA.

Goal 5: Maximize Effectiveness of HPC Solutions

Objective 3: Improve the user experience

NOAA will enhance tools and methods to provide users a consistent user experience across its integrated HPC resources. For example, NOAA will improve workload and software management practices, establish common reporting, and create common documentation and communication tools.

Not all climatologist, oceanographers, and scientists are computer experts. Maximizing effectiveness involves creating a common, yet flexible, user experience that enables NOAA researchers and practitioners to do their best work because of the technology, not in spite of it.

The systems integration team prototypes workflow alternatives and user interfaces on the testbed for NOAA HPCC evaluation. Best practices from other government HPC sites and and other data science can be replicated on the testbed for evaluation. For instance, High

Performance Data Analytics (HPDA) researchers have adopted interactive user environments called notebooks to manage, collaborate, document, and execute their applications on compute clusters. Jupyter, Zeppelin, and Beaker are open source notebook examples that may complement traditional HPC job management software and allow users to easily share and reuse their investigation setups and results. Another example might involve GUIs built to supplement scripting and command line interface to provide reproducible success. These types of activities create momentum to increase NOAA HPC capabilities and accessibility.

Conclusion

NOAA HPCC can realize significant benefits in their provision of operational and R&D services to NOAA missions by implementing cost-effective and targeted approach described in this white paper. Using a configurable HPC Testbed (TaaS), provided by a systems integration team possessing NOAA science expertise guiding the development, use and exploitation of the testbed in partnership with NOAA's HPC experts will lead to many combinations of the following benefits:

- Risk reduction
- New optimized software platforms
- New, novel, vetted, trialed hardware configurations
- Establishing performance metrics and expectations
- Incremental, evolutionary, and revolutionary results
- Shared knowledge within the NOAA HPC user community
- Lessons learned from industry, other government agencies
- Low risk capital expenditure planning
- Leveraged capabilities
- Reduced time to insight
- User community enhancements in training and capabilities development
- Access to the latest in hardware technologies

The success of this approach requires the testbed be:

- provided by an industry partner as a service, not a capital expenditure
- configurable including scale-up, scale-out and hybrid architectures
- dedicated to NOAA HPCC

The success of this approach requires the systems integration team:

- be hardware and software solutions architecture experts, knowledgeable in the evolution of HPC hardware and software
- be intimately knowledgeable of NOAA science and missions and their HPC requirements
- be vendor "agnostic" for solution specifications ensuring best value and most affordable solutions
- have access to a wide array of leading developers and manufacturers of HPC technologies
- be "skunk works" team leveraging small business affordability and "best value" that can be accessed through existing contract mechanisms

The HPC TaaS augments NOAA's mission with a focus on technology evaluation, vetting and insertion. It provides a low-cost budgetary solution to ensure early access to new

technology developments and implementations using the development, systems integration, customer application testing, and production test models.

In-depth discussion and coordination with NOAA scientists, engineers and thought leaders is critical to achieving these results and should including discussions about technologies that may not be clearly understood or may be completely unknown.

FedCentric and MERC offer the concept of a testbed as a service to provide a proven and affordable method to increase R&D exploitation of next generation HPC; and vetting of new, emerging and disruptive technologies for the NOAA HPC user community. This approach allows NOAA to more fully explore technology advances and techniques in a risk-free, cost-effective environment to better understand hardware and software approaches, that when implemented, greatly reduce R&D and operational risk.

For additional information about FedCentric and MERC capabilities contact:

David W. Miller
dmiller@merc-md.com
410-570-7937

Authors:

FedCentric Technologies Inc.
https://fedcentric.com/
FedCentric Technologies LLC – Service Disabled Veteran-Owned Small Business and HUBZone Company

Miller Engineering and Research Corporation
http://www.merc-md.com/
MERC is an Economically Disadvantaged Woman Owned Small Business