

## Current and Future Research Computing Infrastructure

To meet the demands of the research portfolio, the WSU research computing community turns to a number of advanced computational facilities including (1) the small-to-mid-scale HPC facilities of Kamiak at CIRC, (2) smaller departmental and satellite campus clusters, such as the Aeolus experimental heterogeneous cluster in the Voiland College of Engineering and Architecture, (3) regional facilities such as the Pacific Northwest Regional Laboratory, and (4) national supercomputing facilities including leadership class facilities at Oak Ridge, Lawrence Livermore, and Argonne National Labs, as well as NERSC and XSEDE.




Figure 1. Growth of Kamiak, from its inception in the fall of 2016.

exchange for prioritized service on the nodes they own. In addition to investor nodes, the cluster also includes a small set of community-owned nodes that are equally available to all researchers. Under this model, the HPC facilities are open to the entire WSU research community, and any user can run jobs on idle resources that are not otherwise in use, called the "backfill queue". The efficacy of this model is manifested by the strong growth of the user base, the faculty investments, and the corresponding cluster resource utilization (Figure 1).

For applications that must run at scales beyond those feasibly supported in Kamiak, researchers use Kamiak as a testbed for tuning and refining their applications so that they can be run on national supercomputing and leadership class facilities. A key component for enabling the transition to larger facilities is the staff support CIRC provides for application development and portability, as well as assisting researchers in identifying and applying for grants to use those facilities. CIRC also acts as agent to provide appropriate support for the disparate departmental and satellite cluster facilities in order to coordinate activities under an umbrella of integrated and responsive support. This includes enabling network and storage connectivity between Kamiak and its outliers so as to facilitate a seamless high-performance computing environment and a framework for transitioning to Kamiak as appropriate.

To meet the ever-increasing computational needs of the faculty and researchers at WSU, CIRC in collaboration with faculty governance is pursuing the provision of additional HPC resources on two fronts. The first is upgrading Kamiak, which follows the highly successful condominium model, to radically improve its performance and expandability. The second is the proposed acquisition through an NSF MRI of a regional state-of-the-art HPC cluster, called Camas, that will provide vitally needed open-access computational resources for researchers across WSU and its partners at regional institutions.

## I. Upgrading Kamiak to Meet Future Needs

The first step to meet the increasing computational needs of the faculty and researchers at WSU is to design and implement an infrastructure refresh of the current Kamiak HPC cluster. Such a refresh will encompass upgrading the storage, network, and login and administrative nodes, and is anticipated to result in a 20-fold increase in performance, principally due to a 50-fold increase in storage throughput. The design for the upgrades is complete, the funding is in place through Omnibus and OR supplements, and the implementation is expected to be complete in the Spring of 2021. The upgrades will yield a system that is structurally similar to the regional cluster proposed in the MRI. The upgrades will consist of:

- An **all-flash parallel file system** that provides 1PB of usable storage that is scalable and interoperable with cloud storage through an S3 cloud API.
- An **HDR-100 Infiniband network core** with spine switches that operate at 200Gbps and splitter cables for 100Gbps connections to nodes, a 2x speedup.
- **State-of-the-art login and administrative nodes** that are fully virtualized for high availability and optimal resource utilization (a so-called hyper-converged infrastructure), and which will support the spin-up of portals for scientific gateways.

## II. Development and Acquisition of A Regional Infrastructure Supporting the Land-Grant Mission of Resource Management

An NSF MRI (Major Research Instrumentation) grant has been submitted to secure funding for a next-generation HPC cluster, called Camas, that will provide open access to computational resources for researchers across WSU and its partners at regional institutions. The proposed computing system will include the latest generation of Intel Xeon Scalable CPU and Nvidia GPU's, a state-of-the-art parallel file system with an all-flash tier that transparently stages actively used data, and a fast interconnect between compute nodes as well as storage. The cluster will have 105 compute nodes, 10 of which are dual GPU's and 1 of which is a large-memory node, that will provide 4,216 CPU cores and 102,400 Cuda GPU cores. The proposed system follows a modular design that can support expansion up to 160 nodes, and scale-out storage. Camas will also include front-facing portals that support public databases and scientific gateways that will further expand external engagement. The architecture will provide key features to enable data-intensive and AI-focused scientific computing in areas such as evolutionary biology, horticulture, and natural resource management. These features include:

- A **parallel file system** that is scalable and interoperable with cloud storage. By transparently migrating data to the all-flash tier from the adjacent object store through an S3 cloud API, the model provides high performance in a user-friendly fashion that does not require explicitly staging to a burst-buffer.
- **State-of-the-art GPU's** that are 6x faster than previous generations, for AI and machine learning applications.
- **Dense compute nodes** at or above the 40-core range, which accelerate multi-threaded applications and decrease node-to-node communication.
- An **HDR100 InfiniBand network** to enable scalable computing and vastly improve file access. The InfiniBand network will handle both inter-node communication and file access, at speeds 8x faster than that currently deployed.
- **Public-facing portals** for science gateways that provide community access to resources such as genomics databases and applications.

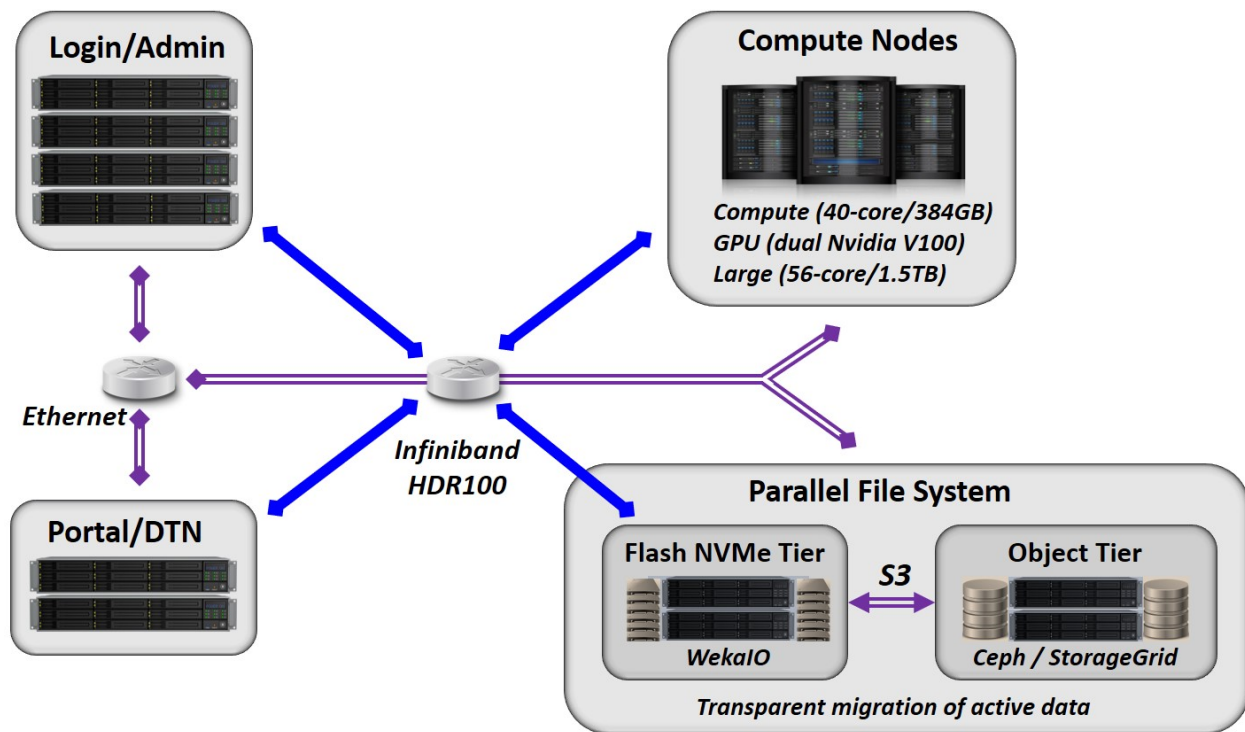


Figure 2. The next generation HPC cluster, Camas.

Figure 2 illustrates the design of the Camas next-generation HPC cluster.

### III. Responsive Support for Diverse Research Computing Environments

It is envisioned that the purview of research computing at WSU will not reside with a single HPC cluster such as Camas, but rather with a set of multiple HPC computing environments spanning WSU's several campuses, as well as in cloud, regional, and national HPC facilities. As alluded to previously, there are a number of reasons why a single cluster, or even several HPC clusters of similar size as can be found in many universities, will not suffice for the research computing needs of WSU.

First, for extreme scale computations, local facilities simply will not be powerful enough to provide responsive results. In such a scenario the local small-to-mid scale facilities provide a stepping-stone for transitioning applications to leadership class facilities. Second, experimental clusters, such as Aeolus in the college of engineering, provide ad-hoc and heterogeneous capabilities (needed for example for chip simulation) that cannot coexist with a conventional HPC cluster. Third, security concerns for processing sensitive information (e.g. Controlled Unclassified Information) require isolated and secure enclaves, in the form of either walled-off clusters or secure cloud HPC services (e.g., AWS GovCloud). Cloud HPC computing, while about 10 times more expensive, is also attractive for its flexibility and ease of configuration in standing up short-lived or specialized projects, as well as providing world-wide access to shared data. It is notable that WSU has an egress-waiver with AWS that allows cost-free movement of data into and out of its services.

To fully support research computing at WSU, and advance the land-grant mission, strategies must be followed that accommodate the reality of multi-cluster environments. These strategies include: 1) the provision of a central campus-wide support framework for computational scientists and system administration, 2) satellite clusters that are maximally configured to seamlessly interact with each other and with Kamiak and Camas, both in job submission, and in storage and

network connectivity, and 3) establishing and maintaining robust internet connectivity to outlier national facilities and cloud services, as well as other campuses.

#### **IV. Strategies for Enhancement and Growth**

There are several strategies that can be followed to increase user participation, grow HPC facilities, and enhance the user experience. The high quality of CIRC support staff, and the excellent service they provide, and CIRC-wide promotion and outreach activities, are systematically attracting new users and building the investor base. Second, CIRC can begin to leverage recently implemented WSU infrastructure, like the 10Gb HSSRC campus-wide research network in Pullman that was recently connected directly to the Internet2 national research network. CIRC can also strategically work with ITS to grow infrastructure in a sustainable manner that benefits computing research. At the same time, CIRC should be an advocate to support the development of dedicated fault-tolerant facilities for housing computing resources, as is customary at *peer land-grant institutions* like Clemson and Purdue that have developed research computing facilities that rank in the 2018 top 500 list of worldwide supercomputers.

When necessary, CIRC supports and will grow activities aimed at identifying and migrating computing research to the most appropriate facilities for the application. In some cases, this means helping faculty write user facility grants for national supercomputing or leadership class computing resources that are supported by NSF and DOE. In other cases, cloud computing may be most appropriate. With the rise of cloud computing and pay-as-you-go providers such as Amazon Web Services (AWS), Microsoft Azure, and Google Cloud, there is increasing pressure to offload the cost of setup and maintenance of local HPC facilities. However, it is the role of CIRC to support the best solution for researcher needs, and thus we actively identify the specific cases where cloud resources provide the best cost-to-benefit ratio. The trial-and-error nature of hypothesis-driven research creates difficulties in estimating costs as the cloud compute resources needed to develop new algorithms, repeating failed analytical attempts, and optimizing algorithms can deplete budgets prior to full completion of an experiment, especially as data size and computational time can scale exponentially. CIRC has identified a select few cases where cloud computing provides distinct advantages – these include the management of HIPAA compliant data.

#### **V. Beyond Computing Infrastructure**

There is no doubt that the infrastructure associated with doing research computing goes beyond actual high-performance computing nodes and their support. Indeed, a research computing ecosystem encompasses a larger class of support infrastructure that can and should be integrated with the capabilities and vision of other support units. Many other units on campus deal with ancillary support infrastructure that is in the best interest of the needs of their faculty. Toward this end, we have identified key domains where CIRC should have active collaborations and engaged strategic planning.

**A. Backup and archival storage.** Though CIRC's mission includes many technical activities, the overarching goal of CIRC is to support research. With limited resources, CIRC must necessarily focus on activities for which only CIRC is positioned to engage in while leveraging existing technology services throughout the university system where they already exist. Within the WSU system, multiple areas already invest in storage platforms, and at least one large storage as a service platform already exists. CAHNRS offers a low-cost network attached storage service for backups, archives, and network attached shared folders which includes dedicated research nodes that are attached to the HSSRC. In the future VCEA may also offer storage as a service. In assessing the suitability of the CAHNRS storage service for use by CIRC and CIRC customers, we examined the services sustainability model, roadmap, and the scope of services and their integrations with CIRC.

**1. Sustainability Model.** The CAHNRS storage platform recovers costs on a per GB/Year basis, which pays for the physical infrastructure and direct costs of the service. Indirect costs, notably support overhead, are distributed across the WSU system through the use of delegated management. Examples include delegation of management functions to the College of Business, the College of Arts and Sciences, and the College of Education. The distribution of management allows faculty to receive assistance and support from their local technical support while preventing the concentration of burdensome workload and personnel costs in any one unit. The consolidation onto a single storage platform creates cumulative cost efficiencies, which allows savings to be passed on to faculty customers or to improve the service incrementally without raising rates.

**2. CAHNRS Storage Roadmap.** The current vision for the CAHNRS storage service follows two concurrent tracks. The first track is the gradual evolution from the current CAHNRS centric service center into a coordinated, multi-college service center that allows for shared investment and management of common infrastructure. Some headway is being made through the increased adoption of the CAHNRS storage service across the WSU system. However, additional work remains to formalize and maintain mutually beneficial partnerships across the colleges.

The second track of the CAHNRS storage roadmap is the introduction of local storage nodes statewide, including and especially at the non-Pullman campuses. CAHNRS is currently working with Spokane to deploy new storage nodes at that campus, and preliminary conversations are underway to do the same at Everett. The staging of storage nodes across the campuses will allow faculty to access storage that is local to them, but which is also a part of a larger network of services that tie into research resources in Pullman. One example of this is the storage node in Spokane which will allow Spokane faculty to save data locally, and have that data automatically stage itself to Pullman for use on Kamiak. Likewise, when Spokane faculty save the output of their work from Kamiak to the Pullman storage nodes, the data will automatically sync to their local storage in Spokane.

**3. Scope and Integration of Services.** Currently, the CAHNRS storage service uses a mix of high availability nodes, and offsite backup nodes to ensure that every file always has an offsite backup and at least two redundant systems on which service can be delivered. Currently, data has been backed up to the CAHNRS collocated servers in the Westin Building in Seattle. However, backup nodes are being installed in Spokane, and possibly Everett in the future. CAHNRS is also evaluating AWS Deep Archive as a future backup and distribution platform for the Highly Available nodes around the state. The CAHNRS storage service is integrated with the WSU HSSRC and state-wide high-speed networks. These networks and distributed node model allow for flexible integrations with CIRC on a per-node basis.

**B. Development of visualization capabilities.** Creating visualizations is an important step in understanding data, and as computing and research evolve, so do the visualization needs of the faculty. To address emerging needs, CIRC will take an iterative approach to build and refine support for visualization technologies that reflects the needs of faculty over time. To begin these efforts, CIRC will ensure that tools requested by faculty to create visualizations as an output of their work on Kamiak are available and that the hardware offerings of dedicated nodes reflect the potential needs for graphically intensive computations. The scope of both the software and hardware offerings will be based on faculty needs communicated to CIRC or captured by surveys.

As visualization needs grow, a second phase in the development of visualization capabilities may emerge as new research lines. Whether researching the use of visualizations in pedagogical contexts or as a system for exploratory research, faculty who establish a research program in visualizations will have unique challenges and needs for computational and storage support. As these programs grow, CIRC will adjust the hardware and storage offerings to support these activities. Additionally, new modalities of human-machine interaction may be necessary to support these research lines. While it is impossible to know the future needs of such programs,

CIRC will work to ensure that the central computational services remain relevant to these programs. Critical to the development of visualization capabilities is the reliance on faculty feedback to guide CIRC activities and investment. The wealth of possible technological avenues to pursue concerning visualizations and human-machine modalities requires faculty direction, lest CIRC invests resources in ways that don't create value for the faculty.

**C. Connecting to regional and national infrastructure.** In order to enhance the competitiveness of WSU research computing and support for data-intensive research, it is critical that WSU engage with regional and national-level Cyberinfrastructure (CI) projects. There are several ways that CIRC can help. First is the integration with existing national computational infrastructure. For example, on campus facilitators can serve as XSEDE campus champions who coordinate with XSEDE to provide local outreach. In addition, job submission can be integrated with the Open Science Grid (OSG) such that on-campus users can directly submit computational jobs to OSG from local infrastructure, thus easing the burden to scale workflows from the local to the national level.

The second is integration with and participation in other national network infrastructure projects. With regards to networking and data movement, the Pacific Research Platform (<http://pacificresearchplatform.org/>) aims to resolve communication challenges to create a seamless communication platform between research institutions across the Pacific Wave (<https://pacificwave.net/>) and the Pacific Northwest GigaPop (<https://pnwgp.net/>) networks. The PNW GigaPop is also WSU's entryway into the Internet2, a key national research network (<https://www.internet2.edu/>). WSU's uplinks into the PNW GigaPop in turn are through Washington's K-20 and Idaho's Iron networks. Fully leveraging such networks allows WSU researchers to more easily move large data with collaborators at other institutions. It also allows agricultural researchers at WSU who maintain farms in both hemispheres to more easily move data.

While WSU is a member of the Northern network Tier Consortium (<https://www.ntnc.org/>) which fosters relationships with national and international network organizations including the Internet2 and the PWN GigaPop, further direct participation with PRP and Internet2 organizations would prove beneficial. Member institutions join the PRP collaboration and participate via active communication and attendance at yearly meetings. Currently WSU does have some engagement with PRP but this is ad-hoc through faculty outreach and some interaction from CAHNRS IT, and not via coordinated institutional engagement. Internet2 is also actively engaged in development of Cyberinfrastructure as it relates to data movement and analysis across the network. Better engagement with Internet2 and faculty with big-data use-cases could foster more WSU involvement in national-level CI projects, even if just providing scientific use cases.

As another example, the NSF recently established several regional "Big Data" Hubs meant to foster development of CI and to address grand challenges important to each region of the US. One of the host institutions for the western region is at the University of Washington. WSU has no formal participation (<https://westbigdatahub.org/>). Improved institutional engagement with the Big Data Hub would provide WSU researchers an avenue for engagement in regional-level collaborations to address challenges in the state and our region and would foster improved state-wide relationships between researchers at UW and WSU. Moreover, NSF and NIH both offer requests for proposal for development of CI such as from NSF's office of Advanced Cyberinfrastructure or recently NIH's Data Commons project. A higher-level of involvement in national projects will position WSU to be more competitive, creative and novel when submitting proposals to these programs.

The University of Washington is more connected to other national Cyberinfrastructure partners such as the ACI-REF (<https://aciref.org/>), recently renamed the Campus Research Computing Consortium (CaRCC) (<https://carcc.org/>), the Pacific Research Platform (<http://pacificresearchplatform.org/>) and others and received funding for its NSF Center in Data

Science. Nationally, CIRC can provide the leadership and management team needed to support a long-term strategic plan that could allow WSU to compete with other land-grant peer institutions such as Clemson and Purdue. For example, CIRC can serve as WSU's point of contact to national groups such as the CaRCC which is currently developing resources such as a cyberinfrastructure maturity checklist to afford greater integration and cooperation in cyberinfrastructure across the nation. Through active participation with such a group, at the national-level, CIRC can position WSU to participate in development of national-level cyberinfrastructure for data science. Lastly, CIRC can directly impact the land-grant mission of the University by providing Cyberinfrastructure needed for big data research in precision agriculture and engineering that can directly impact growers, our state's industry, and individuals.