Recommendation to the University of Washington Regarding Workforce Support of its High-Performance Computing and Data Ecosystem

Hyak Governance Board, June 2016.
[http://www.int.washington.edu/users/mjs5/HYAK/governance/index.html]

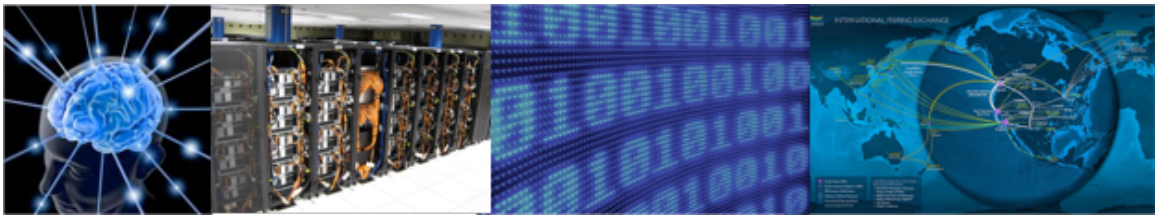

Preamble

High-Performance Computing and Data Ecosystems (HPCDEs) are essential to scientific discovery and to advancing many fields of science, engineering, and medicine. They are key to furthering the social sciences, are critical to an increasing number of industrial sectors and other commercial activities, as well as many other areas that have significant impact upon society. A HPCDE provides the necessary infrastructure to: (*a*) share expertise and experience in high-performance computing and data, scientific computing, and its application across a broad range of disciplines that are critical to the state of Washington and the nation; (*b*) educate the next generation of leaders who will continue to address the most challenging scientific, engineering and societal problems; and (*c*) enhance UW's leadership in research and education by providing faculty, students, researchers and staff with access to, and the capability to use, the HPCDE resources that they require.

The United States is poised to deploy a generation of pre-Exascale leadership-class computers, which are on the critical path to Exascale Ecosystems. These are expected to be heterogeneous-architectured supercomputers embedded in low-latency/high data rate communication fabrics with peak performances in the multi-exaflops range that are capable of producing, storing and curating exabytes of data with a range of storage options and networks sufficient to generate, analyze and store the data, and with real-time capabilities to interface with users. Such systems are essential to meeting the objectives of a wide range of the domain sciences,

engineering applications, medicine and information sciences. The optimal use of these compute resources will require close collaboration between domain scientists, applied mathematicians and computer scientists. Given the complexity and cost of such resources, novice users will learn the necessary tools from running single-core serial applications on ``workstations'', to multi-core parallel applications on small CPU clusters, to larger-scale parallel applications on architecture similar to the leadership computing platforms, before ultimately migrating into the exascale ecosystem (the ``HPC user escalator''). The Hyak phase-1 ecosystem has proven to be a successful resource in this education path (P2E, or "Path to Exascale), enabling students with a limited knowledge of scientific computing to quickly develop the skills necessary to undertake meaningful, productive use of some of the nation's largest supercomputers.

Equally important to the education path of young scholars, is the evolution of codes to optimize scientific productivity. During the last decade, homogeneous compute platforms, comprised of CPUs evolved toward heterogeneous systems with their benchmarked performances dominated by accelerators. Existing application codes did not run or ran poorly on such systems, and a significant effort has been required to translate or reinvent code in a way to exploit the heterogeneity. This trend will continue and increase going forward. At the UW, we have seen a significant growth in the demands on the HPC trained workforce in UW-IT due to this evolution, and this growth is expected to accelerate during time frame of Hyak NextGen and beyond.

Data – storage, movement, production, and analysis, are significant foci in current national supercomputing efforts, and have a defining role in the emerging Exascale ecosystems, requiring new algorithms, code structures, concepts regarding workflows and most other aspects of computing. As a result, the need for increased institutional support to provide the workforce infrastructure that enables UW researchers to remain at the forefront of scientific computing and information technology is now extreme. Further, the workforce need is similarly extreme to support the UW administration, non-scientists, business intelligence. For example, there is a growing need for real-time rapid access to, and manipulation of, large-scale databases across nearly every aspect of contemporary university operations.

We are not alone in identifying this present and future need. Michigan State University has recently created the *Department of Computational Mathematics, Science and Engineering* (CMSE) to lead research and educate the next generation of scholars in data science, computational modeling and high performance computing. CMSE has been created to enhance MSU research across a number of frontiers, and provide a talent stream with applied computation expertise to industry. This forward-looking strategic move by MSU will position them well for the future of higher education and cutting-edge research.

The UW's High Performance Computing and Data Ecosystem

The Hyak High Performance Computing and Data Ecosystem has been in full production since July 1, 2010 — about six years. In that time the Hyak user community has grown by more than ten-fold (10x), increasing in diversity at an even greater rate. User outreach and support demands continue to grow with the expanding diversity among user needs. This increase in demand has resulted in a growth in system size and complexity, requiring increased systems engineering and support efforts. In short, Hyak has matured. It is no longer a startup system, and, as a going concern, staffing levels should increase to effectively support this growing resource.

Hyak has grown from fewer than four hundred compute nodes, with fewer than 3,000 CPU cores, to approximately 900 compute nodes with approximately 11,000 cores today, while system performance has increased by more than 20x. With the addition of NextGen Hyak in late 2016, total system size will grow to more than 14,000 CPU cores. At that time, the UW will have entered an era in which Hyak HPCDE staff will operate a minimum of two different compute clusters, each with a different design (as the ecosystem evolves into the natural cycle of ``standing up'' the next generation of system as the previous one runs out its last few years). Sustaining this evolution requires skilled, and difficult to find (and hire), HPCDE engineers.

In the last six years, Hyak's user community has grown from approximately 20 unique users per month at the outset to hundreds of unique users per month today. And the diversity of users has increased immensely, from just a few natural science domains initially, to the situation today where Hyak supports research in dozens of academic units across the university. More so than the absolute increase in numbers, it is this diversity which presents the greatest outreach and support challenge. The Hyak governance board (HGB) and UW-IT have embraced the "HPC User Escalator" concept in principle, but to be effective, substantial increases in advanced user support staffing are necessary.

The 2013 CASC report on staffing levels at academic HPC centers** is helpful in understanding the situation we have in the context of national HPCDEs. The Hyak HPCDE is an outlier in terms of being understaffed. That report breaks down results by the size of each center, measured in the number of CPU cores. Phase-1 Hyak sits at the boundary between small and medium sized centers. The report's basic measure of relative staffing is **Compute Cores/FTE**, and by that measure, Hyak has **FIVE** times less staff (i.e. 1/5 of the staffing level) than the median of other small centers and at **LESS THAN HALF** the median for medium-sized centers. In terms of advanced user support — the critical limiter in realizing the goals of the HPC User Escalator — the situation is even more challenging. The median staffing levels in this area for small and medium-sized centers, respectively, is 2.0 FTEs and 4.0 FTEs, while at UW, staffing is at the level of 0.5 FTE.

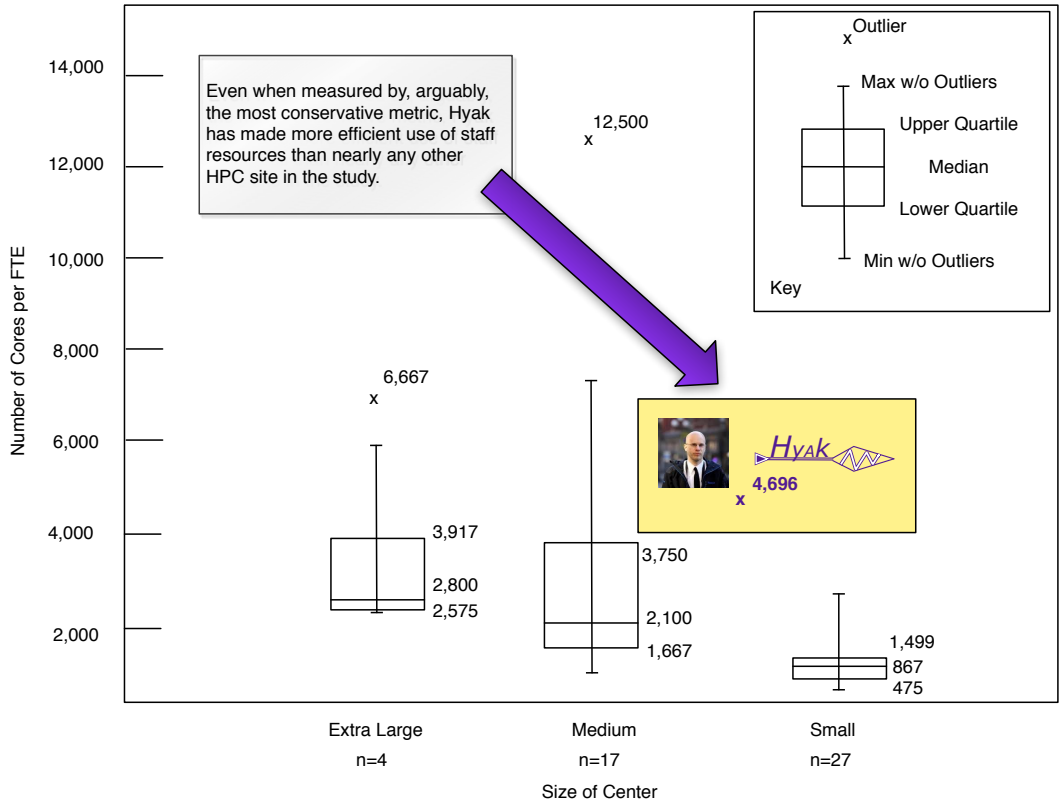

Figure 1: Conventional compute cores supported by each FTE at HPCDEs within the US.

Even in the absence of evidence from other HPC centers, it can be concluded that the present situation of operating the Hyak HPCDE with a single systems administrator is sub-optimal. It exposes the Hyak enterprise to an unacceptable level of risk of service interruption when the administrator is unavailable (vacation, sick leave, etc.), and increases the work stress of that person (increasing the risk of departure of that person). The Hyak Governance Board believes that a minimum of two senior systems managers is required for the operation of such an enterprise.

Recommendation

To realize the full potential of UW researchers and to effectively educate the next generation of scientific leaders in the state of Washington, we require new investments in the Hyak High-Performance Computing and Data ecosystem. These investments are crucial to the UW recruiting and retaining world-class faculty, students and researchers, and are necessary to satisfy the clearly demonstrated increasing demand for the HPCDE established with Phase-1 Hyak.

In light of the increasing number and diversity of users of the Hyak HPCDE and the established ``long tail'' distribution of usage requirements, along with the increase in the size and complexity of the system anticipated with Hyak NextGen, we

recommend that the University of Washington enhance its HPC-proficient workforce supporting the Hyak HPCDE and associated activities by at least two FTEs.

**http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2313089