

Beyond Campus Bridging: A retrospective of Cyberinfrastructure Integration Efforts



J. Eric Coulter, Indiana University; Rich Knepper, Resa Reynolds & Jodie Sprouse, Cornell University; Stephen Bird, Indiana University

xsede.org/ecosystem/ci-integration

Introduction

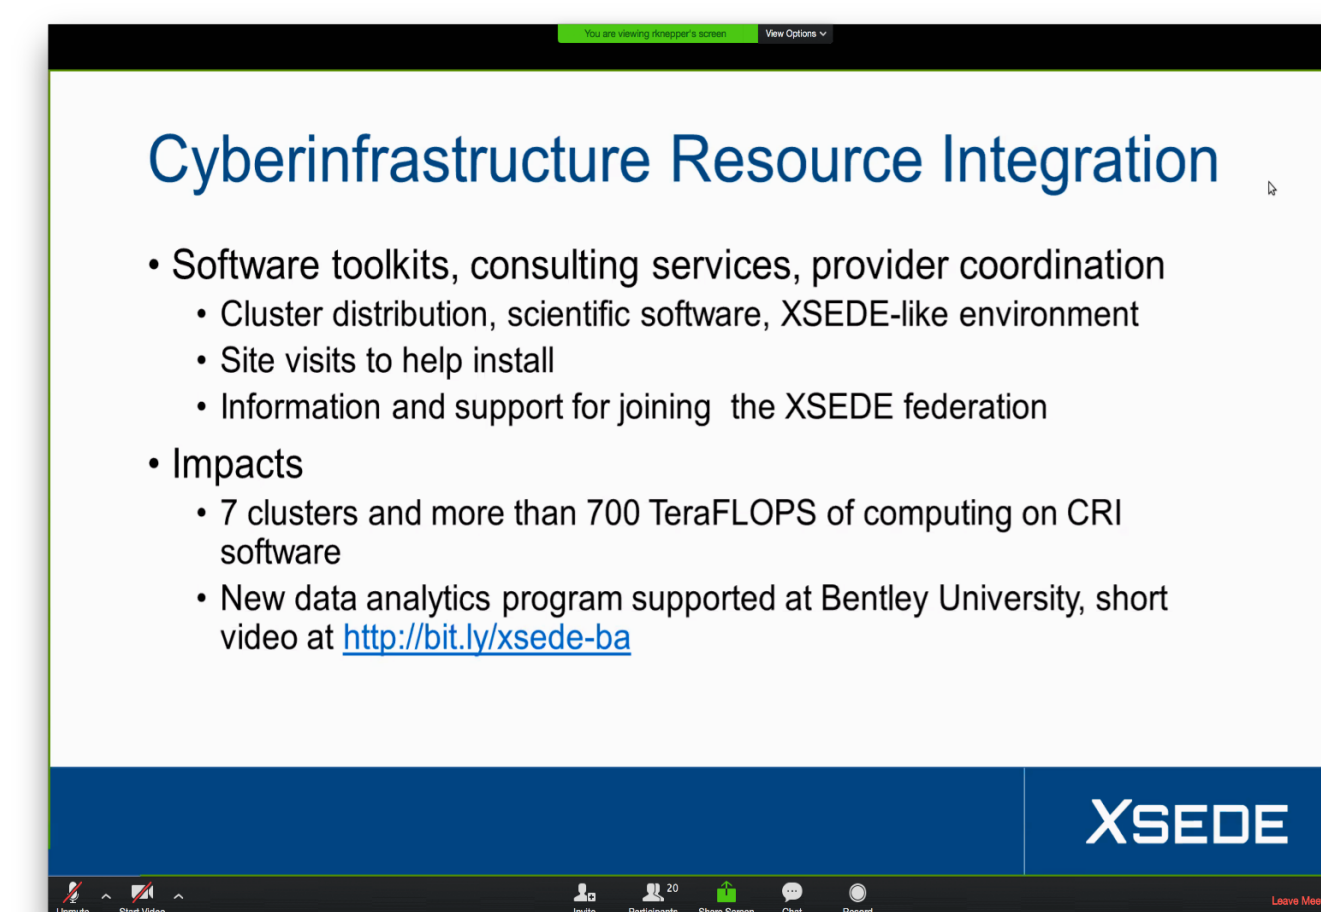
- XSEDE Cyberinfrastructure Resource Integration (XCRI) was originally known as “Campus Bridging,” and focused on facilitating the use of national supercomputing center resources at scale.
- Today, the XCRI team is extending national cyberinfrastructure through the XSEDE Service Provider Forum, campuses, and regional partnerships to maximize support for research initiatives.
- XCRI develops and implements toolkits so system admins can create local environments with national supercomputer features. This makes it easier for local users to transition to national resources.
- System admins also add software locally that makes it easier to allocate resources, manage systems, and produce metrics.
- Community feedback helps XCRI make improvements to its toolkits and processes.

Outreach

- Outreach is essential in explaining the types of services available.



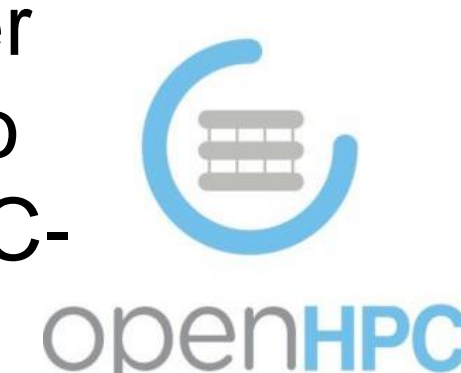
YouTube XSEDE Helps Small Universities Connect with the Nation's CI



- XSEDE Campus Bridging and XCRI has made dozens of presentations, presented 5 tutorials, and made over 20 papers and posters available to the user community.
- XCRI accomplished broader Cyberinfrastructure community participation via the Campus Champions group and OpenHPC Technical Steering Committee, and generated substantial “return on investment” in the form of new collaborations, site visit partners, and a growing group of XCRI adopters.

Toolkits & Services

XSEDE-Compatible Basic Cluster (XCBC) – primary toolkit whose implementation typically involves a site visit from an XCRI field engineer. Easy entry point for system administrators either unfamiliar with HPC systems, or without adequate time to easily implement one. OpenHPC component provides HPC-oriented software on top of basic OS installation.



Toolkits & Services

XSEDE-National Integration Toolkit (XNIT) – standard YUM repository of commonly used scientific applications to ease the burden of ongoing system administration. Easily installed with little assistance from XCRI team members. ~100 site subscribers to date, but XCRI is currently transitioning to an easier to maintain tool for easing the pain of installing scientific software.

Cluster Monitoring Toolkit – typically used alongside the XCBC. Enables easy setup and configuration of common tracking tools for HPC systems. Allows an HPC administrator to quickly install and configure Open XDMoD and Ganglia. Open XDMoD tracks jobs at a very granular level and creates reports on cluster usage to justify and procure funding. Ganglia monitors the current state of the entire system.



XCRI Virtual Cluster Toolkit – allows users to create bespoke HPC-style systems on the Jetstream cloud for experimental, educational, or research purposes. A series of scripts enable one to quickly create an elastic HPC environment on Jetstream, creating efficient use of allocations, and freeing users from the burden of creating and destroying virtual machines manually.



Site visits



Site	Nodes	Users	TFLOPS
Brandeis University	145	100	200*
South Dakota State University	70	10-100	200*
Slippery Rock University	10	10	10
Doane University	33	6	10
University of Cincinnati	36	101	10-100

XSEDE-Compatible Basic Cluster (XCBC) implementation examples (*includes GPUs)

Site Visits – XCRI has embarked on 6 site visits in which OpenHPC XCBC resources were implemented, separate from Campus Bridging years which are well documented in Coulter et al. Simple XSEDE-like Clusters: Science Enabled and Lessons Learned, *XSEDE16*.

Engagements with Research Computing Organizations – XCRI also provides health checks and consulting.

Organization & Metrics

- XCRI maintains these efforts with fewer than 3.5 FTEs spread across 6 individuals. Experienced professionals at funded CI providers provide deep advice on specific technical topics.
- Key metric: number of users supported by partner institutions (rather than total TFLOPs). Also tracked: shorter consultations without implementing local XCBC resources, number of new toolkits released, number of significant updates to software toolkits, and “User Satisfaction” with workshops, tutorials, and site visits.

Maintenance

XCRI migrates to new software when there are obvious signs of a community shift, such as using the Slurm scheduler rather than SGE, or moving to OpenHPC vs. Rocks. Maintenance of existing toolkits is tracked, and the team continuously integrates updates to the underlying software into existing toolkits. We also respond to new needs by updating toolkits based on lessons learned at new sites.

Rolling Out New Tools

- XCRI implemented best practices for new toolkit creation based on those used by the XSEDE Requirements and Analysis Capability Delivery. Community feedback is solicited, when possible.
- Making the installation and configuration of tools easy often requires DevOps skills.
- Uninitiated campus resource providers require regular outreach to identify needs and grow.
- Situating XCRI staff at a variety of home institutions is critical for deep access to institutional CI provider knowledge.



XSEDE and the Aristotle Cloud Federation announced a partnership to implement OpenStack clouds on U.S. campuses. This XCRI service begins in Sept. 2019 and will use Aristotle team documentation.

Partnerships

- XCRI formed a strong coalition with the Science Gateway Research Centers, which provides easy CI access to a huge variety of science domain projects.
- Ongoing partner interactions highlight areas of greatest need (e.g., 4 sites requested aid setting up GPUs after initial site visits).
- XCRI sites are very helpful spreading the word about XCRI services. Several have become XSEDE Level 3 service providers.

Conclusion

- XCRI services and toolkits have successfully reduced barriers to entry to national research computing. Lessons learned from our efforts will strengthen future programs, in whatever form those take.
- To learn more, read the companion paper to this poster: J. Eric Coulter, Rich Knepper, et al. Beyond Campus Bridging: a retrospective of Cyberinfrastructure Integration Efforts, *PEARC '19*.