

Cyber-Infrastructure Support

Objectives

Develop additional services to address un-met demand for expertise and infrastructure in support of UW research computing.

- Consulting - Provide subsidized, fee-for-service assistance in the tools and techniques of scalable computing, including domain-specific application selection, workflow optimization, code tuning, and data management. This consulting would address the broad range of large-scale computational challenges confronted by UW scholars today, enabling not only established High Performance Computing (HPC) users, but also enlarging the cohort of UW researchers prepared to apply computational resources to their research. This consulting would complement, and be offered in partnership with, services provided by the eScience Institute.
- Storage – Provide lower cost, enterprise-class storage for research data by changing the lolo storage cost model to be on-par with “departmental” storage options. The lolo storage system includes separate services for data archiving and data sharing. Both are designed to scale in capacity and performance as an enterprise resource. Both are also situated in a campus Science DMZ which makes the data available for access by large-scale CPU resources (e.g. Hyak) via high performance networks (on and off campus). By reconfiguring the lolo service to use a condo pricing model, researchers would pay lower storage costs overall, while protecting against the data security risks and loss associated with lower cost, non-enterprise class storage options.
- Collaboration - Offer partnership opportunities and easy access to services and communities to improve research computing. By joining consortia such as CASC (Coalition for Academic Scientific Computation), Open Science Grid, and Xsede, the UW will enhance its ability to anticipate and adapt solutions for new use cases and new technologies. UW researchers will also have easier access to new HPC options and services through: 1) the addition of a web site specifically for Scalable Scientific Computing and 2) an update to the UW-IT service catalog to provide a portfolio view of services available to researchers.

Strategic Context

UW is a top tier research university, yet we provide minimal central IT support for researchers despite the increasingly IT-intensive nature of research. During a September 2012 meeting of research-focused universities at Clemson, the ten participating universities (including Harvard, ASU, USC, Wisconsin, and Stanford) compared FTEs devoted to high-performance computing support for research. The UW is on the low end when compared with these peers.

Going back to at least 2005, researchers on campus have emphasized the need for access to CI-resources, including scalable CPU, storage and networking, as well as access to the expertise necessary to exploit these resources. This was a key point in the 2009 in-depth “Conversations with the University of Washington’s Researcher Leaders” report. The Hyak Governance Board recently reviewed the motivation and goals articulated in the formation of the eScience Institute and identified the same list. UW has an opportunity to build on the foundation laid by the first three years of success with Hyak and lolo by making the investments detailed here.

Key Benefits

- Improve the effectiveness and efficiency of UW researchers as their work becomes more IT-intensive.
- Increase competitiveness of UW researcher proposals and faculty recruitment.
- Enable the “long tail” of UW faculty and researchers to also make use of HPC resources.
- Establish additional centrally available skills to provide quality assistance to researchers in choosing among complex and rapidly changing technology options, and promote collaboration and sharing of services and knowledge throughout the organization and beyond.
- Replicate the existing Hyak HPC computing model for storage in order to promote shared and standardized services.
- Enable research data to be easily shared and processed over our highest speed networks.

Risks

Not funding this proposal, could result in the following:

- UW spends more money overall than it needs to as departments duplicate costly storage options.
- Researchers use less secure storage that put their data at risk of loss or compromise.
- Research grant proposals are less competitive due to lack of adequate cyber infrastructure and support, thereby putting our status as a top research university at risk.

Budget Request

UW-IT is requesting additional funding of \$280,000 per year to develop a new service area for support of Cyber-Infrastructure:

- Consulting - Provide consulting resources to help with high-performance computing selection, workflow optimization, code management, and system operation. The 1.0 FTE to be hired would ideally be hired as two ½ time positions - one senior and one junior – to maximize flexibility while controlling costs. This service will be offered for-fee, but at subsidized rates (i.e. not full cost recovery) to be competitive to graduate student labor costs.
- Storage – Subsidize lolo data services to provide a storage cost model on-par with “departmental” storage options. Reconfigure this enterprise-class service to use the same condo model as is used with Hyak and cover its fixed infrastructure costs.
- Collaboration – Pay dues to join Coalition for Academic Scientific Computation (CASC). An improved web presence and participation in other consortia (e.g. Open Science Grid, and Xsede) is already funded by UW-IT.

Budget Request:

Description	Total
Consulting - 1 FTE (½ senior, ½ junior) Sal/Ben	\$170,000
Storage – Subsidy of lolo data services by covering infrastructure costs of the service	100,000
Collaboration – Fees to join CASC	10,000
Total	\$280,000