Special Report for NASA

# Exploring Options for a Bespoke Supercomputer Targeted for Weather and Climate Workloads

Bob Sorensen, Alex Norton, Earl Joseph, and Steve Conway
October 2019

## HYPERION RESEARCH FINDINGS

The primary focus of this study was to gather key insights, through a series of surveys with weather and climate users and potential HPC suppliers, on options available to NASA, and others, to develop a bespoke HPC system specifically targeted for weather/climate research.

Weather and climate survey respondents had a broad list of limitations with current HPCs as well their prospects for planned commercial HPC offerings. Key concerns centered on limitations in memory and storage latency and bandwidth, the lack of diversity in processor option/designs, the current reliance on GPUs that are not well suited to the current weather/climate community workloads, and the trend towards vendor specific interconnect options at the highest levels of computing.

These respondents provided a number of suggestions for future developments, noting that any project to build such a system would need to be a community effort. Currently such projects are seen as Center/Lab specific or based on "grand challenge" programs with relatively short time horizons. Another major theme centered on bringing AI, particularity machine learning capabilities, into the programming mix, but that opinion was not universally supported.

All of the vendors surveyed indicated that they would be interested in exploring the development of a bespoke HPC for NASA. However, most indicated that the option for building a special purpose one off system or even a series of such systems over time could not be economically justified regardless of the amount of NRE unless that particular design had value within the wider weather community or even across a number of complementary verticals. Other key points included the desire for NASA to conduct a robust pre competitive codesign effort with potential vendors, while reexamining and modernizing NASA's existing code base to take advantage of existing and anticipated progress in high end hardware and software.

Recommendations for next steps include a NASA initiative to more accurately assess the range of existing and planned workloads to better meet specific hardware and software requirements, compose benchmarks of mini application suites, testcases, or even full applications to help determine the various strengths and opportunities of any potential HPC design, and organize a pre-competitive codesign conference, or series of conferences, that brings together interested commercial vendors to discuss various options and opportunities. A key goal of these co-design efforts would be to discuss projections of advances in future HPC hardware and software that would better align with NASA requirements but that are in keeping with larger commercial trends.

*Note: This page is intentionally blank.*

## EXECUTIVE SUMMARY

The primary focus of this study was to obtain information, perceptions, and comparisons of the options available to NASA, and perhaps others, to develop a bespoke HPC system specifically targeted for weather/climate research.

The study was divided into two major phases:

1. The first phase centered on a series of interviews held with a wide range of expert researchers and users in the HPC based climate and weather community to gather their thoughts and insights on current and future operation requirements as well as the specific HPC hardware, software, and architectures needed to meet those workloads.
   - For phase one, 15 different weather/climate organizations in the US and overseas were surveyed including ECMWF, LANL, NOAA, ORNL, UCAR, and the University of Delaware.
2. Phase two consisted of taking the results of phase one to generate a list of key HPC specifications and requirements that formed the basis of a second survey of HPC suppliers and independent HPC designers to assess the challenges and opportunities of developing a bespoke HPC to meet the phase one requirements.
   - For phase two, HPC suppliers providing input included Cray Inc, Dell EMC, HPE, and IBM.
   - Vendor response consisted of either written input or conference calls and included participation by teams of in house HPC experts, many with strong backgrounds in the climate and weather community.

### Phase One Findings

Researchers and users in the HPC based climate and weather community had a broad list of limitations with current HPCs as well as for prospects commercial HPC offerings in the next few years. Key concerns centered on limitations in memory and storage latency and bandwidth, the lack of diversity in processor option/designs, the current reliance on GPUs that are not in well suited to the current weather/climate community workloads, and the trend towards vendor specific interconnect options at the highest levels of computing.

Other concerns specifically mentioned included:

- Inappropriate node/CPU designs that lead to low (near 1%) efficiency on current atmosphere and ocean models
- Overall high, and increasing, system capital and operating costs

Respondents offered a number of suggestions for future structural requirements needed to design and build an effective bespoke HPC. They centered on ensuring a well-defined, robust development effort that had applicability across a wide range of climate and weather workloads.

Specific responses noted that:

- The effort to build such a system would need to be a community program. Currently such programs are center/lab specific or based on “grand challenge” programs with relatively short time horizons.
- The system should not be a one off “science fair project”. A robust architectural roadmap would have to be provided that would persist over at least 2-3 system generations (10-12 years).

- The system (innovative or not) would have to be competitively procured. It would have to outperform existing platforms on weather benchmarks, it's TCO would need to be lower, including software porting investments, and it would need to be proven reliable.
- Any proposed solution would have to be configured as a true end-to-end system. For example, increasing the speed of the computer system by 10X without addressing the massive flow of data through the analysis and post processing system would be unacceptable.

Respondents offered some of their insights on future application requirements, noting that the development of a bespoke HPC program presents a major but uncertain opportunity to radically alter the way NASA applications are developed and run. Within such a paradigm shift, respondents noted that emulation in particular seems to have good potential to replace compute heavy physics parameterizations, and that any efforts going forward should reflect a strong shift in the community towards ensemble prediction systems vs simulation.

Another major theme centered on bringing AI, particularity machine learning capabilities, into the programming mix, but that opinion was not universally supported:

- Opportunities were seen to exist for mixing machine learning (ML) with traditional computation techniques likely running concurrently (asynchronously) with the simulation. However, such usage would require a shared memory mechanism for communication between the ML and traditional computation instances.
- Researchers can expect to see growth in ML and increasing requirements for supporting ML software frameworks, such as TensorFlow and Keras.
- GPUs might be (although are not necessarily) a significant platform for ML applications in atmospheric research.
- Other areas such as distributed data processing and observations' error correction were considered good candidates for AI as well.

## Phase Two Findings

For phase two of the project, vendor survey responses varied widely as each vendor brought their unique perspective, experience, and insights to the question set. However, a number of major themes emerged from the discussions.

All of the survey respondents indicated that they would be interested in exploring the development of a bespoke HPC for NASA. However, the level of commitment and the degree to which a system would use special purpose or customized hardware and software varied greatly. Most indicated that the option for building a special purpose one-off system, or even a series of such a system over time, could not be economically justified regardless of the amount of NRE unless that particular design had value within the wider weather community at a minimum or across a number of complementary verticals.

The major concern of respondents interested in working with NASA centered on NASA's need to modernize their codes to better capture improvements in both high end hardware and software. Almost all vendor respondents cited the critical need for code modernization and refactorization, while acknowledging that such a task will not be easy. Comments to that effect were quite emphatic.

- Any available NRE for such a bespoke system should be committed to code optimization. In two or three years, there should be interesting architectures suitable for NASAs workloads, but there will still be a significant gap in software capability.

Although respondents demonstrated a high level of interest in exploring the design of a NASA bespoke HPC, they indicated that determining exactly what such a system would look like could be challenging. Properly configured codesign efforts were universally seen as the best solution to this oft mentioned problem. Specific concerns about collecting NASA requirements were mentioned by many.

- NASA has a wide and varied range of workloads that would need to be accurately characterized and combined into a comprehensive set of hardware and software requirements.

## Recommendations for Next Steps

Based on the information and insights gathered in both phase one and phase two of this study, there are a number of recommendations for next steps for any NASA initiative to plan and procure a series of HPCs that can meet NASA computing requirements in a cost-effective manner.

NASA planners need to more accurately assess the range of existing and planned workloads to better provide specific hardware and software requirements:

- Key considerations here involve, among others, an analysis of processor core counts and memory bandwidth, the impact of special memory frameworks such as HBM and persistent memory, and the use of GPUs and other accelerators.
- Such efforts would be helped considerably by enlisting the insights of one or more major HPC vendor, all of which showed an interest in working with NASA on such a project.

The results of this comprehensive requirement analysis can then be used to compose benchmarks that consist of mini application suites, testcases, or even full applications to help determine the various strengths and opportunities of any HPC design under consideration. Such benchmarks will also be highly beneficial in providing potential vendors with insights on NASA high priority features and performance expectations. Benchmark efforts should also consider a number of new workloads that NASA likely will be increasingly turning to in the near future that include:

- AI based programming not only in machine learning application but also to accelerate the performance of traditional modern and simulation jobs
- Big data jobs, both batch and near real time, that could be a mix of large storage volumes along with high bandwidth data inputs from large arrays of edge connected sensors
- Hybrid on premise/cloud computing platforms to help absorb large fluctuations in workload, to test out and perhaps take advantage of new hardware opportunities quickly, and to tap into the growing base of cloud service provider resident HPC software

NASA, along with other interested weather/climate research organizations, could organize a pre-competitive codesign conference, or series of conferences, that bring together interested commercial vendors to discuss options and opportunities based on the analysis described above. A key goal of these co-design conferences would be to build long term relationships with key vendors to realize a strategic multi generation procurement plan for HPCs instead of targeting a onetime procurement. Other key action items for such meetings could include:

- Working with vendors, many that have significant expertise in profiling and tuning applications for future technologies, to develop a robust methodology for crafting performance predictions from NASA benchmarks.
- Generating relevant projections of advances in future HPC hardware and software that aligns with NASA requirements but that are in keeping with larger commercial trends.

- Enlisting vendor insights and expertise on the best ways for NASA to begun refactoring and modernized their current software base to best take advantage of existing and planned commercial technology and HPC products.
- Working with vendors to ensure that any new NASA software developments are in keeping with the latest trends in HPC application development and related software development tools.
- Seeking opportunities to help leverage any NASA hardware or software developments into the larger weather/climate HPC ecosystems that can help vendors justify the development costs of any targeted hardware or software.
- Identifying and including other HPC verticals that share some common mission elements with key NASA programs to generate wider interest in NASA relevant HPC developments, build economies of scale, and offer additional financial motivations for commercial HPC vendor participation.
- Engaging the growing cadre of cloud service providers who increasingly are on the forefront of new HPC hardware and software techniques, especially in the areas of modular software development and deployment.

Ultimately, respondents from both phase one and phase two of the survey were optimistic about NASA's efforts to confront their existing hardware and software limitations. They indicated that many of the suggestions offered here stand a good chance of ensuring that NASA has access to the kind of HPCs that will be needed to continue their world class research in key climate and weather research.

## TABLE OF CONTENTS

## LIST OF TABLES

## LIST OF FIGURES

## IN THIS STUDY

The primary focus of this study was to gather key insights, through a series of surveys with weather and climate users and potential HPC suppliers, on options available to NASA, and perhaps others, to develop a bespoke HPC system specifically targeted for weather/climate research.

The study was divided into two major phases: the first phase centered on a series of interviews held with a wide range of expert researchers and users in the HPC based climate and weather community to gather their thoughts and insights on current and future operation requirements as well as the specific HPC hardware, software, and architectures needed to meet those weather and climate related workloads. Phase two consisted of taking the results of phase one to generate a list of key HPC specifications that formed the basis of a second survey, this time including HPC suppliers and independent HPC designers to assess the challenges and opportunities of developing a bespoke HPC to meet the phase one requirements.

## Background

NASA's earth system models have unique high performance computing (HPC) requirements, which can differ from standard industry offerings. Moreover, the gap between vendor HPC solutions and earth system models has been growing, such that these models can exploit less and less of the peak computing capability of current HPC systems.

Hyperion Research was asked to study the feasibility, including but not limited to budget, manpower, and expertise, and the potential for a shift in HPC design to better meet the requirements for NASA earth system models. Efforts center on the potential for the creation of a bespoke supercomputer targeted for weather and climate workloads. Key baseline assumptions for the study included concentrating on US HPC vendors, excluding exotic hardware technologies such as quantum computing, minimizing concerns with long-term archival storage, and deemphasizing solutions based primarily on cloud architectures.

## Methodology

Phase one surveyed users and was designed to explore and gather insights into the operational and related HPC requirement needed to support the weather and climate research community. As such, the survey sought to delve into a broad range of weather and climate related current and future operational requirements along with questions that sought to map those future operation requirements in HPC hardware and software specifications.

## Phase One: Weather/Climate Operational Requirements and Related Questions

The following set of questions were used as the basis of the phase one survey.

### Current and Future Operational Requirements

1. What are the current operational requirements within your facility for weather prediction (number of models, runs per day per model, model per day runs times etc.) and seasonal and decadal predictions (model/grid size, run times)?

2. What are future (5-10 year) operational requirements for the workloads described above and what will be some of the key changes in operational requirements between current and future systems (e.g., model complexity and resolution, adding analytics to simulation)?
3. What are the key operational requirements needed to drive improved prediction capabilities (e.g. multiscale physics, active chemistry, increased ensemble members, higher grid resolution, coupled models, time to solution)?
4. What are some of the key limitations in the current and planned base of commercial HPC developments and systems that will most adversely affect your ability to meet future operational requirements?
5. What are some of the key architectural, hardware, or software changes that you are willing to accept (or would like to see from vendors) to help ensure the future availability of HPCs that can meet your future operation requirements?
   - Can you discuss how you would make the decision to move to a new, innovative solution that could require you to make major changes to your applications?
   - What specific software or other elements would you be willing to change to realize a performance gain in your operational capability, e.g. tradeoffs such as moving to a different programming language (C, C++, Python or other), running in a vitalized computer, running advanced analytics, using a hybrid system consisting of traditional CPUs and offload devices, or a 100% GPU system.
     - Likewise, to what extent would you be willing to refactor your software (100%, 80%, 60%, 40%, 20%, or even 0%) or adopt new algorithms to gain performance and/or accuracy?
     - What performance improvement would you need to do any of these?
   - What would be the key issues and concerns to be covered in a formal co-design process for a new custom HPC for NWP?

### Mapping Future Operational Requirements into HPC Specifications

6. How will the above requirements have an impact on your weather applications and the HPC systems needed to best support them?
7. Has your organization already specified your future HPC requirements? Does a formal plan exist?
8. What are the key technical architectural specifications needed in an HPC system to meet your future (3-5 years) operational requirements discussed above?
   - Consider key system architectural and related performance requirements including total core count, types of processors, processor to accelerator ratios, node and board configurations, total power consumption, floor space requirements, cooling specifications, total memory and storage size, byte/flop or related benchmarks (LINPACK, HPCG, STREAMs as well as kernel, dwarf or mini benchmarks, etc.), typical job and user mix, use of mixed precision workloads, others.
   - What are the overall application performance improvements (i.e. speedups of 5%, 10%, 50%, etc.) being sought over systems that are currently on the market and can any insights be made as to the acceptable range of tradeoffs in price for added performance?
9. What are the key technical hardware specifications needed in an HPC system needed to meet the future (3-5 years) operational requirements discussed above? In addition, please describe how familiar you are with these technologies and the extent you have looked at them in the past, are doing work with them now, or planning on doing work with them soon.

- Processor(s): type, performance requirements, key features.
- Accelerator(s) (GPU, FPGA etc.): type, performance requirements, key features
- Memory (including cache): DDRx, high bandwidth memory, NVMe DIMMs, etc. size, type, and bandwidth, latency
- Interconnect: topography, bandwidth, latency, critical communication patterns (broadcast, scatter, gather, etc.)
- Storage: solid state, spinning disk for scratch, campaign and archival storage, capacity, bandwidth, latency etc.
- Other items?

10. What are the key technical software specifications needed in an HPC system needed to meet the future (3-5 years) operational requirements discussed above?
    - Key elements of the software stack needed to ensure ease of use in porting existing the code base
    - Key elements of the software stack needed to ensure ease of use for development of new software
    - Key elements of software stack needed to ensure interoperability among various critical codes across the weather/climate modeling community
    - Key performance characteristics of select software elements including, OS, job schedulers, compliers, middleware, software development tools, etc.
    - Key application workloads and their associated hardware requirements
    - Expectations of new application composition based on new software areas such as big data analytics, deep learning, and hybrid on-prem/cloud configuration.
    - Expectations about bitwise reproducibility and rounding errors that occur when optimizations are made, and their current and future needs for single, double, or half or mixed precision variables
    - Other items?

A key priority of the phase one survey was to cast as wide a net as possible to include a broad range of both operation and research personnel across the weather and climate community. Working in concert with the sponsor, invitations were sent out to identified experts in the HPC related climate and weather community. In addition, many invitees provided suggestions for additional experts to contact. In total, 56 climate and weather experts were invited to take part in the study from North America and Europe.

The range of invited experts included those from: US weather agencies including NOAA and NCAR, DOD factions from the Navy, Army, and Air Force, universities including Illinois, Texas, Colorado, Ohio State, Princeton, the University of Delaware, UCAR, NSF, DOE labs from LANL, ORNL, and LLNL, foreign experts from ECMWF, Oxford and CanMet, and NASA.

Ultimately, 10 different weather/climate organizations in the US and overseas were surveyed for their top considerations to meet their HPC based computation requirements in the next 3-5 years representing:

- ECMWF
- LANL
- NASA AMES

- Naval Post Graduate School
- NOAA
- ORNL
- Oxford
- UCAR
- University of Delaware
- University of Washington Applied Physics Laboratory

Many responses consisted of the collaborate efforts within a single organization and contained combined answers from multiple climate/weather experts, sometimes across organizational boundaries. Some survey respondents included RFPs and other technical materials outlining their current and future procurement requirements.

- Within the set of respondents, both operational (sort term weather forecasting) and research (long term climate prediction) were represented, but most organizations surveyed were predominately research vice operational facilities.

Prior to initiating phase two (surveying the various HPC vendors and HPC architects) the results of phase one were used to generate a two page specification sheet that summarized the various HPC software and hardware requirements. That specification sheet is presented later in this report. The specification sheet consists of a list of the hardware, software, and architectural elements that were mentioned most often or most emphatically by the set of respondents. No attempt was made to deconflict or generalize the respondent inputs beyond what was necessary for the sake of brevity and clarity.

## Phase Two: Turning NASA/Weather Requirements into Supplier Recommendations

Phase two consisted of taking the bespoke HPC specification sheet described above along with a survey set for HPC vendors that sought to determine how they could best respond to the results from phase one.

The following set of questions were used as the basis of the phase two survey.

1. In the form of a pre-bid scenario, and based on similar experiences in past procurements:
   - If you were given a medium amount of NRE, e.g. tens of millions of dollars, how would you best address the above requirements?
   - If you were given substantial NRE/R&D, e.g. $100 million, how would you best address the above requirements?
   - If you were given substantial NRE/R&D, e.g. $100-200 million, plus an agreement to purchase at least $500 million in systems, how would you best address the above requirements?
2. What do you see as the main limiting factors in responding to these requirements? Including, but not limited to:
   - Overall ecosystem and talent
   - Legacy applications (e.g. cost and talent for porting to the new design)
   - Algorithms
   - Programming languages

- Supporting software framework
- The time to implement all changes, including rewriting software and applications, and verification

3. What do you see as the "low hanging fruit"?
4. How would you prefer to address these requirements?
5. What would be the cost and technical feasibility in providing upgrades for these systems over time, and how long would that commitment be in force?
6. What would be the approach you would use to determine the expected performance improvements in NASA/weather community workloads because of the options outlined above?
7. Would you be able to sell HPCs meeting these requirements into market sectors beyond weather/climate? To what extent?
8. Other approaches

As with phase one, the goal of phase two was to cast as wide as net as possible for respondents. The supplier survey was sent to a range of potential suppliers of HPC in one form or another and included: AWS, Cray, Dell EMC, Google, HPE, Intel, Lenovo, Microsoft, NVIDIA, Oracle, and Penguin.

Survey respondents were told that their specific answers would remain confidential and folded into larger conclusions and reported with attribution in the final report. As such, this report lists the suppliers surveyed but does not make available any specific respondents' names or affiliations. They were also guaranteed that any input provided would not carry any obligation or future commitment to NASA or any other agency. Finally, the survey was sent to a number of cloud service providers with the caveat that NASA was not looking specifically for a complete cloud-based solution but would consider a hybrid on-premise/cloud solution.

For phase two, the following HPC suppliers provided detailed inputs:

- Cray Inc.
- Dell EMC
- HPE
- IBM

In addition, a number of respected independent HPC experts also provided supplementary background and insights.

## KEY FINDINGS OF WEATHER/CLIMATE EXPERT DISCUSSIONS

The following sections contain summaries of the phase one survey results that include coverage in current and future operational and research capabilities, limitations of current and planned commercial HPCs, and future architectural, hardware, and software considerations.

### Phase One Results: Current and Future Operation and Research Capabilities

The following are highlights of survey results concerning current operational and research capabilities as well as future operational and research requirements.

Sample current operational capabilities cited included:

- GFS (soon FV3GFS) 4Xday for 10 days at 28km resolution
- North American Model 4X/day for 84 hours at 12km resolution
- High Resolution Rapid Refresh model hourly (24X/day) for 36 hours
- Short term research on next generation NWS capabilities (FV3, FMS, MOM6)

Sample current research capabilities cited included:

- Earth system models
  - 30 simulated years per day (11,000x faster than reality) @ ~100 km resolution (all components)
  - 5 simulated years per day (1800x faster than reality) @ 25 km (atmosphere & land surface) + 100 km resolution (ocean and sea ice)
- Climate simulations
  - Between 1-10 simulated years per wall clock day for all resolutions/configurations with high resolution (25km atm, 10km ocean)
  - Typically ending up at the lower end of that range while coarser resolutions target a higher throughput and larger ensemble

Sample future operational requirements cited included:

- For climate studies, 5-10 simulation years per computational day considered minimum
- Grid resolutions for cloud resolving (~1km) or cloud permitting (~3km) --> 1,000-10,000 in a decade
- Running 3KM global models operationally, and regional models at 1KM or finer scales

## FIGURE 1

### Future Research Requirements: Examples and Configurations

| | *Example Application* | *Configuration* |
|---|---|---|
| **Weather** | Tropical Cyclones | 3km refined mesh, coupled ocean, forecasts |
| **Climate** | Hydrologic Extremes | 3km refined mesh, forecast and climate simulations |
| **Polar** | Arctic Prediction | 3km refined mesh, coupled ocean, land, sea ice, land ice. Forecast and climate simulations |
| **Geospace** | Space Weather Prediction | 25km global atmosphere to the ionosphere, forecast. |
| **Chemistry** | Urban/Regional Air Quality Prediction | Urban: <1km regional forecast. Regional: 3km refined global mesh, climate and forecast |

Source: Hyperion Research, 2019

## Phase One Results: Limitations of Current and Planned Commercial HPCs

Survey respondents had a broad list of limitations with current HPCs as well as for their expectations for planned commercial HPC offerings in the next few years. Key concerns centered on limitations in memory and storage latency and bandwidth, the lack of diversity in processor option/designs, the current reliance on GPUs that are not well suited to the current weather/climate community workloads, and the trend towards vendor specific interconnect options at the highest levels of computing. Other concerns included:

- Inappropriate node/CPU designs that lead to low (near 1%) efficiency on current atmosphere and ocean models
- Overall high, and increasing, system capital and operating costs

## Phase One Results: Future Climate and Weather Workload Considerations

Respondents offered a number of suggestions for future organizational requirements needed to design and build an effective bespoke HPC. They centered on ensuring a well-defined, robust development effort that had applicability across a wide range of climate and weather workloads. Specific responses noted that:

- The effort to build such a system would need to be a community project. Right now, such projects are Center/Lab specific or based on grand challenge programs with relatively short time horizons.
- The system should not be a one off science fair project. A robust architectural roadmap would have to be provided that would persist over at least 2-3 system generations (10-12 years).

- The system would have to be competitively procured. It would have to outperform existing platforms on benchmarks, it's TCO would need to be lower including software porting investments, and it would need to be proven reliable.
- Protecting the investment in new modeling software would have to be addressed in the overall plan. This suggests developing or adopting a domain specific language (DSL), such as Kokkos or directive based systems, such as OpenMP 4.X. Any DSL would need to have cross center buy in/support.
- Any proposed solution must be configured as a true end-to-end system. For example, increasing the speed of the computer system by 10x without addressing the massive flow of data through the analysis and post-processing system would be seen as unacceptable.

Most respondents were optimistic that future performance gains were possible with straight line extensions of existing technology bases but were concerned that algorithm and related software questions could slow down the process. One respondent noted that a well-designed forecast model could theoretically be written capable of achieving 10% of peak processor performance with perhaps 80% scaling efficiency:

- It would emphasize uniform calculations and grids where calculations could be done in parallel for most of the modeling dynamics.
- Physics based calculations would use a mix of AI and traditional compute.
- An increase in the price performance is needed, but more efficient processors would mean lower power consumption providing the basis for some tradeoff.

Others noted that for a custom solution, potential users would need to see a substantial performance improvement (>10x) over a general purpose machine and long term support for both the architectural design and the required programming model/software stack. One respondent stressed that some of the biggest changes always come from algorithmic changes and mapping algorithms to underlying hardware properties, so exploration of new algorithms for many components needs to be continuous.

Noteworthy cautionary comments suggested that the performance/workload data necessary to specify memory and cache characteristics for future systems is a complex and often risky exercise, and that the biggest impediment to accepting new hardware types or software is the plethora of legacy systems and training of scientists.

## Phase One Results: Future Architectural Considerations

Survey respondents had a long, ambitious, and somewhat diverse list of architectural considerations for any future bespoke HPC. Specific callouts noted that:

- Heterogeneous systems should consist of mostly CPU only nodes with a smaller fraction of CPU-GPU nodes (for example, 4 GPUs per node,) otherwise the interconnect can become a bottleneck.
- Both the near term and long term trends in the HPC sector will continue to add more parallelism within the node, currently with a GPU like approach.
    - The current difficulty here is that to achieve the throughput numbers needed, systems will be pushed out to the strong scaling limit where there is not enough work per node to utilize the on node parallelism within a GPU, so they lack the computational intensity at this strong scaling limit.

- Better use can be made of on node parallelism at lower node counts, but at lower absolute throughput (though running ensembles in this mode can provide a more efficient use of an allocation).
- Node architectures need to have specific, likely static resources allocated for non computation operations like running the O/S, handling communications, etc.
- A key gating performance factor will be memory bandwidth per core, a measure that is at best flat, or in some cases declining, for DRAM based CPU systems.
  - Vendors are adding more DDR channels to sockets, but they're adding cores just as fast.
  - However, the adoption of HBM memory systems (the principal reason GPUs are faster than CPUs right now) over the next 5 years could change this calculus.
  - The application base's dependence on memory bandwidth is such that one can predict the performance of well implemented PDE solvers on systems just based on the bandwidth.
- Processors will likely have hundreds to thousands of compute cores that will require more fine-grain parallelism in the applications.
  - This will drive requirements for more memory bandwidth per processor, lower latency communications, parallel I/O with high bandwidth storage and file systems.
  - Future compute nodes will likely contain 10-50 sockets (CPU) or complex architectures like the NVIDIA DGX-2 system with 16 GPU, 300GB/s memory bandwidth.
- Data analysis will be much more I/O intensive, including a ~100x increase in the number of jobs handled by the job scheduler and will increase the importance of model initialization overhead to throughput performance.
- Application scalability (which is largely on the model developers to achieve) is another key gating factor for performance improvements.
  - To support scalability, the interconnect needs to be balanced relative to on node memory bandwidth and the MPI tasks must be mapped to minimize off node communications.
  - Current figure of merit for a balanced system is ~12:1 (on node vs off node BW).
- Traditionally, double precision is used in climate and single for weather.
  - Some work has been done to develop mixed precision climate capabilities but not much has not been put into production.
  - However, there is no identified single case (yet) where half precision is useful for anything related to weather and climate even in machine learning.

## Phase One Results: Future Performance Considerations

Survey respondents offered up a wide range of specific performance requirements. Overall, performance goals for future performance were varied, but two key themes emerged: aggressive performance gains will be needed to meet future operational requirements, and such gains are needed to generate interest among the potential user base for bespoke systems.

Specific performance improvements included:

- A minimum of 50% "computing cost neutral" speed up in combination with enhanced scalability (i.e. fixed hardware, power, cooling, costs)
- Multiple 10x improvement to get to global cloud resolving scales, so more interest in substantial improvements (2-10x) and not just incremental improvements (10-50%)
- A goal of between 1,000x to 10,000x increase in computational capability within 10 years

Respondents noted that to attract interest and be truly disruptive, a new HPC would need to be substantially faster than contemporaneous systems when it is ready for production service. Experience in the past suggests the following performance speed up heuristic for the fraction of model developers to be interested:

- 2x: 10% interested
- 4x: 20% interested
- 10x: 50% interested
- 100x: 90% interested

Others suggested that assuming a five year time frame to get a new system and refactor code, a new system would need at least 32x the performance of current systems for a broad swath of the community to really get wide ranging interest in porting to a new system.

Despite an emphasis on HPC performance increases from survey respondents, there was only minimal mention of targeted benchmarks to gauge performance on existing and planned key applications such as the development of a suite of mini benchmarks or using HPCG and Streams as a proxy for projecting throughput of systems.

Respondents also provided some key inputs on various hardware elements of a desired bespoke system.

Processors

- Ideally target one core capable of an exaflops, targeting a system that will hold O(10^5 - 10^6) cores over the next 5 years, capable of O(10s of petaflops)
- Main processor considerations cited included price performance (say HPCG gigaflops/$) and power performance (HPCG gigaflops/watt). As codes are memory bandwidth limited, favor SKUs with fewer processors and midrange clock speeds
- Based on current processor technology where processors are not getting faster, over 300K CPU cores (10,000 GPUs) will be required to run 3KM models operationally

GPUs

- Lack of robust OpenACC implementations for GPUs make the first step “try before you buy” difficult to get past
- Experience with GPUs indicates that most likely the dynamical cores will have to be completely rewritten, or at least strongly refactored
- GPUs will be needed for compute (physical modelling), machine learning, and visualization
- Onboard GPU system memory is not particularly constraining for applications
- GPUs are attractive to infrastructure providers because of the 4x HPCG GFlops/watt. However, GPUs were difficult to program until very recently
- GPUs with HBM outperform CPUs with DDR there’s a big difference between 100 GB/sec and 900 GB/sec. It appears that DDR performance is playing out (BW/watt) and HBM is the long term way to go

FPGAs

- FPGAs are at least a decade away from being a viable accelerator technology for widespread use in the atmospheric research and forecasting
- No experience with FPGAs at this time

- FPGAs are not a factor in our current planning, because of their unprogrammability

Memory and Storage

- Processor-in-memory capability offer the ability to perform simple operations in a memory tier without needing to move data elements through the various cache levels within a processor, i.e. zero out data elements, scale data elements by constant value or add a constant value to data elements, etc.
- In CPU systems with multiple DDR memory channels, on node memory capacity has ceased to be a constraint it is hard to configure a 2-socket node with less than 2 GB/core
- Generally, populate the available DDR slots to maximize bandwidth
- NVMe DIMMS might add capabilities for data intensive applications, but have not seen pricing or tested it in a system
- Node resident SSDs are needed for analysis nodes
- Migration of data between storage tiers (NVMe, HDD, campaign storage, tape, and cloud) in this system must be thought through carefully

Interconnect

- On capable, balanced interconnects (e.g. Cray Aries, precisely designed and implemented IB, etc.), models are not interconnect-limited but, rather load balance-limited
  - Real interconnect limitation is more defined by the limits of the message passing library implementations than the hardware characteristic elements
  - Because most numerical methods are local, and the message passing dominated by MPI-send/recv based halo exchange, the system interconnect can be often oversubscribed
- Fat nodes connected with high bandwidth inter node communications and robust file system performance are needed
  - I/O would require ~20-30TB of output for a single 3KM model run (10 day forecast, output every 3 hours of forecast time)
  - Other models executing simultaneously (regional models, ensembles, etc.) will double the I/O requirement
- Scatter/gather operations are critical for I/O ops, nearest neighbor communications for finite volume, and finite element or spectral element models
  - Some low b/w broadcast mostly for input, and minimal collectives (e.g. max, min, sum, etc.) are needed, but these should be avoided in the algorithms as much as possible

Respondents offered a wide list of software considerations as well. Two key themes were mentioned by a number of respondents: a wide range of existing applications need to be rewritten to take advantage of new hardware and software trends, and there needs to be a strong base of software development tools to support those activities. Many expect that this process will be complex and met with some resistance.

Existing Applications

- Applications need to be completely rewritten.
- Most modeling software is Fortran90+, but there is a healthy amount of C and C++ code as well.
- In general, existing applications are poorly architected, and largely designed for traditional, low core count CPU processors.

- Going forward, portability is key.
- With GPUs the refactoring/model development time required to develop these applications is non-trivial.
- Researchers would relax requirements like bitwise reproducibility (in favor of newer statistical reproducibility techniques) and other current practices that could otherwise hamper advancing the capabilities on more revolutionary architectures.
- Most codes may require applications be written in another language or a combination of traditional Fortran (good for performance) and an object oriented language like C++ (good for programmability, maintainability).

Development Tools and Middleware

- A reliable, robust optimizing compiler is requirement zero.
- Python is taking off as the scripting language of choice among early career scientists.
- The importance of having efficient, accurate, vectorizing math libraries cannot be over emphasized.
- Need reliable performance analysis and debugging toolsets, assuming that existing code can be ported
- Better Support for MPI and OpenMP is required.
- Minimizing O/S interference with application computation and communication is essential.
- Must have capable job handlers.
- Any toolsets must handle concurrent model components easily.
- MPI needs to support accelerator coprocessors if those are part of the system design.
- Application profilers and debuggers capable of profiling parallel, multi node applications (i.e. ones with MPI support) are essential.

## Phase One Results: Highlights on Future Applications

Respondents offered some of their insights on future application requirements, noting that the development of a bespoke HPC program presents a major but uncertain opportunity to completely change the way applications are developed and run. Within such a paradigm shift, respondents noted that emulation in particular seems to have good potential to replace compute heavy physics parameterizations, and that any efforts going forward should reflect a strong shift in the community towards ensemble prediction systems vs simulation.

Another major theme centered on bringing AI, particularity machine learning capabilities, into the programming mix, but that opinion was not universally supported. Notable comments included:

- Opportunities exist for mixing machine learning with traditional computation techniques likely running concurrently (asynchronously) with the simulation.
  - However, such usage would require shared memory mechanism for communication between the ML and traditional computation instances.
- Researchers can expect to see growth in ML and increasing requirements for software frameworks supporting machine learning, such as TensorFlow and Keras.
- GPUs might be (although are not necessarily) a significant platform for machine learning applications in atmospheric research.
- Other areas such as distributed data processing and observations error correction, are good candidates for AI.

Regarding the question of bitwise reproducibility, respondents noted that at least one mode of execution where bitwise reproducibility can reasonably be expected is preferred. Additional inputs on this issue:

- Reproducibility of a given layout configuration (i.e. same source code, compiler + options, libraries, MPI layouts, threading options, model inputs, etc.) could be used to identify some system problems (e.g. various node and communication network issues).
- Rounding error analysis for various instruction schedules created by a compiler would allow for estimations of impacts without the need to run a series of perturbed ensembles.
- Bitwise reproducibility requirements exist in the sense of repeated runs of the same numerical experiment on identical hardware, runtime, compiler, etc.
  - This is important for debugging. In modeling, these requirements are relaxed, although would like bit for bit reproducibility for restart.
- At least one respondent noted that they would relax requirements like bitwise reproducibility (in favor of newer statistical reproducibility techniques) and other current practices that could hamper advancing the capabilities on more revolutionary architectures.

Finally, the last central theme covered the need and relative complexity of refactoring existing software. One respondent indicated that it would require at least a 2x performance improvement for a full refactoring but would consider refactoring individual routines and sub components to achieve 10-30% performance improvement. Another respondent provided the following analysis to estimate the cost of refactoring 100,000 lines code with existing hardware options, using the estimation methodology of the COCOMO model. The EAF roughly describes the level of difficulty. The fraction of code touched is based on experience for CPU/GPUs but is more speculative for FPGAs.

**Table 1**

**Estimating Code Refactoring Costs**

| | EAF | Fraction of Code (%) | Refactoring Cost ($M) |
|---|---|---|---|
| CPU | 1.12 | 5 | 0.28 |
| GPU | 1.55 | 20 | 1.67 |
| FPGA | 3.5 | 60 | 11.9 |

Source: Hyperion Research 2019

## KEY FINDINGS OF PHASE TWO: HPC VENDOR/EXPERTS DISCUSSIONS

Phase two of this study consisted of taking the results of phase one to generate a list of key HPC specifications that was combined with a list of specific questions to form the basis of a second survey, this time of HPC suppliers and independent HPC designers to assess the challenges and opportunities of developing a bespoke HPC to meet the phase one requirements.

As with phase one, the goal of phase two was to cast as wide as net as possible for respondents. The supplier survey was sent to a range of potential suppliers of HPC in one form or another and included AWS, Cray, Dell EMC, Google, HPE, Intel, Lenovo, Microsoft, NVIDIA, Oracle, and Penguin. The survey was also sent to a number of cloud service providers with the caveat that NASA was not looking specifically for a complete cloud-based solution but would consider a hybrid on premise/cloud solution.

For phase two, the following HPC suppliers provided detailed input:

- Cray Inc.
- Dell EMC
- HPE
- IBM

In addition, a number of respected independent HPC architects were contacted and their inputs are included without attribution in the survey results as well.

### Bespoke HPC Spec Sheet (From Phase 1)

Along with the specific set of phase two survey questions outlined in the methodology section, potential respondents were also sent a two page specification sheet that summarized the various HPC software and hardware requirements gleaned from the phase one survey.

- It is important to note that the specification sheet consists of a list of the major elements that were mentioned most often or most emphatically by the set of phase one respondents.
- No attempt was made to deconflict or generalize those responses beyond what was necessary for the sake of brevity and clarity.

Based on the phase one survey described above, the following specifications were deemed the most important.

<u>Processor</u>

- Consider range of processor options (x86, ARM, Power, etc.)
- Single, double precision is adequate, mixed precision not critical
- Support for hardware scatter/gather
- High memory bandwidth especially memory bandwidth per core
- Better channel per core ratio
- Factor of 5X best current x86 to justify software rewrites

GPU/Accelerators

- Heterogeneous system with mostly CPU only nodes and with a smaller fraction of CPU-GPU nodes that does not overwhelm the interconnect
- FPGAs not seen as near term solution

Memory/Storage

- Well organized cache hierarchy
- Current workflow uses lots of small block I/O, so need tiered solid state storage
- Migration of data between storage tiers (NVMe, HDD, campaign storage, tape, and cloud) in this system must be thought through carefully

Interconnect

- Limited use of vendor specific interconnects
- High I/O bandwidth
- On capable, balanced interconnects, models are not interconnect-limited but rather, load-balance limited

Architecture

- Strong scaling across entire system. Currently lack computational intensity at the strong scaling limit – not enough work per node to justify GPUs, etc.
- Interconnect needs to be balanced relative to on node memory bandwidth and the MPI tasks must be mapped to minimize off node communications
- As codes are memory bandwidth limited, favor SKUs with fewer processors and midrange clock speeds

Software

- Data analysis will be much more I/O intensive, inducing a ~100x increase in the number of jobs handled by the job scheduler
- Long term support for programming model and software stack
- A reliable, robust optimizing compiler is requirement zero
- Application profilers and debuggers capable of profiling parallel, multi node applications (i.e. ones with MPI support) are essential
- Portability is key
- Requirements for on prem/off prem cloud interoperability

System Level Considerations

- High flops/byte ratio (well above current industry averages)
- Provide end to end design. For example, increasing the speed of the computer system by 10x without addressing the massive flow of data through the analysis and post-processing system would have to be addressed
- Current figure of merit for a balanced system is ~12:1 (on node vs off node BW)
- A robust architectural roadmap would have to be provided that would persist over at least 2-3 system generations (10-12 years)

Benchmarks

- No significant peak or LINPACK numbers mentioned
- Need multiple 10x improvement meet requirements, so more interested in substantial improvements (2-10x) and not just incremental improvements (10-50%) over counterpart COTS HPCs
- HPCG and Streams are the proxy for projecting throughput of systems
- Assume five year time frame to get new system, and refactor code needs 32x the performance of current systems for a broad swath of the community to really get broad interest in porting to a new system

## Phase Two Results: Vendors' Perspectives: A Range of Opinions and Suggestions

Vendor survey responses were widely varied as each vendor or related HPC expert brought their own unique perspective, experience, and insights to the question set. Key points made by almost all vendors included: a solid interest in pursuing such an activity, albeit each with a unique perspective on the effort and their response, a strong interest in conducting a robust pre competitive codesign effort to better match NASA workloads with HPC capabilities, and a widely held acceptance that NASA needed to reexamine their existing code base to take advantage of existing and anticipated progress in high end hardware and software.

## Phase Two Results: High Interest in Exploring a Flexible Bespoke HPC Design

All of the survey respondents indicated that they would be interested in exploring the development of some form of a bespoke HPC for NASA. However, the level of commitment and the degree to which a system would use special purpose or customized hardware and software varied greatly. Most indicated that the option for building a special purpose one-off system or even a series of such system over time could not be economically justified regardless of the amount of NRE that accompanied such an effort unless that particular design had value within the wider weather community at a minimum or across a number of complementary verticals.

- Some respondents indicate that no amount of NRE would justify the development of a one-off systems as the opportunity costs would always be too high.

Not all comments were, however, as negative:

- Most respondents noted that with increasing user demand for systems that can support a suite of diverse workloads, such as big data and machine learning jobs along with traditional modeling and simulation, the options to bring together the exact hardware and software configuration to meet specific workloads have never been better.
- Most respondents indicated that they would be willing to work with NASA to develop a framework that met the bulk of NASA requirements but that would also have general applicability to a wide base of potential users.
- Respondents indicated that such efforts to blend NASA and more general market requirements would in the long run benefit NASA procurements as they would be much more in step with more varied, timelier, and ultimately less costly HPC mainstream.

- One respondent noted that any efforts to develop system software that provided data services to manage multiple tiers of storage with different access and performance characteristics would be useful in a wide variety of market segments.

A key point made repeatedly by respondents was that flexibility in HPC design through either a reliance on open systems components or carefully chosen proprietary designs could significantly ease the burden for vendors attempting to craft a system that could meet NASA requirements.

Specific insights on HPC design flexibly included:

- Hardware flexibility can be best achieved through the right choice of memory fabric so that any processing element can be used and scaled as long as it can speak to the fabric. Gen-Z was cited as a promising open system bus that can support a wide range of memory access modalities.
- Processor selection was viewed as a topic open for discussion. Intel and AMD X86, ARM, and related ARM vector processors such as those available from Fujitsu were all named as potential processor choices.
  - ARM presents some attractive options for a custom design that would allow for targeted core count to memory bandwidth ratios.
- Regardless of processor choice, a careful selection of processor SKU would enable the tuning of memory and interconnect bandwidth per core to maximize these when customer workloads or benchmarks show that these are performance limiters.
  - The careful selection of processor SKUs to provide needed memory bandwidth would enable NASA to exploit new, rapidly emerging storage technologies including high bandwidth memory (HBM), on package memory, and nonvolatile memory as they become more available.
- One vendor would consider partnering with a processor company to optimize their design to include on chip (HBM) but only if there was substantial NRE available to support such work.

Finally, almost all respondents stressed the need for further exploration into the use of GPUs and other accelerators into any future bespoke HPC. Their potential for significant performance gains in some traditional modeling and simulation jobs, as well as the promise of bringing AI techniques into some of NASA jobs, was seen as simply too great for GPUs and accelerators not to be considered.

## Phase Two Results: Widespread Concerns About Existing Code Base -- Modernization is Key

By far, the major concern of respondents interested in working with NASA centered on NASA's need to modernize their codes to better capture improvements in both high end hardware and software. Almost all respondents cited the critical need for code modernization and refactorization, while acknowledging that such a task will not be easy. Comments to that effect were quite emphatic.

- Any available NRE for a bespoke system should be committed to code optimization.
  - In two or three years, there should be interesting architectures suitable for NASAs workloads, but there will still be a significant gap in software capability.
- There are a lot of codes out there that have not been refactored and as many of the people who maintain these codes retire or move on, many of them likely will never be refactored.

- Ultimately, the barrier to achieving higher performance will not be due to hardware and the associated slowdown of Moore's Law, but instead due to the resistance to refactoring code. Indeed, the big gains in performance will increasingly come from software.
- The most effective use of any NRE for this system would be updating the programming model to take advantage of higher core count processors.
- Of all limiting factors in such an effort, the cost, time and talent needed to recast legacy code to use new technology can be the overarching issue.
- There will only be escalating costs to support an increasingly obsolete programming paradigm, and the first step to addressing that reality is to 'bite the bullet' and embrace a comprehensive software rewrite plan.
- Spend the first set of dollars on software development, not hardware.

Despite widespread vendor concern, respondents did offer some positive comments indicating that software modernization efforts could lead to interesting new algorithms.

Other comments included:

- Cloud service providers were cited as leading the way in driving continual code refactoring as they provide new software releases on a regular basis, allowing for the quick harnessing of any new technology gains, and NASA could learn much from their processes.
- Some respondents noted that code modernization efforts would reduce the need to hire or replace the declining base of uber legacy code experts such as those capable in Fortran, with a range of programmers with skills in newer programming paradigms.
  - Warnings were issued however, that hiring such experts can be expensive as they currently are in great demand in both the HPC and general IT community.

## Phase Two Results: Upfront Codesign Efforts are Essential

Although respondents demonstrated a high level of interest in exploring the design of a NASA bespoke HPC within certain bounds, they indicated that determining exactly what such a system would look like could be challenging. Properly configured codesign efforts were universally seen as the best solution to this oft mentioned problem. Specific concerns about collecting NASA requirements were mentioned by many:

- NASA has a wide and varied range of workloads that would need to be accurately characterized and combined into a comprehensive set of hardware and software requirements.
- Some respondents expressed concern that there would be 'too many cooks in the kitchen' to develop a well-defined set of vendor requirements for this community.
- For any system under consideration, vendors indicated they would be chasing a moving target, and vendors would want a better understanding of what the end goals were and how best they can meet those requirements.

On performance metrics, respondents indicated that benchmarks like LINPACK and HPCG were not sufficient to guide them in an overall system design. Many indicated that benchmarks using mini application suites, testcases, or even full applications would be the best way to determine the various strengths and opportunities of any HPC design under consideration.

- Questions to be addressed included: what are your current problems, what are you trying to accomplish, what steps have you taken to profile your ecosystem for bottlenecks, etc.?

- Respondents stressed the need for better analysis of existing NASA workflows and key applications to trace the use and movement of data to design systems that maximized the value of data while minimizing its movement.
- One respondent noted that design requirements should not only consider the top three or four workloads but should also look to address those for secondary workloads in the four through ten range to ensure that the system will be able to handle the complete site workload.

All respondents strongly indicated that the first and perhaps most important step in the development of a bespoke HPC should be efforts centered on pre-RFI vendor and NASA codesign or related team activities.

Specific suggestions included:

- Host a design conference open to all vendors that fosters frank dialogues about existing and planned codes, hardware and software requirements for those codes, and any other special considerations unique to the NASA workloads.
- Any codesign activities should not target current technology or products, but instead consider what will be available three years hence.
- Efforts should be made to engage the wider weather/climate community in any such codesign activities as well as to extend outreach to those outside the sector who are also looking to acquire systems that are not 'straight out of the catalog."
- Engage vendors to allow them to assist in analyzing workflows and applications to help them better refine their system capabilities to meet those needs.
- Offer expert vendor staff to engage with NASA application developers and computational scientists to develop representation test case and appropriate benchmarks to better understand the constraints and opportunities for performance improvements.
- Work with relevant NASA experts to collect necessary information about existing and planned application workloads to better characterize the potential of different anticipated technology developments.
- Use new highly productive programming languages that have been developed outside of the traditional HPC communities and new distributed computing methods that are revolutionizing many tasks. However, reworking workflows to use these new hardware and software technologies will require substantial effort in both programming and validation.

## RECOMMENDATIONS FOR THE NEXT STEPS

Based on the information and insights gathered in both phase one and phase two of this study, there are a number of recommendations for next steps for a NASA initiative to plan and procure a series of HPCs that effectively meet NASA computing requirements in a cost effective manner. Such efforts can be helped considerably by enlisting the insights of one or more major HPC vendors, all of which showed an interest in working with NASA on such a project.

Acknowledging that NASA has a wide and varied range of workloads that would need to be accurately characterized and combined into a comprehensive set of hardware and software requirements, NASA planners need to more accurately assess the range of existing and planned workloads to better provide specific hardware and software requirements. Key considerations involve an analysis of processor core counts and memory bandwidth, the impact of special memory frameworks such as HBM and persistent memory, and the use of GPUs and other accelerators.

The results of this comprehensive requirement analysis can then be used to compose benchmarks that consist of mini application suites, testcases, or even full applications to help determine the various strengths and opportunities of any HPC design under consideration. Such benchmarks would also be highly beneficial in providing potential vendors with insights on high priority features and performance expectations.

Benchmark efforts should consider a number of new workloads that NASA likely will be increasingly turning to in the near future that include:

- AI based programming not only in machine learning application, but to also to accelerate traditional modern and simulation jobs
- Big data jobs, both batch and near real time that could be a mix of large storage volumes along with high bandwidth data inputs from large arrays of edge connected sensors
- Hybrid on premise/cloud computing platforms to help absorb large fluctuations in workload, to test out and perhaps take advantage of new hardware opportunities quickly, and to tap into the growing base of CSP resident HPC software.

NASA, along with other interested weather/climate research organizations, could next organize a pre-competitive codesign conference or series of conferences that bring together interested commercial vendors to discuss options and opportunities based on the analysis described above. A key goal of these co-design conferences would be to build long term relationships with key vendors to generate a strategic multi-generation procurement plan for HPCs instead of targeting a onetime procurement.

Other key action items for such meetings could include:

- Working with vendors, many that have significant expertise in profiling and tuning applications to systems using future technologies, in helping to develop a robust methodology for making performance predictions from NASA benchmarks.
- Generating relevant projections of advances in future HPC hardware and software that will better align with NASA requirements but that are in keeping with larger commercial trends.
- Enlisting vendor insights and expertise on the best ways for NASA to begin refactoring and modernizing their current software base to best take advantage of existing and planned commercial technology and HPC products.

- Working with vendors to ensure that any new NASA software developments are in keeping with the latest trends in HPC application development and related software development tools.
- Seeking opportunities to help leverage any NASA hardware or software developments into the larger weather/climate HPC ecosystem to help vendors justify the use of targeted hardware or software.
- Likewise identifying and including other HPC verticals that share some common mission elements with NASA to generate wider interest in key developments, build economies of scale, and offer additional financial incentives for commercial HPC vendor participation.
- Engaging the growing cadre of cloud service providers who increasingly are on the forefront of new HPC hardware and software techniques, especially in the areas of modular software development and deployment.

Ultimately, respondents from both phase one and phase two of the survey were optimistic about NASA's efforts to confront their existing hardware and software limitations. They indicated that many of the suggestions offered stand a good chance on ensuring that NASA has access to the kind of HPCs that will be needed to continue their world class research in key climate and weather research and operations.

## About Hyperion Research, LLC

Hyperion Research provides data driven research, analysis and recommendations for technologies, applications, and markets in high performance computing and emerging technology areas to help organizations worldwide make effective decisions and seize growth opportunities. Research includes market sizing and forecasting, share tracking, segmentation, technology and related trend analysis, and both user and vendor analysis for multi-user technical server technology used for HPC and HPDA (high performance data analysis). We provide thought leadership and practical guidance for users, vendors and other members of the HPC community by focusing on key market and technology trends across government, industry, commerce, and academia.

## Headquarters

365 Summit Avenue
St. Paul, MN 55102
USA
612.812.5798
www.HyperionResearch.com and www.hpcuserforum.com

---