RMACC Cyberinfrastructure Plan

The Rocky Mountain Advanced Computing Consortium (RMACC) is a collaboration among academic and research institutions located throughout the intermountain region. Our mission is to facilitate widespread effective use of cyberinfrastructure throughout the region by:

- Educating graduate and undergraduate students, faculty, researchers, facilitators, and industry partners on the use of computational science, high performance computing, advanced networking, virtualization, and data management.
- Coordinating multi-institutional efforts to advance research, practice, and education in computational science in order to address important regional problems.
- Bringing together a broad range of research computing staff, researchers, faculty, and industry partners with a depth of experience and expertise not available at any single institution and facilitate their collaboration in multi-disciplinary and multi-institutional teams.

Since advancing the regional cyberinfrastructure is part of the goal of the consortium, RMACC institutions are focused on improving the regional cyberinfrastructure, especially people, networking, shared compute and data resources going forward.

1 Broadening participation

RMACC has organized an annual conference with diverse participation from students, researchers, industry and cyberinfrastructure professionals. The last two years, we have had more than 270 participants, with over 100 students participating as well as a networking track that is part of the NSF Campus CI outreach project. The top four ranked students in the poster competition participate at SC conferences paid for by the RMACC. The four students that participated in the SC15 conference came from a variety of backgrounds but had similar experiences. They were all able to share their unique backgrounds and knowledge with other attendees and simultaneously learn new skills, research methods, and meet with students and researchers from around the world.

The system administrators meet four times a year to update on campus problems and solutions, and are in the process of creating a community of research computing sysadmins in the Rocky Mountain region that draw on each other's expertise. Similarly, RMACC has become one of the regional XSEDE regions for the campus champions program. This will enhance the collaboration of cyberinfrastructure professionals across the region to improve access to local, regional, and national cyberinfrastructure and to further enable workforce development.

2 Networking

2.1 Advanced Networking

CSU: Several years ago, CSU began to establish dedicated 10 Gbps connections to research groups, based upon their need for ultrahigh speed. However, subsequently, we abandoned this approach in favor of building a “Research DMZ” in the main data center where we provide space, power, cooling and electrical connectivity free of charge for systems and applications which have ultrahigh data transport needs and have been approved by an oversight committee to participate in the Research DMZ. Five of these connections have been deployed, including to the

Cray HPC system and the Globus file server. All of these use the shared research DMZ that augments generic networking needs to the buildings, so as to avoid a bottleneck in performance for research applications.

CU-Boulder: The science DMZ connects the compute resources to a "home" and "projects" file system located in the Computing Center (Comp), provides access to the PetaLibrary from locations on and off campus, brings individual dedicated 10 GigE circuits to various campus locations, and connects storage located at the National Center for Atmospheric Research (NCAR) through the Front Range GigaPoP (FRGP). CU-Boulder participates in Internet2 and is an active member of the FRGP. Additionally, RC offers private VLANs to institutes and departments participating in the science network. CU-Boulder deployed DYNES equipment that is integrated into the current science network. The DYNES equipment was installed in a networking rack in the Telecomm building next to the equipment that supports our connection to the BiSON network operated by the FRGP which then connects to both the National LambdaRail network and the Internet2 ION network through the Western Regional Network. Several research entities (e.g. High Energy Physics, National Snow & Ice Data Center) use this network for off-site data mobility across the R&E networks.

University of Utah: The University of Utah maintains a fully redundant campus backbone network that serves the administrative, academic, and research needs of the main campus and the needs of the University of Utah Hospital and Clinic infrastructure. The University maintains multiple security zones in order to balance the data security and compliance requirements of the various constituencies. Multiple layers of security controls apply to various zones. The campus backbone supports technologies such as Multi-Protocol Label Switching (MPLS), Virtual LANs (VLANs), Virtual Routing instances, and both air-blown fiber, and conventional optical fiber. The campus backbone fully supports both IPv4 and IPv6 routing (BGP, EIGRP, OSPF, and OSPFv3) and multicast routing. The campus backbone supports multiple 10-Gbps and 40-Gbps backbone links today, and the campus uplinks to the Utah Education Network (UEN) via 100Gbps. The University of Utah campus connects to the new off-campus Downtown Data Center and to its upstream Internet providers by the University/UEN metropolitan optical (DWDM) ring that extends over the greater Salt Lake City area. This Metro Optical network supports multiple 10-Gbps and 100-Gbps wavelengths, up to a capacity of 8.8 Terabits per second. UEN currently connects via this optical network to the Internet2 Network at its intermountain regional node in western Salt Lake City with a 100-Gbps link and supports the University of Utah with multiple 10-Gbps connections and a 100-Gbps connection.

2.2 IPV6 Deployment

CSU has been experimenting with IPv6 since I2 made a block of IPv6 addresses available in 2006. As a result, we are poised to implement a transition plan and to deploy it fully when required.

CU-Boulder is currently running IPv6 in dual-stack mode on the campus backbone (core and distribution routers) but are not routing any user subnets yet. On the border, we have IPv6 BGP peering with our regional (FRGP). Our current IPv6 allocation is an assignment that belongs to Internet2 (via FRGP) but we plan to petition ARIN for a block dedicated to CU-Bolder. Early deployment and testing is expected to take place inside the science DMZ.

Utah—The UofU's campus backbone has supported IPv6 for over a decade and the growth in IPv6 enabled endpoints continues to increase. The UofU utilizes a combination of firewalls,

unicast Reverse Path Forwarding (uRPF), and router Access Control Lists (ACLs) to protect the environment and to prevent spoofed traffic from originating from any on-campus network. .

2.3 BCP 38

CU-Boulder: We employ the following measures for the prevention of IP spoofing (BCP 38/RFC2827).

1. On all campus distribution routers, unicast reverse path forwarding (uRPF) is configured on all gateway interfaces with the following syntax: ip verify unicast source reachable-via rx allow-default.
2. On the border routers we configure explicit packet filters to discard spoofed CU address space from entering the network border and private (RFC1918) address space from leaving the network border.
3. On the science DMZ border router platform (Arista), uRPF is supported.

CSU: We employ the following measures for the prevention of IP spoofing (BCP 38/RFC2827).

1. On all campus distribution routers, unicast reverse path forwarding (uRPF) is the standard configuration for all gateway interfaces with the following syntax: ip verify unicast source reachable-via any allow-default. Exceptions are allowed where dual-homed Linux hosts cause conflicts with this setting.
2. On the border routers, we configure explicit packet filters to discard spoofed CSU address space from entering the network border, and to discard both private (RFC1918) and non-CSU public addresses from leaving CSU via the network border.
3. The ScienceDMZ uses a dedicated interface on the border router, with a packet filter preventing private (RFC1918) and non-science DMZ addresses from leaving CSU via the network border.

Utah: The university actively monitors traffic with various security tools and subscriptions to global blacklists in order to identify and eliminate spoofed or nefarious traffic. These tools also help provide an audit trail for BCP38 compliance, as well as regulatory compliance such as HIPAA, FERPA, FISMA, PCI, etc. The Utah Education and Telehealth Network reinforce BCP38 compliance at a regional level.

2.4 Security

CU-Boulder: We use BRO to monitor our ScienceDMZ and require two factor authentication on all systems managed by Research Computing.

CSU: We monitor traffic on our ScienceDMZ via border router/firewall logs, and are planning an implementation of BRO for the ScienceDMZ in the 2016/17 academic year. Two-factor authentication is currently being rolled out to central IT staff and researchers.

Utah: We collaborate with the campus Information Security Office (ISO) on the implementation of new campus security policies and procedures that are focused on localizing risk and performing careful risk assessment reviews on a regular basis. We are also enhancing the existing computing environment supporting secure use and storage of data containing personal health information and also supporting other forms of research data restricted under other compliance regimes (e.g., FISMA, FERPA, Export Control). We also work closely with the campus networking team and ISO to implement best current practices to prevent IP and e-mail spoofing as well as other forms of fraud and misuse involving University research computing resources.

2.5 Middleware

CSU was an early participant in InCommon, and was the first institution in the region to deploy Shibboleth, as both an Identity Provider and a Service Provider. CSU has widely extended its use of Shibboleth over the past two years, to both external and internal entities, and was the first in the world to deploy shib for Electronic Books Library (EBL). We have just completed the upgrade to shib ver. 3.0. We are considering participating in Internet2's TIER activity that is emerging nationwide to provide new and more robust IAM services.

CU-Boulder is an Identity Provider under InCommon. RC has worked with Globus to enable InCommon login to Globus Online using the CU-Boulder credentials. Additionally, RC and the Libraries worked the DMPtool https://dmp.cdlib.orgto accept CU-Boulder credentials. RC is working with the OIT identity management group to accept credentials from other institutions. Right now, RC accepts two factor authentication from NCAR, CSU, and CU-Denver including the Anschutz Medical Campus.

Utah—The University of Utah participates in Shibboleth and InCommon and has active development in expanding its role based Authentication, Directory, and Provisioning services. The UofU is expanding its multi-factor authentication across its academic, research, and healthcare. The UofU also has active work with collaborators from Clemson, CoSign and other groups in developing NSF funded Fedushare which will have the ability to bring federated services to the command-line within HPC and stand-alone servers

3 Storage

CSU: We have an excellent approach, our Large-scale Storage Appliance (LSA), to meet elerging storage needs. We have added such storage appliances to our Cray HPC, our digital repository, and to more than fifteen additional systems. Our plan for general storage for research data is two-fold: 1) we have developed a central virtualized server/storage/backup infrastructure that offers storage capacity at 30¢/GByte/year (not backed up) or 55¢/GByte/year (backed up), and 2) we have offered workshops to campus IT support staff to build their own LSAs. In the summer of 2012, we offered this workshop ('Build your own Large-scale Storage Appliance') to seven institutions of higher education in Colorado, and it was very successful. Moreover, several LSAs have been deployed locally on campus too.

CU-Boulder: Reliable, high-performance, professionally managed raw data storage is a fundamental building block of CU Boulder's research cyberinfrastructure. Simply stated, while users of computing (e.g., via their own clusters, through a campus condo cluster, an NSF computing center, or commercial clouds) can tolerate periods of downtime or limited availability, raw storage must be highly reliable and persistently available.

The research computing group is providing a scalable (multi-petabyte) storage infrastructure with different service levels. Different researchers can tolerate different risk storing their data. We are providing the following services:

- Active Storage
- Archival Storage
- Replication
- Copy of data on tape

Utah: Utah specializes in the deployment of cost effective, engineered, and reliable storage solutions to meet the diverse research computing needs on campus. Currently over 14 PB of active disk, mostly purchased by researchers at the cost of the hardware, are deployed in tiers

ranging from high performance parallel file systems (Lustre, Isilon), to RAID6 for user project space and home directories, and most recently to Ceph based object storage solutions. Various policies on archive, replication, and multiple copies are possible among the tiers depending on the willingness of the researcher to pay, for example for extra copies or shadowing. CHPC at the UofU also has deployed dozens of data transfer nodes at 10 Gbps and 40 Gbps for enabling data transfer needs on the UofU ScienceDMZ to bypass the campus firewall for more efficient data transfers, typically using Globus. Handling of restricted or protected data is also a strength and CHPC has maintained a siloed protected environment of HPC compute, virtual machine compute, and storage resources, with two-factor authentication and enhanced security, auditing and control.

4 Regional RMACC supercomputer

The Summit supercomputer, funded by NSF under Grant No. ACI- 1532236, is currently in acceptance testing. The system has peak performance of over 400 TFLOPS. The 380 general compute nodes have two Intel Haswell CPUs with 12 cores each, 128 GB of RAM, and a local SSD. Additionally, the system has 10 GPU nodes containing two NVIDIA K80 GPUs, 5 high-memory nodes with 2 TB of main memory; and a second phase of deployment is scheduled in December 2016 20 Xeon Phi ("Knight's Landing") nodes with 72 real cores supporting 288 threads for development and benchmarking. All nodes are connected through a high-performance network based on Intel's Omni-Path with a bandwidth of 100 Gb/s and a latency of 0.4 microseconds. 1 PB of high-performance storage is provided using the IBM GPFS file system. This system is available to CU-Boulder researchers and collaborators, as well as 10% of cycles are provided to members of the Rocky Mountain Advanced Computing Consortium.