

Campus Cyberinfrastructure Plan - University of California-Riverside

Introduction

The University of California, Riverside (UCR) stands as a vanguard of transformative research and scholarship, deeply rooted in historical advances in agriculture but now reaching across an expansive disciplinary horizon. With a commendable 70% growth in research funding over the past decade, UCR is steadfast in its commitment to its mission pillars, encompassing distinctive transformative research, an empowering learning environment, inclusivity, public good, and sustainability. UCR's research trajectory, from its pioneering advancements in agriculture to domains like arts, humanities, education, and STEM, demonstrates our global academic imprint. Positioned uniquely within the UC system, we offer solutions to regional challenges, while also resonating on a global scale. Our recent accolades, from joining the Association of American Universities (AAU) to being designated a Hispanic serving institution (one of only 4 in the AAU), further attest to our academic and societal impact. Central to UCR's endeavors is the Information Technology Solutions (ITS) team. Working synergistically with the High Performance Computing Center and as well as departmental and college-level IT teams, ITS anchors UCR's comprehensive research computing infrastructure.

As global challenges become multifaceted, UCR's vision for the research CI is designed to be an integrative, collaborative hub, facilitating interdisciplinary undertakings. To elucidate, our research CI strategy pivots around three core tenets:

1. Platform - This represents the foundational infrastructure, geared towards streamlining data, computational tasks, and user identity across diverse research computing landscapes.
2. Portal - Envisioned as a one-stop assistance center, the Portal will empower researchers to swiftly navigate, identify, and harness the resources they require.
3. Enablement - More than just a strategy, Enablement embodies UCR's ethos of continuous growth. By offering sustained mentorship, up-to-date resources, and tailored guidance, we aspire to uplift our researchers, ensuring they remain at the forefront of their respective fields.

In essence, as UCR charts its research future, our Campus CI, steered by the ITS team and its partners, aspires to be the cohesive tapestry weaving together diverse research streams. Through strategies such as a unified storage layer and enhanced network access, we are set to redefine the boundaries of the global research ecosystem.

Service and Consultation

Research Computing Facilitation

Research computing at UCR is served by the UCR Information Technology Solutions (ITS) organization as well as the distributed campus IT partners and the High Performance Computing Center (HPCC). ITS, with close to 180 FTE, provides support across multiple service areas including enterprise IT, wired and wireless networking, cybersecurity, identity management and research computing. Within ITS, the research computing team (which includes student workers as well) coordinates and supports campus-wide and extramural research computing initiatives. The goal of the ITS org is to provide seamless access and a unified front to the diverse research computing options available to our researchers and to provide thought leadership of research computing for the internal as well as partner research ecosystem.

Compute

HPCC

The High-Performance Computing Center (HPCC) provides state-of-the-art research computing infrastructure and training accessible to all UCR researchers and affiliates at low cost. Housed in a 600 sqft server room in the Genomics Building (Rm 1120A), the HPCC boasts over 6,500 CPU cores (60% Intel, 40% AMD) with 512-1024 GB RAM per node, a diverse GPU set including 6x K80, 2x P100, and 8x A100, and a robust 3PB GPFS storage system scalable to >50PB, all interconnected via a 56Gbps IB network and managed by the Slurm queueing system. Currently, it supports over 150 research groups with more than 800 active users. Its resources are also heavily used for instructing undergraduate and graduate classes in a wide range of computational, statistical, life science and engineering disciplines.

Nautilus

Through the National Research Platform (NRP), we offer our researchers access to Nautilus, a ~300 node distributed cluster. NRP encompasses a suite of programs, infrastructure, and guidelines tailored for scalable growth. Nautilus stands as NRP's premier computational, storage, and networking asset, equipped with CPUs, GPUs, and FPGAs. This system is segmented into specialized subsystems adept at handling diverse tasks from data science and simulations to machine learning and artificial intelligence. Additionally, it facilitates data retrieval via a federated, national-scale content delivery network.

Cloud

The Information Technology Solutions (ITS) team champions a multi-cloud strategy, fostering partnerships with leading cloud service providers: Google Cloud Platform (GCP), Amazon Web Services (AWS), and Microsoft Azure. Taking a consultative approach, ITS collaborates closely with our research community, understanding specific project needs, cost constraints, and use cases. In partnership with ITS, researchers can expect to utilize the cloud for tailored workstations, scalable high-performance computing clusters, flexible cloud storage, direct access to advanced AI & ML tools, and streamlined DevOps tools for seamless software development and project delivery.

Storage

UCR provides faculty with a range of storage solutions, each intricately tied to specific compute resources, resulting today in a siloed storage infrastructure. These offerings span from HPCC-linked storage and NFS for collaboration to dedicated backup services. While specialized college and departmental clusters address niche research needs, there's a recognized growth in the need for expansive, agile storage systems adept at vast data processes.

Networking

UCR has made significant investments in its electronic communications backbone to bolster research computing. The current campus network, robustly integrated with 10Gbps and 1Gbps connections to campus facilities, interfaces seamlessly with the California Regional Optical Network (CalREN/CENIC) through multiple 10Gbps switches. A standout feature is the Air Blown Fiber (ABF) network, spanning over a million linear feet, offering dynamic fiber-on-demand capabilities tailored to the evolving demands of research. Encompassing more than 50,000 endpoints, our holistic network architecture, strengthened by 10Gbps routers, embodies a fusion of reliability and expansive reach. The responsibility of managing this infrastructure is shared amongst various UCR teams, creating a synergistic environment that prioritizes the optimal functionality and security of the Science DMZ private research network. This vigilant

management employs a spectrum of tools, from VPNs to advanced firewalls, and utilizes the capabilities of both perfSONAR and tailored monitoring systems.

Science DMZ

Established with the financial backing of a previous NSF grant, the Science DMZ stands as a monumental stride in bolstering data transfer velocities and fostering interdisciplinary collaborations. At its heart lies the Cisco Nexus 7710 switch, specifically tailored to meet the data-heavy demands of our research ecosystem. Operating in its exclusive realm, detached from other networks, the DMZ boasts an armor of router-based ACLs and is powered by a dedicated fiber infrastructure. This network demonstrates unmatched adaptability, championing virtual networking, capitalizing on SDN-capable electronics, and optimizing a 100 Gbps direct linkage to the CENIC/CalREN network. With the perfSONAR monitoring seamlessly woven into its fabric, the DMZ guarantees unwavering performance metrics, while the DMZ Advisory Committee ensures that its mission aligns holistically with UCR's encompassing research aspirations.

IPv6 Integration

Taking a leap towards future-ready network configurations, all research VLANs on campus are now IPv6 enabled. This means that the IPv4 Gateways, which have been the bedrock for our research computing networks, now also possess an IPv6 gateway IP and route directly to CENIC on the HPRN (Science DMZ) network.

Data Center and Colocation

UCR provides researchers with colocation support both within and outside the campus. These facilities cater to researchers' needs by accommodating a range of systems, including computational clusters, storage systems, file/application servers, and backup systems. Every UCR colocation facility ensures physical security, maintains optimal environmental conditions, is equipped with UPS and generator backup, and benefits from round-the-clock monitored supervision.

Cybersecurity, Privacy, and Identity

The ITS team at UCR offers several integrative services to bolster our research computing infrastructure, with cybersecurity and identity management being paramount. The Information Security Office (ISO) not only offers foundational security services to the campus but also provides specialized research security services. These foundational services encompass tools and measures like intrusion detection, SEIM, vulnerability management, and multi-factor authentication. Specifically for research, the ISO has established a Research Information Security Plan template and collaborates closely with the Central IT Research Support team on cybersecurity matters. The ISO's strategic vision - the UC Riverside Secure Trust program, is anchored in five pillars, including Identity and Access Management and Data Security, and is complemented by an evolving research information security initiative. Concurrently, our identity and access management provides researchers with unified, multi-factor access across various platforms via the Duo security system. As a member of InCommon and the UCTrust federation, UCR also supports InCommon Silver level assurance, facilitating single sign-on access for researchers to federal sites.

External Collaborations

Beyond our primary infrastructure and additional departmental clusters, we diligently facilitate both access and training for our researchers to harness the full potential of computing resources offered by our external collaborators, such as the San Diego Supercomputer Center, Open Science Grid, and XSEDE.

We also link our researchers to NSF/NIH-endorsed cloud pooling endeavors like Cloudbank and STRIDEs. Our overarching aim is to align our researchers with the computing resource most suited to their specific requirements.

Sustainability

UCR's commitment to a sustainable cyberinfrastructure is anchored in its robust governance and strategic planning. The IT Research Advisory Board, deeply embedded within UCR's intricate governance structure, functions collaboratively with diverse stakeholders on campus, from faculty researchers to the library. Their mandate is expansive, ensuring that research remains secure, meets the stringent regulatory requirements, and that cyberinfrastructure services are optimally utilized by the research community. Moreover, this board is seamlessly integrated with the campus IT Strategy Council, a consortium of senior campus leaders, including the Provost, CIO, and several Vice Chancellors. This holistic approach is further complemented by Information Technology Solutions' adherence to a diligent 5-year renewal and replacement cycle for all IT assets, ensuring UCR's infrastructure remains at the pinnacle of modernity. The future-focused Campus 2030 Strategic Plan underpins our operational sustainability. It envisions an augmentation in research funding, investments in core and renovated research facilities, expansion in doctoral and postdoctoral programs, increased national recognition for UCR, and strategic investments in on-campus research centers.

Research Cyberinfrastructure Vision - Platforms, Portal and Enablement

UCR recognizes the abundance of cost-effective research computing resources available both on and off-campus. In line with the growing inclination of funding agencies towards grants utilizing existing or cloud-based computing facilities, our strategic vision emphasizes connecting faculty and researchers to these resources, be it on-campus, at national labs, supercomputing centers, or public cloud providers. While there might be instances where Research Computing might develop or host solutions on-campus, our primary approach leans towards leveraging or expanding existing resources. The rapid evolution in technology and discovery is molding our trajectory. UCR envisions a holistic, campus-wide research support system, integrating advanced research technology solutions with comprehensive training and facilitation. We firmly believe in an interdisciplinary, data-driven and AI augmented future for research. To truly capitalize on the potential of advanced computational tools, researchers require a diverse computing infrastructure. Our research Cyberinfrastructure, supported by UCR ITS, is a blend of computing platforms unified by a communications and collaboration layer. Our goal is to augment this architecture, focusing on three pivotal aspects: **Platforms, Portal, and Enablement**.

Platforms

The "**Platforms**" facet is geared towards constructing integrative layers that seamlessly knit together our multifaceted resources. By championing the principles of interoperability and portability across software, data, and identity, we aim to sculpt a versatile and dynamic ecosystem, tailored to cater to the multifarious and evolving requirements of our research community. In the near term, we aim to establish a hybrid cloud infrastructure capable of addressing sudden computational demand surges. Central to this vision is the integration of expansive storage across diverse platforms via a shared data fabric based on Ceph in order to facilitate the research data lifecycle in a seamless fashion across all resources. In addition, we are also planning to deploy an on-premises GPU cluster that leverages our data fabric. This GPU cluster will not only enhance our computational capabilities but also provide a robust platform for machine learning and data-intensive applications. To ensure seamless workload interoperability across environments, we are prioritizing the adoption of Containerization and Infrastructure as Code (IaC), complemented by other orchestration tools. Looking forward, our strategic intent is to develop a more automated and integrated

research ecosystem. We anticipate a framework where workload initiation and scaling are both seamless and automated. By employing sophisticated tracking systems, we aim to optimize platform utilization, thereby maximizing research output for UCR and our partner ecosystems. To increase accessibility to our shared storage and compute fabric across the entire campus, we look to improve network connectivity across campus from 10Gb links to greater speeds.

Portal

The **"Portal"** dimension of our strategy zeroes in on accessibility and discoverability. By orchestrating a centralized hub that meticulously catalogs our available resources, we ensure that our researchers can effortlessly pinpoint and harness the tools and platforms pivotal to their work. The key goal would be to add an "assistance" layer on top of our resources to increase accessibility. Our immediate objectives are centered around enhancing the accessibility and relevance of our computing resources for researchers. We aim to construct an assistive layer that can efficiently retrieve and extract pertinent computing information. This layer will serve as a bridge, matching researchers to the precise resources tailored to their needs, encompassing both internal and external assets. As we project further into our strategic roadmap, our ambition is to amplify the capabilities of this assistive mechanism. Leveraging the power of GenAI, we intend to automate the processes of knowledge extraction and analytical workflows. This advanced system will not only utilize the proprietary research data housed at UCR but will also integrate insights from a broader spectrum of open-source scientific libraries and datasets.

Enablement

Augmenting the Platform and Portal is our **"Enablement"** initiative. Far more expansive than mere training, "Enablement" encompasses a holistic suite of services – from hands-on training sessions and expert consultations to dedicated implementation support, all delivered by our seasoned team of research computing specialists. This multi-pronged approach ensures that our researchers are not merely equipped with state-of-the-art tools but are also adeptly guided, from conceptualization to realization, in their research endeavors. Collectively, the Platform, Portal, and Enablement strategy heralds a transformative leap towards a unified, agile, and empowered research milieu at UX.

Conclusion

The University of California, Riverside (UCR) is a world-renowned research institution committed to transformative research and scholarship. As UCR continues to grow and evolve, its cyberinfrastructure must keep pace with the ever-changing needs of its researchers as well as the needs of its diverse student population and the region of California that it serves. The UCR Campus Cyberinfrastructure Plan outlines a vision for a holistic, campus-wide research support system that integrates advanced research technology solutions with comprehensive training and facilitation. This plan focuses on three core tenets:

1. Platforms: Building integrative layers that seamlessly knit together UCR's multifaceted resources, fostering interoperability and portability across software, data, and identity.
2. Portal: Creating a centralized hub that meticulously catalogs available resources, enabling researchers to effortlessly pinpoint and harness the tools and platforms critical to their work.
3. Enablement: Offering a comprehensive suite of services, including training, consultations, and implementation support, to ensure researchers are not only equipped with state-of-the-art tools but also adeptly guided in their research endeavors.

By implementing this plan, UCR will position itself at the forefront of research innovation, empowering its researchers to make groundbreaking discoveries and address the world's most pressing challenges.