

A. INTRODUCTION & EXECUTIVE SUMMARY

Vision: Cyberinfrastructure as the key to research acceleration

The University of Nevada, Reno (UNR), the land-grant institution in the State of Nevada, attracts over \$100M annually in sponsored research projects and was recently ranked as a Carnegie Highest Activity (R1) Research University where significant student participation in research drives next-generation workforce development. Essential to this status is coordinated implementation of a cutting-edge cyberinfrastructure that enables institutional success and accelerates the research process.

“Cyberinfrastructure” is defined as a fabric of highly connected systems for information and data acquisition, visualization, computing, and storage with associated human expertise serving end-to-end scientific and engineering workflows to improve scholarly productivity and enable breakthroughs not otherwise possible.

Advances in technology have shifted the global research paradigm from reliance on individual experience to the need to generate and access large datasets, perform advanced analytics, and produce unique outputs across interdisciplinary teams within and outside institutions. Success in research going forward depends entirely on the efficiencies of data movement, processing, and management – which all require coordinated cyberinfrastructure, expert technology facilitation personnel, and sustainable institutional investment.

Faculty members, staff, and students at UNR perform internationally recognized research in diverse fields such as materials engineering, wildfire, biotechnology, neuroscience, intelligent systems, climate, and water resources. In order for these and other research efforts at the University to maintain the highest standards of scholarship, increase in scale and efficiency, and lead in scientific discovery, the expertise, tools, and capacity of a robust and well-designed cyberinfrastructure (CI) remain essential.

Cyberinfrastructure governance and development at UNR

UNR has invested in the building blocks for an effective culture of collaborative CI development across the Offices of Information Technologies (OIT), Research & Innovation (R&I), and the campus academic community. The University established a campus Cyberinfrastructure Committee (CiC) that is led by tenure-track faculty members, reporting jointly to the Chief Information Officer (CIO) and the Vice-President for Research and Innovation (VPRI). OIT received central institutional funding to hire a research-active, doctoral-level Director of Cyberinfrastructure starting in 2018, and institutional High Performance Computing (HPC) was moved underneath this position in 2020. This evolutionary direction ensures a research-facing perspective behind CI strategy and administration, vital during a time of campus transition from a focus on administration and teaching to R1 sustainability. The CiC and CI Director work together to ensure that CI resources at UNR meet clear institutional needs, develop in alignment with state and national research priorities, and are grown into a sustainable CI Program with a broad range of accessibility.

This planning document is an evolving strategy tool for UNR OIT, R&I, CiC, and the institution at large to communicate vision, align objectives, guide effort, apply resources, and demonstrate direction. While the primary responsibility to draft and develop the institutional CI Plan falls to the CI Director, the faculty-led CiC provides additional input, discussion, and ultimately approval and consensus after soliciting feedback from the campus at large. Major updates to this document occur every two years, associated with internal and external survey tools, or as necessary.

Technical implementation of CI at UNR is led from within OIT, as the institutional technology architectures of networking, datacenters, and computing are highly centralized. However, the University CI Program is envisioned as a shared responsibility and contiguous group of human and machine resources from the college division level up to central enterprise administration, as faculty member-facing support roles need to remain coupled with research domain applications.

CI Plan Contents

Sections A-B: Introduction, Vision, & Strategic Objectives

Sections C-D: Current Infrastructure Overview & Projects (*subject to semi-annual status updates*)

Appendices A-B: Detailed CI Program Goals & Elements

Appendix C: CI Committee By-Laws

Appendix D: Program Development Strategic Timeline

B. CYBERINFRASTRUCTURE STRATEGIC OBJECTIVES

UNR has identified a number of long-term strategic CI objectives that are essential to achieving greater research capability, regional technology leadership, and institutional risk reduction for research data and processes. Existing, fragmented CI at the institution today can only provide a fraction of needed capability, and so these objectives serve as a basis for a comprehensive UNR CI Program that will serve the entire UNR research community and positively impact the vision of accelerated research and discovery.

Community of expertise: A core cadre of research technology experts that develop and maintain critical research infrastructure, mentor student assistants, facilitate research workflows, contribute to regional and national cyberinfrastructure communities, and introduce new technologies into the UNR ecosystem.

Connectivity: Diverse and flexible research network connections with highest-possible speed and capacity between on/off campus research endpoints and datacenters, other Nevada System of Higher Education (NSHE) institutions, regional research and engineering partners, and national research fabrics.

Research data storage: Reliable, scalable storage systems accessible to all researchers to serve broad categories of general-purpose, high-performance scratch space, and backup/archival. Should include cloud-native storage presentation and provision for protection from both disaster and unauthorized access.

Consolidated data centers: Modern colocation facilities for both centrally-managed and researcher-maintained hardware systems that provide reliable power supply, adequate cooling capacity, and secure physical space for a range of server technology needs and support levels.

Computing: Capability and expertise across a range of remotely-accessible computing services vital to most research workflows, including: high-performance compute (HPC), interactive/visualization computing, persistent service nodes, throughput computing, and lab-in-a-box servers.

Cybersecurity: Facilitation and implementation of methods and practices that ensure compliance and privacy for a range of sensitive research workflows, that do not adversely affect the ability to collaborate, acquire/share data and information, or otherwise hamper other institutional CI systems performance.

Training, education, and transformation: Empowering campus IT professionals, faculty researchers, and students to make best use of available CI resources on and off campus, while preparing the institution for disruptive and emerging technologies and practices.

Scalable management: Integration of semi/automated systems and management tools designed to scale and smooth research workflows, including access, resource provisioning, and data exchange.

National community alignment: Presence and activity in regional and national research CI communities that link practices and distributed resources across political/geographical boundaries.

FAIR Data Management: Assisting institutional development of policy, practices, and technical capability that enables Findable, Accessible, Interoperable, and Reusable (FAIR) principles in regards to data and metadata capture, organization, documentation, and curation.

Access: Provide a range of common services and technology at “no charge” to all institution researchers, with higher levels of fee-based service and resources available on-demand to funded projects.

Sustainability: Build and sustain a comprehensive CI Program with a robust and dedicated budget responsive to institutional trends and insulated from market perturbation – using a diversity of state allocations, facilities and administration overhead, division-level investment, fee-for-use, CI-focused grants, and external donations and endowments.

C. CURRENT INFRASTRUCTURE OVERVIEW

Central research computing

UNR OIT manages multiple **datacenter facilities**, with consolidation of research computing resources into two colocation spaces: the on-campus Research Computing Data Center (15 kW x 16 racks), and off-premise with industry partner the Switch Tahoe-Reno Data Center (10GbE uplink; 36 kW x 10 racks) .

UNR OIT maintains a local **batch HPC system** (“Pronghorn”) and attached storage. Located in the Switch Tahoe-Reno datacenter, Pronghorn is available (2018) to all UNR and other NSHE-affiliated researchers primarily as a fee-for-service model. Pronghorn is comprised of 108 nodes with 32 processors and 256GB memory (11 nodes containing 4 Nvidia Tesla P100 NVLink GPUs each) connected with a 100 Gb/s Intel Omni-Path fabric. Pronghorn’s 2.2 PB storage contains two DDN Spectrum Scale GPFS appliances.

Virtual lab workstations and interactive computing are centrally hosted have transitioned from pilot (2015) to production (2020). UNR OIT offers a dedicated administration service for faculty-purchased hardware (“SilverWolf”) to provide robust, on-demand remote Windows multi-user workstations that are accessible from any platform anywhere that can run a web browser. Private cloud multi-OS virtual machines for temporary use are also available on a limited basis through the enterprise WolfCloud cluster.

UNR OIT provides limited **data storage** access to an enterprise scale-out NAS solution with over 160 TB of tiered file storage on a redundant platform that supports de-duplication. Redundant SAN’s with a dedicated fiber channel storage network provide 78 TB of block storage and is co-located in campus data centers with backup generators, UPS, and environmental monitoring.

For **cloud-hosted data storage**, the UNR NevadaBox and OneDrive services offer limited file storage and sharing hosted by third-party cloud providers. This provides cloud storage to students and faculty members with a secure sign-in, ability to store sensitive data, and the ability to share and collaborate with outside entities.

Networking

UNR is connected to **Internet2** via the Nevada System of Higher Education System Computing Services (NSHE-SCS), which maintains a 100 Gbps connection to the Internet2 backbone national network. The campus connected in 2019 to the SCS WAN via a 100 Gbps link with 20 Gbps redundancy. IPv6 routing is available from the SCS WAN, but is not generally employed on campus as of 2020.

The campus network core is a diverse series of five routers (Brocade VDX) fully meshed in a VCS fabric with 40GbE links, allowing for data to move through the core at 160 Gbps. These are dispersed around campus to fiber aggregation nodes with UPS and generator back-up. The campus is segregated into 17 regions each having two 10GbE links using single mode fiber to the diverse core allowing for redundancy. Each building has connectivity via fiber with a mix of 1GbE and 10GbE links. Wireless connectivity is integrated as a core part of the total campus network, with wireless access in all research and instructional spaces. Network status and performance is monitored by **SolarWinds**, **Intermapper**, and **perfSONAR**.

A separate physical layer for **core research-facing networking** was established in 2019, providing dedicated 40 Gbps downlinks to strategic locations on campus, and 100 Gbps uplinks to both the campus core network (inside the border firewall; “Research DMZ”) and NSHE-SCS (outside the border firewall; “Science DMZ”). The two security zones are software partitioned on the same switching/routing hardware.

There are other university-related networks managed by research units, primarily the region-leading wide-area **Research and Hazards Monitoring Network**, which is a mix of landline, microwave backhaul, and local wireless networks that spans the state of Nevada and beyond. It provides transport for sensor data and remote connectivity managed by the Nevada Seismological Laboratory (NSL) and ALERTWildfire, with other research partners on the UNR campus and elsewhere in California, Oregon, and Idaho.

Research data management

Research data management is not centralized at UNR (2020). UNR Libraries offers guidance and consulting on data management practices. UNR R&I piloted a digital repository (ScholarWorks) in 2015 to enhance the

research support available to the university's faculty members. ScholarWorks assists in collecting, preserving, and distributing the university's intellectual output, with a focus on theses, papers, and other final products.

Cyber security

UNR's strategic approach to cyber security is a policy-driven data classification methodology, combined with strong technical safeguards and proactive user engagement. A robust border network control and monitoring system is in place (2016) using a combination of layer 7 application firewalls and network inspection using fiber taps and SDN switches to distribute flows to a CERT NetSA SiLK capture system. Big data analytic environments built to grow to 100Gbps were updated in 2018 to handle anticipated demands of **the Research & Science DMZs**. UNR subscribes to the **InCommon** certificate service to increase utilization of encryption for all online services. UNR maintains an inclusive Identity Management System that allows auto-provisioning for students, faculty members and staff and accommodating guests, affiliates, and visiting scholars. Shibboleth is the primary authentication gateway for all federated services. A campus wide Active Directory (AD) environment provides a multi-platform authentication and authorization system to all constituents. **Eduroam** authentication was enabled on the campus in 2016.

Human infrastructure

Both central and distributed IT staff have been incrementally increased to support the growing UNR research community. Security staffing doubled in 2015-2018 from 2.5 FTE to 5.0 FTE. Central IT HPC engineering personnel grew to 3.0 FTE in the same time period. The College of Engineering division hired dedicated IT support personnel (1.0 FTE) that are shared (60/40) with OIT Cyberinfrastructure (OIT CI).

A Director of Cyberinfrastructure position in central IT was hired in 2018. This doctoral-level position is funded internally with support from the central administration, and maintains a research-focused relationship with an academic department. The position reports directly to the CIO, and is tasked with defining and pursuing the institutional vision for CI, building significant external engagement, and assisting with high-priority research projects. The HPC team (2.5 FTE) moved under OIT CI in 2020. OIT CI also supports a software infrastructure developer (0.5 FTE), a networking facilitator (0.25 FTE) and a Graduate Assistant.

The CiC was formed from a University Technology HPC subcommittee in fall 2016, reporting jointly to the VPRI and the CIO. A member of this committee also retains a seat on the University Technology Council to provide coordination and communication. In addition to oversight of UNR HPC resources, the CiC provides faculty-led coproduction of the CI vision along with the CI Director, identifying and prioritizing development and training needs and opportunities across the campus (see *Appendix C: CiC Bylaws*).

Training and support in the use of CI resources on and off campus remains limited for both faculty members and students (2020), and this is linked directly to the need for additional human infrastructure with technical expertise, research experience, and institutional knowledge. A pilot GRAD778 course (2019-20) in Research Computing is organized and sponsored by faculty members, the Graduate School, and UNR OIT.

Resulting challenges

The current (2020) infrastructure described above, while incrementally helpful, does not adequately address the needs of a rapidly expanding faculty population or alleviate their increasing administrative workloads in the pursuit of Carnegie R1-level leadership in research and teaching. Each of these infrastructure areas requires substantial improvements in support, scale, training, and access.

UNR's CI Program and Research aspirations face the primary challenges of: 1) historical under-investment in technology statewide; 2) lack of systematic engagement with the national CI community of practice; 3) the explosion of the Internet of Things and digital data sources; 4) recent emergence of national practices and standards for end-to-end data management; and 5) a tight employment market for technology and data professionals in the west coast region.

UNR's CI vision includes transforming key organizational strategies and practices to address and overcome these challenges as a Carnegie Highest Activity Research University and flagship higher-education institution in Nevada. A working timeline of necessary investment to coordinate a critical mass of CI resources and evolve the CI Program is in *Appendix D: Program Development Timeline*.

D. CURRENT CI PROJECTS & EFFORT (2020-21)

CI Program Development

Development of an institutional CI Program continues, including updating this document (2-yr cycle), internal and external engagement, proposal writing, high-priority project assistance, consolidation of institutional CI resources, and general pursuit of goals detailed in Appendix A: Detailed Program Goals (2020-2024).

Deliverables include presentation of a CI Program scope (summarized in this document and appendices), submission of proposals to CI-oriented funding programs, internal engagement on CI issues with research division and institutional leadership, representation of UNR as a member of regional and national research CI organizations, and external engagement with local industry and civic partners in CI-related projects.

2020 milestones:

- Consolidate and align HPC operations with campus CI Program objectives
- Expand CI workforce development inside OIT with GRA and project-funded FTE
- Engage UNR research community with campus CI workshop and updated CI Plan
- Align OIT Enterprise and CI architecture planning and mutual support capability
- Participate in building-block projects funded by extramural CI-focused programs (e.g., NSF-CC*)
- Actively participate in collaborative CI platforms (e.g., PRP/NRP pilot)
- Continue support of priority projects with CI industry and civic partners
- Collect institutional data on UNR researcher needs, technology gaps, and CI successes

2021 milestones:

- Establish CI Program in institutional strategic planning and investment
- Expand CI workforce development outside OIT with partner divisions
- Engage internal research community with consolidated, clear, and stable CI services profiles
- Begin HPC phased planning/replacement with design for new research computing architecture
- Participate in building-block projects funded by extramural CI-focused programs (e.g., NSF-CC*)
- Pilot facilitated access for UNR researchers to collaborative national CI platforms
- Continue support of priority research projects with CI industry & civic partners
- Establish measures of CI impacts on institutional productivity and development

Research and Science Networks

Creation of fast, dedicated network circuits, security solutions, and pilot automation for friction-free research connectivity, including data transfer, instruments, servers, and sensor networks. SDN-WAN pilot and integration. PerfSonar monitoring, advanced port-level telemetry, data transfer applications, and dedicated DTN infrastructure. *Funding: NSF-OAC CC* (#1827186: \$495k, 2021; #2019164: \$995k, 2023).*

Colocation services

Establishment of facilities, procedures, and policy for organized colocation of researcher computing equipment in OIT-managed facilities. Colocation access, appropriate use, and eligible equipment decisions are set by OIT data center managers. *Funding: existing OIT FTE and operational budgets.*

Computing services

Maintain and operate the HPC cluster Pronghorn, evolve recharge center and annual allocation mechanisms. Focus on HPC adoption, priority projects, and student research services. Evolve and streamline administrative workflows, user features, and research impact measures. Develop lifecycle plan for next institutional capital investment. *Funding: existing OIT operational and CI FTE budget.*

Evolve research-facing remote/virtual computing and storage services in coordination with OIT Operations and architectures. Facilitate access to these and related applications for division technology leads and individual researchers. Review, pilot, and recommend a range of architectural solutions. *Funding: existing OIT FTE and operational budgets; priority project/extramural grants.*