The Pacific Research Platform: Making High-Speed Networking a Reality for the Scientist

Larry Smarr
Calit2
UC San Diego
lsmarr@ucsd.edu

Camille Crittenden
CITRIS & Banatao Institute
UC Berkeley
ccrittenden@berkeley.edu

Thomas DeFanti
Calit2/Qualcomm Institute
UC San Diego
tdefanti@ucsd.edu

John Graham
Calit2/Qualcomm Institute
UC San Diego
jjgraham@ucsd.edu

Dmitry Mishin
San Diego Supercomputer Center
UC San Diego
dmishin@ucsd.edu

Richard Moore
Cygnus Consulting
r3moore@ucsd.edu

Philip Papadopoulos
San Diego Supercomputer Center
UC San Diego
ppapadopoulos@ucsd.edu

Frank Würthwein
San Diego Supercomputer Center
UC San Diego
fkw@ucsd.edu

ABSTRACT

While NSF's recent Campus Cyberinfrastructure investments have catalyzed an enormous leap in campus networking capabilities, it remains necessary to integrate these capabilities into the routine workflows of researchers transferring data among their remote collaborators, data repositories, and visualization facilities. The Pacific Research Platform (PRP) is a program to develop a science-driven, data-centric "freeway system" by federating campus Science DMZs into a regional Science DMZ. Research collaborations across PRP partners serve as use cases for deploying the hardware/software and addressing security/policy/social issues to achieve high-performance data transfers between researchers. The regional PRP represents a manageable scale for initial deployment efforts and capturing lessons learned, but the PRP also informs national deployments – i.e., an eventual National Research Platform (NRP). The recent *First NRP Workshop* included many campuses and networking organizations and addressed scientific requirements, scaling approaches, and socio-technical issues to extend the PRP to a national level; a second NRP Workshop will be held in August 2018. In this paper, we describe efforts to build the PRP, the social and technical approaches, the domain science applications of PRP's capabilities, and results of the recent NRP workshop.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.
PEARC '18, July 22–26, 2018, Pittsburgh, PA, USA
© 2018 Association for Computing Machinery.
ACM ISBN 978-1-4503-6446-1/18/07…$15.00
https://doi.org/10.1145/3219104.3219108

CCS CONCEPTS

Networks~Network performance analysis

• *Networks~Network manageability*

KEYWORDS

Cyberinfrastructure, end-to-end performance, network, science engagement.

ACM Reference format:

Larry Smarr, Camille Crittenden, Thomas DeFanti, John Graham, Dmitry Mishin, Richard Moore, Phil Papadopoulos and Frank Würthwein, 2018. The Pacific Research Platform: Making High-Speed Networking a Reality for the Scientist. *Proceedings of PEARC'18, Pittsburgh, PA USA, July 22-26 2018 (PEARC'18). ACM, 8 pages.* https://doi.org/10.1145/3219104.3219108.

1 CONTEXT AND OBJECTIVES

Since 2012, the NSF has made substantial investments in campus cyberinfrastructure [1] by awarding more than 200 competitive grants to over 100 universities (Fig. 1) to upgrade campus networks to 10–100 Gbps and to deploy Science DMZs [2] at their borders. The NSF-funded project "CC*DNI DIBBs: The Pacific Research Platform (PRP)" (award #1541349) emerged from the unmet needs of academic researchers for a data sharing architecture *between these campus Science DMZs* with end-to-end 10–100 Gbps connections, enabling region-wide virtual co-location of data with computing. The PRP is designed to serve researchers who need to exchange data between their campus labs, their collaborators at other institutions, supercomputer centers, remote instruments, data repositories, and visualization environments.

Prior end-to-end programs tended to focus on the "ends" being campus gateways or national facilities, rather than the investigators' on-campus computers and the remote data,

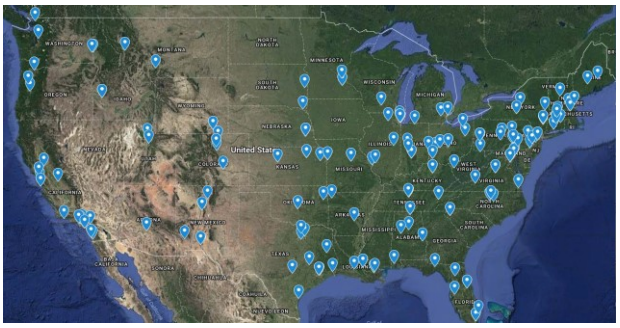


Figure 1: CC* 210 awards across 44 states 2012-2017

compute, or collaborations they sought. Facilitating end-to-end connections is difficult, because the path between the two ends traverses multiple network management domains: campus, regional, and national. No single organization owns the responsibility for providing *researchers* with high-bandwidth disk-to-disk performance. That is precisely what the PRP set out to change.

The Pacific Research Platform (PRP) is a project to develop a science-driven, data-centric "freeway system" by federating campus Science DMZs into a regional Science DMZ. Just like the highways prior to freeways, network connections already exist between scientists, their collaborators, national facilities, and data repositories. However, just as the end-to-end high-speed, no-stop-lights freeway system transformed the economy and societal mobility, a data freeway system will be transformative for 21st-century research.

The PRP initially began as a California-based effort, focused on the ten campuses of the University of California, plus Caltech, Stanford, the University of Southern California (USC), San Diego State University, the San Diego Supercomputer Center (SDSC),

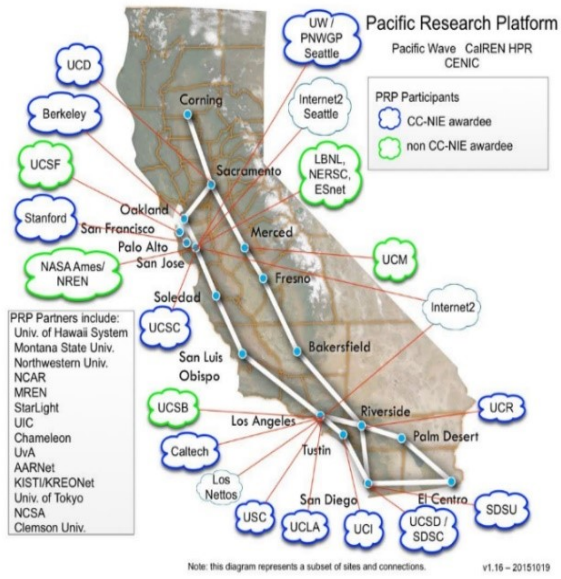


Figure 2: Pacific Research Platform Partners

Lawrence Berkeley National Laboratory, and the NASA Ames Research Center, all connected by the Corporation for Education Network Initiatives in California (CENIC) research and education network. A regional effort was considered a manageable size for the project in terms of the number of participants – large enough to support significant scientific collaborations across the partners, but small enough to enable all-to-all personal interactions and solve the numerous socio-technical issues facing the project. This scale also matches the crucial role of the regional network (in this case, CENIC) in addressing end-to-end performance issues. As shown in Fig. 2, the PRP partnership has been extended to include many campuses beyond California, including national facilities (e.g., NCAR, NCSA, TACC), three other regional networks (Pacific Northwest Gigapop (PNWGP), Front Range Gigapop (FRGP), and the Metropolitan Research and Education Network (MREN)), and international partner sites (Netherlands, Japan, Korea, Australia, Guam, and Singapore).

The success of the project requires a broad base of support among its partners, well beyond the cadre of network engineers. All campus CIOs and more than 50 domain scientists within the partnership have committed to working with the PRP team to solve the end-to-end performance issues for multi-institution collaborations and applications. The success of the PRP has sparked interest from other regional networks and campuses, leading to the PRP and CENIC co-hosting *The First National Research Platform Workshop* held in Bozeman, MT in August 2017. At that meeting, a strong national community emerged, documented in the report published on the PRP website [3]. We anticipate that the PRP's procedures, tools, and lessons learned can be scalably extended to a national/international capability.

2 SOCIAL APPROACHES

We estimate that substantially less than 10% of the labor and infrastructure investments necessary for the PRP's success are actually funded by NSF's grant; that is, the PRP relies heavily on leverage. Hence, it is a social project in many respects, coordinating the efforts of many via a "coalition of the willing," each with their own funding, objectives, management/staff, policies, and security. In this section, we describe the structures and methods used to accomplish the project's social elements.

First, an initiative like the PRP must be driven by the requirements of domain scientists, *then* facilitated by IT and network engineers. As described earlier, high-performance end-to-end networking needs to reach not only the campus gateway, but all the way to the scientist's lab or workflow. Each application brings its own unique questions that must be answered: Where are the collaborators or data repositories? What are the bandwidth and data volume requirements? What is the data access and computing workflow? What application software drives the network? Researchers and their staff understand the applications and requirements, and they must work collaboratively with IT and network engineers within their own campus, with those of their collaborators, and with the regional or national networks connecting their campuses to

solve end-to-end problems. When we recruited scientists for the PRP, we requested that they commit their staff to work with all parties to solve problems, i.e., to collaborate rather than point fingers, the latter a familiar posture when many autonomous parties have responsibilities for pieces of a complex system.

In addition to commitments from application scientists, a crucial stakeholder is the campus CIO. Campus networking infrastructures – and security requirements (e.g., firewalls) that often impede performance – are generally the responsibility of the campus CIO. For every campus in the PRP, we obtained a written letter from the CIO committing to deploy their Science DMZ and dedicate the necessary staff time to support scientists on their campuses. (Most partners already had NSF CC* awards, so this commitment was typically already made to NSF.) Each campus has its own autonomous management, staff, policies, and infrastructure. The campus CIO and their staff must be willing to support their faculty with high-speed end-to-end performance, to address the inevitable infrastructure challenges that arise, and to invest the staff time necessary to overcome them.

These initial commitments from PRP partners, combined with the enthusiastic support and a commitment of substantial resources from the CENIC regional network, laid the foundation for the PRP's success. That foundation has made governance and execution of the project much simpler and effective. The project's leadership team, consisting of the PI (Smarr) and co-PIs (Crittenden, DeFanti, Papadopoulos, and Würthwein), meets regularly to establish project priorities, plan initiatives, review progress, allocate resources, and coordinate external communications. The project's Advisory Board (Tom Andriola (chair), UC Office of the President; Richard Alo, Jackson State; Ruth Marinshaw, Stanford; and Horst Simon, LBNL) represents diverse stakeholders in advising the leadership team.

In executing the program, we have established two coordination groups that meet regularly across the partnership. First, the network engineer coordination group includes technical representatives from all campus partners and the regional networks. Focused on reviewing current operations, it brings in new campus partners and domain scientist collaborations and works through specific performance issues. This group also interacts with technical collaborators within and outside the PRP partnership who are doing closely-related work such as ESNet, Globus, NSF's CI Engineers/Cyberteams, and OSG and XSEDE staff. The second coordination group, the science engagement team, follows the engagement model developed by ESnet [4] and provides an interface between domain scientists and the project/partner technical staff, ensures that domain scientists are successfully using the system, conducts outreach to potential new users, and captures users' requirements. Members have conducted several workshops with specific data-intensive domain applications and have participated in the CrossConnects science engagement workshop series sponsored by ESNet [5].

Along with these standing cross-partnership groups, *ad hoc* teams are formed to support specific research collaborations and deploy applications with high-speed end-to-end performance. These teams include PRP-funded technical staff, representatives of the domain scientists, network engineers and IT staff from the participating campuses, and representatives of the CENIC regional network. A key lesson learned is that getting scientists to use the PRP in their day-to-day workflow involves much more effort than simply telling them to "set up a Science DMZ and use Globus." While there is a valuable learning curve from prior experiences, inevitably some level of customization is required to tailor the network/hardware/software environments to meet the specific needs of each distributed research collaboration; that requires teams that span the end-to-end environments.

Finally, workshops have been a key element of the project's social approach. We have held several PRP project workshops, with an emphasis on bringing domain scientists together with network engineers and IT experts [6, 7]. One workshop focused on big data in the earth sciences [8], another on applications in electron microscopy, and another on high-resolution video, especially in combination with UAVs. We plan to hold additional science-focused workshops (e.g., machine learning), and, as stated earlier, we have supported others' domain-specific workshops. Finally, we hosted the *First National Research Platform Workshop* [3], which will continue as an annual event.

3 TECHNICAL APPROACHES

The PRP team has conducted foundational work in high-speed end-to-end networking for more than a decade, with NSF-funded projects such as the OptIPuter [9], Quartzite [10], and PRISM@UCSD [11], which established an early 10-80 Gbps Science DMZ across UCSD and neighboring research labs. As noted, the PRP builds on existing infrastructure, technology, and people by utilizing the networking capability of CENIC and other regional networks and partner campuses and relying on technical concepts, tools, and software developed by many others (e.g., ESNet, Internet2, Globus, and others). In addition, the PRP has made advances of its own in systems and tools.

The PRP could not be successful without the CENIC regional network and the expert staff it has committed to support member campuses with the PRP. CENIC's California Research and Education Network (CalREN) is a multi-tiered, advanced network-services fabric serving the majority of research and education institutions in the state. The CalREN backbone includes roughly 8,000 miles of CENIC-owned and managed fiber, last-mile fiber, and hundreds of optical components. The backbone has three components: optical (Layer 1), Ethernet (Layer 2), and routing (Layer 3). Recent improvements include upgrading the core backbone to 100GE [12]. From the PRP's inception, CENIC management has enthusiastically supported it, and CENIC staff have been an integral part of the technical team.

Achieving high-performance connectivity in the past has often required using Layer 2 networks, with detailed point-to-point configurations, monitoring and debugging. The PRP team decided at the outset that in order to scale up to a large number of sites, even at a regional level, we should target Layer 3 networking. Over the past few years, we have demonstrated that with modern routers and sufficiently large buffers, we can achieve high-performance at Layer 3.

A foundational element of the PRP is the Science DMZ [2]. The Science DMZ concept, developed by ESNet, was included by NSF as a necessary element of the CC* awards and has been widely adopted by many campuses, national labs, and international organizations. From ESnet's website [2],

> *The Science DMZ is a portion of the network, built at or near the campus or laboratory's local network perimeter that is designed such that the equipment, configuration, and security policies are optimized for high-performance scientific applications rather than for general-purpose business systems or "enterprise" computing. ...A Science DMZ integrates four key concepts into a unified whole that together serve as a foundation for this model. These include:*
>
> - *A network architecture explicitly designed for high-performance applications, where the science network is distinct from the general-purpose network,*
> - *The use of dedicated systems for data transfer (data transfer nodes or DTNs [13]),*
> - *Performance measurement and network testing systems that are regularly used to characterize the network and are available for troubleshooting, and*
> - *Security policies and enforcement mechanisms that are tailored for high performance science environments.*

The PRP is built on a specific adaptable hardware/software data transfer node (DTN) that we developed and call a FIONA (Flash I/O Network Appliance). FIONAs are commodity, fast Linux PCs located on Science DMZs and have the copper or optical fiber network adapters and fast storage to achieve full-speed (1–100 Gbps) large data transfers and to collect measurements (see Fig. 3 for an early FIONA diagram).

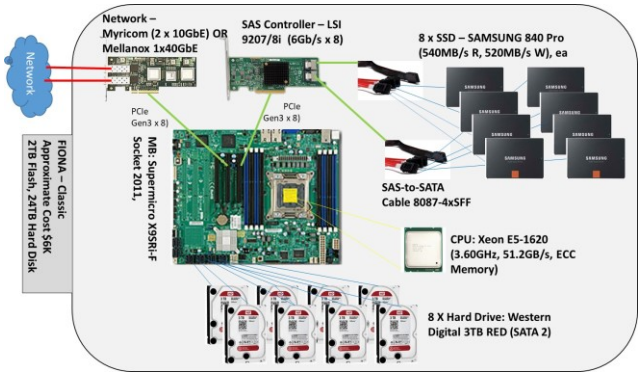


Figure 3: "Classic" FIONA Design (~2015)

While many servers can be configured as DTNs, we have found standardization on the FIONA design to be helpful as we deploy across many partners. FIONAs take some essentials from servers (ECC memory protection and server processors), optimize cost/performance/durability for storage components and are configured for syncing storage and the network interface cards. Software-only approaches for data protection (ZFS and, in some variants, Ceph) reduce cost and decrease dependency on special hardware, but do not compromise achievable disk-to-disk performance or data protection. We have designed and published specifications for variants of the FIONA design [14], including a pocket-sized 1 Gbps FIONETTE costing less than $250, a recent 10 Gbps FIONA for ~$1,000, a standard 40 Gbps FIONA costing ~$8000, a 100 Gbps FIONA at $20,000, and a 8-GPU FIONA8 costing ~$17K. Some FIONAs have 160 TB or more of disk space to provide access to local scientific databases; when amortized over four years, this incremental storage cost is ~$10/TB/yr. Opinions differ on whether DTNs should be used *only* for data transfer functions, or whether they can better serve multiple purposes, including computing (e.g., GPU-equipped systems); while the system's priority must be keeping up with data transfers, our experience is that these systems can often be multi-purposed.

A number of software applications to drive data transfers on DTNs are available (e.g., see [15-17]). The most widely used is Globus Online, an easy-to-use cloud-based software-as-a-service built on GridFTP that lets a user submit transfer requests using a web browser or command line interface and manages the data transfer process in a "fire-and-forget" model [18]. For routine testing of the PRP (e.g., to build the MaDDash monitoring displays), we use GridFTP itself. However, the PRP is science-driven and users have different needs for their specific workflows; FIONAs easily support applications that might use FTP, FDT [19], UDP [20], UDT [21], XrootD, SCP, and direct network mount (with suitable security). Scripps Institution of Oceanography users have made second copies of data from FIONAs into Google's cloud using tools like gDrive.

A basic monitoring and diagnostic tool is ESnet's Monitoring and Debugging Dashboard (MaDDash) mesh visualization tool [22]. MaDDash generally displays results from perfSONAR tests, but we have modified the tools to also import results from GridFTP. Using perfSONAR or other chron jobs, we measure GridFTP 10 GB disk-to-disk file transfers (10 GB is big enough to get TCP/IP on 10/40 Gbps networks up to speed) and iperf3 memory-to-memory transfers four times a day, traceroute every

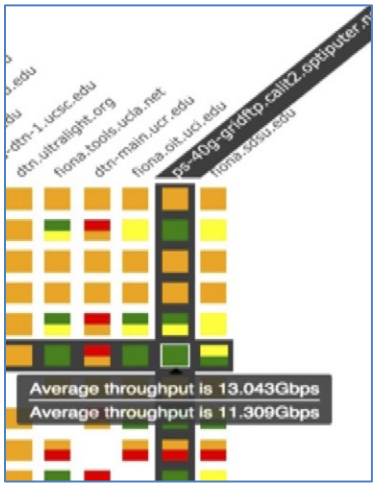


Figure 4: PRP MaDDash visualization early 2016. Green >5Gbps, Yellow 1-5Gbps, Red <1Gbps, Orange 0.

ten minutes and OWAMP (packet losses rate) constantly between FIONAs and archive the results. The results of these end-to-end network performance measurements are represented as a color-coded matrix between FIONA end-point pairs, with the upper/lower halves of each connection representing the node of the measurement initiating the test and the primary/alternative Measurement Archive (MA) if available. The direction of the transfer/test is node on Y axis –> node on X axis. For online displays, the numerical details can be accessed online by hovering the mouse cursor over a square and interactively zooming the plots available by clicking the squares. MaDDash displays are effective diagnostic tools that clearly mark the patches of networking problems once those appear – and serve as a wall of fame/shame to help motivate individual sites to stay green. A portion of an early PRP GridFTP MaDDash display from 2016 is shown in Fig. 4, with very little green showing. In a more recent MaDDash display from 2017, shown in Fig. 5, green squares dominate the matrix with only a few isolated cases of poor performance, showing how much performance has improved between FIONAs and how many more PRP sites are connected.

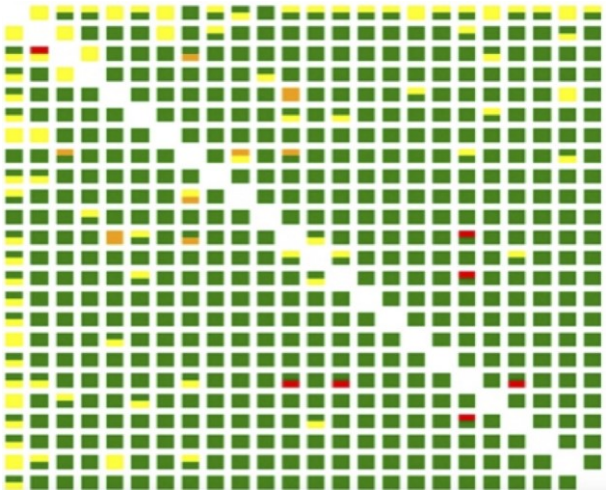


Figure 5: PRP GridFTP Dashboard, 24 sites in late 2017, in order L to R: UvAmsterdam, KISTI-Korea, Clemson, UHM, MREN, UIC, UCAR, MSU, UWash, Utah, UCD, UCSF, Stanford, UCM, UCSC, UCB, Caltech, UCLA, USC, PacificWave-LA, UCR, UCI, UCSD, SDSU. Throughput is shown for both directions. Although the PRP was designed for regional deployment, excellent results to Korea, Clemson, Chicago, etc. show distance invariance.

Another MaDDash example is shown in Fig. 6, which represents a subset of the PRP sites supporting a specific climate science collaboration. Scientists in that collaboration only need to monitor the performance of the sites they use; custom displays can be built to support various stakeholder groups.

Another key tool developed and used by PRP engineers is the traceroute visualization tool (Fig. 7) that identifies bottlenecks in the network segments. Traceroute, of course, is based on Van

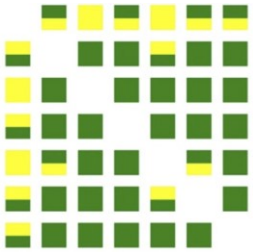


Figure 6: PRP GridFTP MaDDash details of Climate Science group data transfer speeds between FIONAs at UIC, UCSD, UHM, UCI, USC, NCAR, and SDSU

Jacobsen's work [23], but the interactive visualization developed by the PRP staff has proven helpful to automatically overlay the end-to-end test results and mark the underperforming network segments based on the combined data. The visualization tool can either retrieve data from the CENIC perfSONAR Esmond database (preferred) or make measurements directly (slower, but used in cases where some data is not available in Esmond([24].

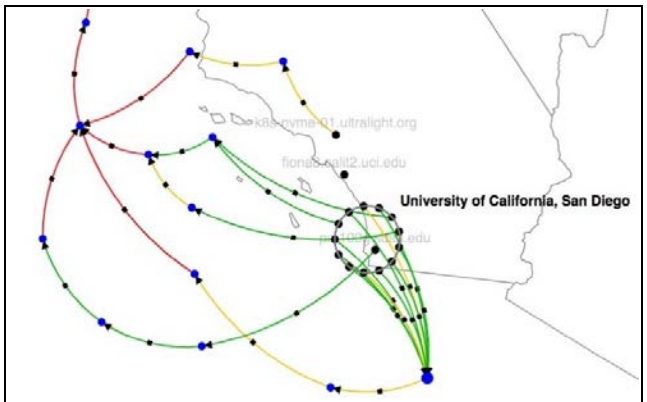


Figure 7: Traceroute Visualization Tool. Red routes are badly underperforming; green routes are performing well. Note: it is a graph, not a fiber route map. Only endpoints are geographically bound.

Currently the perfSONAR installation runs the traceroute tests periodically, and data is stored in the database. The most recent version of the tool is running in a Kubernetes [25] container and is fully integrated with other tools and Kubernetes API, which allows minimal manual setup, while the older version was run in a Jupyter notebook [26] and required manual runs to get the visualization display. In addition, interactive controls are available on the visualization web page for laying out the graph to allow adjustment and better visualization of graphs within a wide range of node numbers, densities, and analysis workflows. Each end user needs to know that the PRP network paths they are using have the bandwidth and latency and absence of packet loss that the PRP technical team predicted. Since unexpected changes can happen on high-speed networks, these types of

visual tools allow the end user *and* the networking specialists to easily check the PRP links and spot underperforming ones [27].

While MaDDash shows at a glance if all 24 PRP campuses are well connected, the traceroute visualization tool interactively shows where the bad link is and helps identify who is responsible for making the Layer 3 routes work as expected by separately marking bad throughput, retransmits, and latency segments. Again, in a complex system with many people responsible for end-to-end performance, diagnostics to identify where problems lie are invaluable.

In our experience, MaDDashes are quite labor-intensive to construct and maintain. The test endpoints must be created, a testing plan and a measurement archive must be defined - all coordinated through person-to-person interactions. To improve efficiency, the PRP has deployed a wide-area, distributed cluster on which software containers can be launched, monitored, and managed [28]. The containers encapsulate the code (e.g., transfer tool, user application, augmented transfer protocol, etc.) and can be quickly moved and executed on networked machines.

Based on extensive testing on automated mesh configuration, we selected Kubernetes (K8s) as our container orchestration system [25]. K8s was developed by Google to run software containers across its distributed compute/storage data centers, automating deployment, scaling, and management of containerized applications; it is now open source via Google's donation to the Cloud Native Computing Foundation [29]. In 2017, K8s became the *de facto* standard in industry [30] for container orchestration. We are adopting K8s in the PRP, as well as its counterpart for an open-source cloud-native object storage system, Rook/Ceph [32], for optimal integration with applications, which requires large, persistent, container-accessible storage.

Rook automates the buildout of Ceph [32], dramatically simplifying the installation and management of the storage. Ceph was developed at PRP partner UC Santa Cruz (UCSC), and PRP staff have worked extensively with Shawn McKee (U Michigan) and Harvey Newman (Caltech) to deploy and test Rook/Ceph under K8s. In the PRP, sixteen FIONAs with non-volatile memory express (NVMe), SSDs and large spinning drives added to accelerate the transfer of data, are now part of the Rook/Ceph K8s storage pool and use the latest versions of perfSONAR and MaDDash for the fully-automated network measurements. These software tools are automatically deployed on every new node as it joins the cluster so that no manual configuration is needed. All the results are sent to a central measurement archive, displayed in the dashboard, and visualized in the traceroute graph tool.

Security across the PRP is critical. We are following the improvements being made in leveraging all available Linux mechanisms by containers, accessible through the flexible K8s policies. Users' access to the system and code access to infrastructure is established, controlled and monitored based on requirements, trust level, and other variables, and is flexible enough to accommodate any foreseeable special request. A web portal integrated with the K8s API has been developed that provides easy interface for users to control their collaboration environment policy settings, while its backend monitors the cluster and automatically adjusts settings for those environments. At the same time, it allows root administrators to enable custom policies by manually adjusting the K8s policies. All Linux security features are engaged in providing the required isolation level of users' containers while allowing monitoring tools (e.g., perfSONAR) full access to the underlying infrastructure. All physical systems – FIONAs, K8s cluster nodes, databases, Rook/Ceph storage, web-servers, *etc.* – are maintained with security updates, log processing, and other routine practices. Furthermore, given the volume of data transfers that cross the network, we analyze the human/applications interfaces, policies, and procedures when working with existing campus and network intrusion-detection systems (e.g. Bro [33]).

Within our infrastructure, the PRP uses CILogon for federated identity management for user authentication in K8s through the web portal and K8s API extensions developed by PRP. We also plan to exploit physical Trusted Platform Modules (TPMs) [34] in K8s physical nodes in combination with BIOS-based secure boot to ensure that software running on K8S nodes is unchanged. (TPMs are widely available as standard hardware components, but are not widely utilized in HPC/academic environments). CoreOS retired its work on the TPM admission controller for K8s, but made it available for future upstreaming [35]. We plan to bring this code forward to the current K8s release 1.9.2+.

4 RESEARCH COLLABORATIONS ADVANCED BY THE PRP

One of the initial steps in formulating the PRP was to identify research collaborations that would benefit from end-to-end high-performance networking and then to recruit those domain scientists to commit their resources to collaborate with PRP teams to implement their application on the PRP. Twelve groups across five disciplines were initially recruited, including:

- Particle physics: Large Hadron Collider Run 2 and Open Science Grid (institutions include UC San Diego (UCSD), UC Irvine (UCI), UCSC, Caltech, UC Davis (UCD), UC Santa Barbara (UCSB), and UC Riverside (UCR)).
- Astronomy and astrophysics: (a) Telescope surveys such as Large Synoptic Survey Telescope (LSST), Dark Energy Spectroscopic Instrument (DESI), Intermediate Palomar Transient Factory (iPTF) (Lawrence Berkeley National Laboratory (LBNL), UCD, Caltech, UCB, SLAC National Acceleratory Laboratory, UCI, UCSC, and the University of Washington (UW)); (b) Galaxy evolution team (UCI and UCSC); (c) Gravitational wave astronomy (Caltech, UCSD).
- Biomedical applications: (a) Cancer genomics hub/browser team (UCSC, San Diego Supercomputer Center (SDSC)); (b) Microbiome and integrative 'omics team (UCSD, UCD, Caltech, UCSF); (c) Integrative structural biology (UC San Francisco (UCSF), SDSC, and the National Energy Research Scientific Computing Center (NERSC)).
- Earth sciences: (a) Data analysis and simulation for earthquakes and natural disasters (UCB, UCSD, UCD, UC Los Angeles (UCLA), UCI, USC, Stanford, Oregon State

University (OSU), and UW); (b) Climate modeling (National Center for Atmospheric Research (NCAR) data portal to collaborators); (c) California/ Nevada regional climate data analysis (UCSD/SIO, NCAR); (d) CO2 subsurface modeling San Diego State University (SDSU), and National Science Foundation (NSF) supercomputer facilities).

- Scalable visualization, virtual reality, and ultra-resolution video (UCSD, University of Illinois at Chicago (UIC), University of Hawaii at Manoa (UHM), UCD, UCI, UC Merced (UCM), and the University of Virginia (UvA)).

These groups represent not only diverse domain applications, but also present a variety of use cases and technical requirements for the PRP. Since the PRP started, most of these groups are now successfully using the high-performance PRP for their applications. For various reasons, some of these projects have lowered their priority for PRP engagement. At the same time, through our science engagement efforts, new applications have come to the PRP. For example, Clemson professor Alex Feltus and his collaborators (e.g., [36]) have deployed a FIONA, with a 2018 goal for SciDAS to do a demo from a Tripal database instance where the PRP uses SciDAS dashboard DTNs to move data over the PRP into a genomics workflow at Washington State University (WSU). They would then extend this DTN-to-DTN approach to include other Tripal collaborators at U Connecticut, U Tennessee, and NIH's National Center for Biotechnology Information. In another example, the High-Performance Wireless Research & Education Network (HPWREN, [37]), which supports numerous missions including wildfire detection and monitoring in southern California, uses the PRP to link their wireless network to servers at multiple PRP campuses, providing disaster recovery, data redundancy, and web access during natural disasters.

We have also realized that effective science engagement, or "research engagement" more broadly, requires a specialized set of skills that has not been traditionally cultivated in IT professionals or students. Research engagement specialists must have sufficient knowledge of network engineering to understand the concepts and tools, while also having sufficient expertise in domain-specific areas to work comfortably with faculty, postdoctoral researchers, staff and students in a given lab. We also recognize that improving high-speed data transfer capabilities would facilitate research and outreach in disciplines beyond those we initially identified, for example, in archaeology, digital humanities, and the arts. Researchers in these domains may also be more open to collaborating with cyberinfrastructure engineers than those in more established data-intensive disciplines, who may be committed to existing workflows. To effectively scale the PRP from its original prototype, we must 1) foster the next-generation workforce that will "translate" the highly technical framework of the PRP to labs and sites that will benefit from its architecture, and 2) seek out and embrace those historically non-technical disciplines for whom a high-speed network might transform their scholarship, artistic practice, or community engagement.

We recently commissioned and published a report on science engagement, which addresses the science work cycle, measurements of success, useful modalities, characteristics of science engagers, and scaling science engagement [38].

5 SCALING UP FROM THE PRP TO A NATIONAL AND INTERNATIONAL RESEARCH PLATFORM

From the outset, it was anticipated that the PRP's lessons learned on a regional scale would inform efforts to scale its capabilities to a national and international platform. During the course of the project, many institutions and science teams have expressed interest in joining the PRP or re-creating its capabilities. To facilitate this process, we convened the *First Annual National Research Platform Workshop* in Bozeman, MT in July 2017 [2]. About 135 people attended from 30 states and five countries; less than 25% of the attendees are involved in current PRP efforts.

The workshop objectives were to (a) engage potential NRP teams and collaborators and facilitate in-person interactions, (b) gain a common understanding of the science engagement process, (c) explore and capture data storage, computing, and networking requirements across scientific domains, (d) identify common science-driven technical requirements for the NRP, and (e) exchange technical ideas for scaling the NRP. A number of findings were identified from the workshop and are described in detail in the workshop report [3]. Due to space limitations, we include here only the workshop recommendations:

- The NSF should not pursue just incremental changes over the next 3-5 years, but rather transformative changes. U.S. campuses' cyberinfrastructure capabilities have advanced greatly under NSF's Campus Cyberinfrastructure (CC*) programs, laying the foundation for an emerging NRP.
- Specifically, we recommend that NSF:
  - Continue funding its Campus Cyberinfrastructure program, including more opportunities for campus cyber-engineers and cyberteams.
  - Issue a call for proposals to address the technical and sociopolitical issues that will need to be addressed in an NRP; this could be an excellent opportunity for DOE and NSF to work together on a joint program.
  - Encourage science engagement facilitators to collaborate across campuses to build the NRP.
  - Support more regional Science DMZs to be formed from existing/new CC* campus grants. Multi-campus networking organizations should be encouraged to take the initiative to create regional DMZ proposals. Parallelism rather than central control is key to scaling.

6 PATH FORWARD

The PRP grant is mid-way through its five-year term. We have completed PRP Version 1, with the primary objective to achieve high-speed connectivity across our current partners, and are moving ahead with plans for Version 2 (V2). The V2 technical highlights include transitioning to IPv6 where feasible, experimenting with SDN/SDX technologies, and developing additional tools and procedures (e.g., the K8s, Rook/Ceph, Jupyter work). In parallel with these technical developments, we

will continue to engage researchers and add use cases within the PRP, work with the non-CENIC PRP regional network partners (e.g., PNWGP, FRGP, and MREN) to extend capabilities to their institutions, and continue outreach efforts and collaborations. A new collaboration with the Cognitive Hardware and Software Ecosystem Community Infrastructure (CHASE-CI) machine learning community [39] actively uses PRP. We are working with HPWREN on their extension to Orange County and beyond. We are collaborating with The Quilt [40] and Internet2 to engage additional regional networks. The second annual National Research Platform workshop (2NRP) in August 2018 is expected to have broad national/international attendance and will focus on common hardware/software platforms, use cases (successful or not), scaling the monitoring infrastructure, security issues, and outreach to campus decision-makers. Prior to the 2NRP, there will be a 4-day training workshop specifically aimed at MSIs, HBCUs, and EPSCoR universities on Science DMZs, DTNs and FIONAs, measurement tools, and Kubernetes.

NSF has recently demonstrated its interest in establishing high performance end-to-end capability at a national scale with a new section in their Campus Cyberinfrastructure (CC*) solicitation 18-508 to "establish a national entity of expertise and resources in improving end-to-end network performance across the NSF research and education communities." We look forward to NSF's continued investments toward a National Research Platform.

ACKNOWLEDGMENTS

The PRP is supported by the NSF CC*DNI DIBBS program under Cooperative Agreement ACI-1541349; additional NSF workshop support was provided under award #1540112. Seed funding and ongoing support have been provided by the University of California Office of the President, the Corporation for Education Network Initiatives in California (CENIC), and the Qualcomm Institute (QI), the UCSD Division of the California Institute for Telecommunications and Information Technology (Calit2). This project is only possible with substantial "volunteer" efforts amongst its partners (see Figure 2). We particularly appreciate the support by CENIC and ESNet, including the invaluable contributions of John Hess and Eli Dart. The project team includes Dan Gillette (UC Berkeley (UCB/Center for Information Technology Research in the Interest of Society (CITRIS)), Chris Hoffman (UCB), Tom Hutton (UCSD/SDSC), Brandie Nonnecke (UCB/CITRIS), Valerie Polichar (UCSD/IT), Patrick Schmitz (UCB), Molly Wofford (UCSD/QI), Mark Yashar (UCB/CITRIS), and Andrea Zonca (UCSD/SDSC).

REFERENCES

[1] NSF Campus Cyberinfrastructure program, including solicitations 18-508, 16-567, 15-534, 14-521 and predecessors. Retrieved from https://www.nsf.gov/funding/pgm_summ.jsp?pims_id=504748.
[2] ESNet. Science DMZ: A Scalable Network Design Model for Optimizing Science Data Transfers. Retrieved from https://fasterdata.es.net/science-dmz/.
[3] The First National Research Platform Workshop: Toward a National Big Data Superhighway, August 7-8, 2017, Bozeman MT. Retrieved from http://prp.ucsd.edu/workshop-reports/NRP%20Workshop%20Report_FINAL-23-Oct-17.pdf/view.
[4] ESNet. Science Engagement. Retrieved from https://www.es.net/ science-engagement/.
[5] ESNet. CrossConnects Workshop Series. Retrieved from https://www.es.net/science-engagement/programs-and-workshops/crossconnects-workshop-series/.
[6] Building the Pacific Research Platform: A Workshop Towards Deploying a Science-Driven Regional 'Big Data Freeway', October 14-16, 2015, La Jolla, CA. Retrieved from http://prp.ucsd.edu/workshop-reports/2015-prp-workshop-report/view.
[7] Pacific Research Platform Version 2 (PRPv2) Workshop: Summary of Workshop Presentation, February 21, 2017, La Jolla, CA. Retrieved from http://prp.ucsd.edu/workshop-reports/prp-v2-workshop-report-2-21-17-final.pdf/view
[8] Big Data and the Earth Sciences: Grand Challenges Workshop, May 31-June 2, 2017, La Jolla, CA. Retrieved from http://prp.ucsd.edu/workshop-reports/BigDataWorkshop2017_Report_FINAL_082417.pdf/view.
[9] Larry Smarr, Andrew A. Chien, Tom DeFanti, Jason Leigh, Philip M. Papadopoulos, 2003. The OptIPuter. Comm. of the ACM, 46, 58-67.
[10] Phil Papadopoulos, B. Dunne, Larry Smarr, J. Ford, and S. Fainman. Quartzite: A Campus-Scale Hybrid Networking Infrastructure. Optical Fiber Communication/ National Fiber Optic Engineers Conference 2008, 24-28 Feb. 2008, pp 1–3.
[11] Viewing Research Bandwidth through a New Prism. Retrieved from ucsdnews.ucsd.edu/pressrelease/ viewing_research_through_a_new_prism.
[12] CENIC Network. Retrieved from https://cenic.org/network/network-overview.
[13] ESNet. Science DMZ: Data Transfer Nodes. Retrieved from https://fasterdata.es.net/science-dmz/DTN/.
[14] UCSD PRP. Summer 2017 FIONA build. Retrieved from http://prp.ucsd.edu/solution-resources.
[15] Brian Tierney, Ezra Kissel, Martin Swany, and Eric Pouyol. 2012. Efficient Data Transfer Protocols for Big Data. IEEE 8th International Conference on e-Science, 8-12 October, 2012, DOI 10.1109/eScience.2012.6404462. Retrieved from https://www.es.net/assets/pubs_presos/ eScience-networks.pdf.
[16] Se-Yeung Yu, Nevil Brownlee, Aniket Mahanti. 2013. Comparative performance analysis of high-speed transfer protocols for big data. IEEE 38th Conference on Local Computer Networks, 21-24 October 2013, DOI 10.1109/LCN.2013.6761252. Retrieved from http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=6761252.
[17] ESnet. Data Transfer Tools. Retrieved from https://fasterdata.es.net/data-transfer-tools/.
[18] Globus. Data Transfer with Globus. Retrieved from https://www.globus.org/data-transfer.
[19] CERN. Fast Data Transfer – FDT. Retrieved from monalisa.cern.ch/FDT/.
[20] KeyCDN. UDP File Transfer – Explained. Retrieved from https://www.keycdn.com/support/udp-file-transfer/.
[21] UDT: UDP-based Data Transfer. Retrieved from http://udt.sourceforge.net/.
[22] ESnet. MaDDash: Monitoring and Debugging Dashboard. Retrieved from software.es.net/maddash/.
[23] Van Jacobson, Internet Hall of Fame Innovator. Retrieved from internethalloffame.org/inductees/van-jacobson.
[24] CENIC. CENIC perfSONAR Dashboard. Retrieved from https://ps-dashboard.cenic.net/.
[25] Kubernetes: Production-Grade Container Orchestration. Retrieved from https://kubernetes.io.
[26] JupyterHub. Retrieved from http://jupyterhub.readthedocs.io/en/latest/.
[27] UCSD PRP. To see the latest PRP traceroute visualization, visit the Kubernetes traceroute tool page at https://traceroute.nautilus.optiputer.net/.
[28] CIO Magazine. What are containers and why do you need them? Retrieved from www.cio.com/article/2924995/software/what-are-containers-and-why-do-you-need-them.html.
[29] The Linux Foundation. Cloud Native Computing Foundation. Retrieved from www.cncf.io/.
[30] Techcrunch, Retrieved from techcrunch.com/2017/09/20/kubernetes-gains-momentum-as-big-name-vendors-join-cloud-native-computing-foundation/.
[31] Rook. File, Block and Object Storage Services for your Cloud-Native Environments. Retrieved from https://rook.io/.
[32] Wikipedia. Ceph (software). Retrieved from https://en.wikipedia.org/wiki/Ceph_(software).
[33] Bro. The Bro Network Security Monitor. Retrieved from www.bro.org.
[34] Trusted Computing Group. Trusted Platform Module (TPM). Retrieved from https://trustedcomputinggroup.org/ trusted-platform-module-tpm-summary/.
[35] Github. CoreOS/Kubernetes-dtc-admission. Retrieved from https://github.com/coreos/kubernetes-dtc-admission/.
[36] William Poehlman, Alex Feltus, Stephen Ficklin, 2017. Large-Scale Genomics Experiments Enabled by OSG Connect Resources. 2017 OSG All-Hands Meeting, 6-9 March 2017, La Jolla, CA. Retrieved from https://indico.fnal.gov/event/12973/session/28/contribution/62.
[37] High Performance Wireless Research & Education Network. Retrieved from hpwren.ucsd.edu.
[38] Trisha Gorman 2018. Science Engagement: Best Practices and Scalability in the 21st Century. Retrieved from http://prp.ucsd.edu/Science%20Engagement-PRP-NRP_Report-FINAL%203-20-2018b.pdf .
[39] CHASE-CI NSF award 1730158, retrieved from https://nsf.gov/awardsearch/showAward?AWD_ID=1730158 .
[40] The Quilt Advanced Regional Networking, retrieved from https://www.thequilt.net.