

**NEREN Seminar:  
“Bridging the Gap: Sharing Computing Resources  
Across Campuses”**

**Sponsored by Intel, Red River, Dell  
Mainline Information Systems and  
Partners Juniper and Cohesity**

**Friday, April 5, 2019, 9:30AM– 3:00PM  
Complimentary breakfast begins at 9:00 a.m.  
(Registration/Breakfast – 9:00AM – 9:45AM)  
Complimentary lunch begins at 12 noon  
Providence Marriott, 1 Orms Street, Providence, RI 02904**

**NEREN Seminar –Friday, April 5 2019, 9:00AM-3:00PM**

In collaboration with UMass Amherst and the Massachusetts Green High-Performance Computing Center (MGHPCC), NEREN presents the sixth in a series of day-long seminars devoted to proposing and advancing ideas for regional collaboration in research computing and networking. By attending in person, you will gain the added benefit of interpersonal networking with peers from across the Northeast region.

Please register by March 25, 2019, to attend in person or remotely by contacting Laurie Robinson, NEREN Program Administrator, at [laurie@neren.org](mailto:laurie@neren.org) or by phone: 401-523-5107. When registering, please indicate whether you are attending in person or remotely. Also, please contact Laurie if you have questions. Or register online:

[https://www.eventbrite.com/e/neren-seminar-bridging-the-gapsharing-computing-resources-across-campuses-tickets-55048617898?aff=utm\\_source%3Dweb\\_email%26utm\\_medium%3Demail%26utm\\_campaign%3Dnew\\_event\\_email&utm\\_term=eventurl\\_text](https://www.eventbrite.com/e/neren-seminar-bridging-the-gapsharing-computing-resources-across-campuses-tickets-55048617898?aff=utm_source%3Dweb_email%26utm_medium%3Demail%26utm_campaign%3Dnew_event_email&utm_term=eventurl_text)

## AGENDA – Friday, April 5, 2019

9:00 – 9:45 a.m. Registration/Continental Breakfast/Networking Sponsored by Intel and Red River.

9:45 – 10:00 a.m. Welcoming/Opening Remarks by NEREN, Inc.

**Please note that the Webcast/Phone options will be available for the presentations only and not the Working session in the afternoon. If joining by webcast, please use the following link:**

<http://demo.mediasite.oshean.org/Mediasite/Play/17067cfd1d8f48118196e35a615e4ba91d>

10:00 – 10:40 a.m. Presentation #1 – “*Eastern Regional Network (ERN) Update*”

**Presenter: Dr. James Barr von Oehsen, Associate Vice President of the Office of Advanced Research Computing (OARC) at Rutgers, The State University of New Jersey**

The goal of the Eastern Research Network (ERN) is to simplify multi-campus collaborations and partnerships that advance the frontiers of research and innovation. Through a partnership of educational institutions, research facilities, regional network providers, and Internet2, the ERN is committed to providing layered and transparent access to shared data and computing facilities. The talk will discuss the research and education needs that motivate a regional research platform, along with the engineering and community-building challenges that must be addressed in order to build it. The talk will also review progress on a resource federation proof of concept that was developed over the past year.

10:40 – 11:20 a.m. Presentation #2 – “*The GENI Project*”

**Presenter: Abraham Matta, Ph.D., Professor and Chair of Computer Science, Boston University**

This talk introduces GENI, the Global Environment for Networking Innovation. GENI is a virtual laboratory that spans compute resources distributed over the United States. It enables research and education in networking and distributed systems, with applications in domain science. The main components of GENI will be introduced, along with use cases to demonstrate its deep programmability for supporting a wide range of research and education efforts.

11:20 – 12:00 p.m. Presentation #3 – “*Building Community CI on the Pacific Research Platform: Looking Towards The National Research Platform*”

**Presenter by Video: Thomas A. DeFanti, Ph.D., Research Scientist at UCSD’s Calit2/Qualcomm Institute and Professor Emeritus of Computer Science, University of Illinois**

The Pacific Research Platform is a NSF project to connect ScienceDMZs in 25 universities and national labs. The first roll out, PRPv1, provided purpose-built ‘FIONA’ rack-mounted PCs that were data transfer nodes (DTNs) tuned to exploit end-to-end 1G, 10G, 40G and 100G connections to CENIC’s 100G HPR and national/international connections. Last year, PRPv2 was launched. It is adding 160-192TB of storage to many FIONAs in campus DMZs. PRPv2 uses Google’s open-source Kubernetes to manage Nautilus, our 14-campus community-based expandable CPU/GPU/Storage/Visualization cluster (as of 3.1.19, 2792 CPU Cores on 75 Hosts with 1.7 PB Storage and 300 GPUs that offer >25M core/hrs/day). Another NSF grant, CHASE-CI, is adding FIONAs with a total so far of 300 GPUs into Nautilus. High-availability computing and secure-storage FIONAs are also offered. Many researchers nationwide and worldwide are adding capability to Nautilus. Recently, *Towards The National Research Platform* was funded for three years by NSF to replicate and federate PRPv2/Nautilus nationwide via selected regional networks, offering advanced measurement, monitoring, IPv6 and security to resource sharers, as well as orchestrating a merger with The Open Science Grid.

12 noon – 12:30 p.m. Lunch Break

12:35 – 1:15 p.m. – Presentation #4 – “Dive into Apache Kafka®”

**Presenter: Paul Earsy, Senior Sales Engineer, Confluent**

In the last few years, Apache Kafka® has been used extensively in enterprises for real-time data collecting, delivering, and processing. In this talk, we'll dive into some of the key internals that help make Kafka popular including:

- Companies like LinkedIn sending greater than one trillion messages per day to Kafka. Learn about the underlying design in Kafka that leads to such high throughput.
- Many companies (e.g., financial institutions) are now storing mission critical data in Kafka. Learn how Kafka supports high availability and durability through its built-in replication mechanism.

Modern businesses have data at their core, and this data is changing continuously. How can we harness this torrent of continuously changing data in real-time? The answer is stream processing, and Apache Kafka® is a core hub for streaming data.. This talk will provide a brief introduction to Apache Kafka® and describe its usage as a platform for streaming data. It will explain how Kafka serves as a foundation for both streaming data pipelines and applications that consume and process real-time data streams. It will introduce some of the newer components of Kafka that help make this possible, including Kafka Connect, a framework for capturing continuous data streams, and Kafka Streams, a lightweight stream processing library. Confluent, founded by the creators of open source Apache Kafka®, provides the streaming platform that enables enterprises to maximize the value of data. Confluent Platform lets leaders in industries such as retail, logistics, manufacturing, financial services, education, government and technology and media, move data from isolated systems into a real-time data pipeline where they can act on it immediately. Backed by Benchmark, Data Collective, Index Ventures and LinkedIn, Confluent is based in Palo Alto, California.

1:20 – 2:20 p.m. Breakout Session – “ERN Planning Session”

2:20 – 3:00 p.m. Closing Remarks Christopher Misra, Interim Vice Chancellor IT and CIO, UMass Amherst

**To register, or if you have questions, please email Laurie Robinson, NEREN Program Administrator, at [laurie@neren.org](mailto:laurie@neren.org) or by phone: 401-523-5107. Or register online:**

**[https://www.eventbrite.com/e/neren-seminar-bridging-the-gapsharing-computing-resources-across-campus-tickets-](https://www.eventbrite.com/e/neren-seminar-bridging-the-gapsharing-computing-resources-across-campus-tickets-55048617898?aff=utm_source%3Dweb_email%26utm_medium%3Demail%26utm_campaign%3Dnew_event_email&utm_term=eventurl_text)**

**[55048617898?aff=utm\\_source%3Dweb\\_email%26utm\\_medium%3Demail%26utm\\_campaign%3Dnew\\_event\\_email&utm\\_term=eventurl\\_text](https://www.eventbrite.com/e/neren-seminar-bridging-the-gapsharing-computing-resources-across-campus-tickets-55048617898?aff=utm_source%3Dweb_email%26utm_medium%3Demail%26utm_campaign%3Dnew_event_email&utm_term=eventurl_text)**

**GETTING THERE -- DETAILS**

**Getting There in Person:**

The Providence Marriott is easily accessible by car, train or bus and is conveniently located minutes from Downtown Providence. To those traveling by train or bus, the Providence Marriott is within short distance from the train and bus stations in Providence, Rhode Island.

## **Helpful Telephone Numbers/Contact Information**

John Griffin, (UMass Amherst) 413-545-9939 or [jgriffin@umass.edu](mailto:jgriffin@umass.edu)

John Goodhue (MGHPCC) 617-834-5601 or [jgoodhue@mghpcc.org](mailto:jgoodhue@mghpcc.org)

Laurie Robinson, (NEREN) [401-523-5107](tel:401-523-5107), [laurie@neren.org](mailto:laurie@neren.org)

For Technical Questions: Jim Carr, Consultant, (OSHEAN), 401-447-5600, [jim@oshean.org](mailto:jim@oshean.org)

## **Meeting Location/Wireless Internet Access**

Providence Marriott does have a guest wireless network.

## **Parking/Public Transportation**

There is plenty of free parking available at the Providence Marriott

## **Getting There Remotely – Phone/Webcast**

<http://demo.mediasite.oshean.org/Mediasite/Play/17067cfd1d8f48118196e35a615e4ba91d>

**If you are planning to participate in the Webcast, please contact Laurie Robinson– [laurie@neren.org](mailto:laurie@neren.org) or register online:**

**[https://www.eventbrite.com/e/neren-seminar-bridging-the-gapsharing-computing-resources-across-campus-tickets-55048617898?aff=utm\\_source%3Dweb\\_email%26utm\\_medium%3Demail%26utm\\_campaign%3Dnew\\_event\\_email&utm\\_term=eventurl\\_textby](https://www.eventbrite.com/e/neren-seminar-bridging-the-gapsharing-computing-resources-across-campus-tickets-55048617898?aff=utm_source%3Dweb_email%26utm_medium%3Demail%26utm_campaign%3Dnew_event_email&utm_term=eventurl_textby)**

Please note that the Webcast will be recorded, and we will provide details following the seminar. Thank you to David Marble, President and CEO of OSHEAN for assisting with the Webcast/Video Streaming!

## **About our Presenters**

**Dr. James Barr von Oehsen**

**Associate Vice President of the**

**Office of Advanced Research Computing (OARC) at**

**Rutgers – The State University of New Jersey**

Dr. James Barr von Oehsen is the Associate Vice President of the Office of Advanced Research Computing (OARC) at Rutgers, The State University of New Jersey. He has worked as a computational scientist for over 16 years. As Associate Vice President of OARC, he is responsible for providing strategic leadership in advancing Rutgers University's research and scholarly achievements through next generation computing, networking, data technologies, and creative learning environments. Prior to joining Rutgers, he was employed by Clemson University Computing and Information Technology (CCIT) as the Executive Director of the Cyberinfrastructure Technology Integration (CITI) group. Before joining CITI, he was affiliated with the Center for Advanced Engineering Fibers and Films (CAEFF), an NSF funded engineering research center, where he directed the research and development of parallel finite element code for modeling polymer processes. While at CAEFF he was also responsible for building the advanced computing environment (including storage, web services, databases, network, and visualization). His current interests include high performance computing, high throughput computing, mathematical modeling, parallel programming, campus level distributed cloud environments, hardware architectures, cloud services, and next generation creative learning environments.

Dr. von Oehsen and the OARC team are supported by University partners and Faculty Advisory Committee.

## **About the Presenters (Continued)**

**Abraham Matta, PhD.**

**Professor and Chair of Computer Science, Boston University**

Abraham Matta received his Ph.D. in computer science from the University of Maryland at College Park in 1995. He works on the design of network protocols and architectures based on a range of computer science principles (such as inter-process communication, decomposition, and recursion), mathematical techniques (such as probabilistic analysis, queuing theory, optimization, and control theory), and performance evaluation tools (such as simulation and emulation). Application domains include the Internet, wireless, mobile, sensor, and disruption-tolerant networks, cloud and distributed systems. He has published over 100 peer-reviewed technical papers. He received the National Science Foundation CAREER award (1997). He won a patent (2011) and two best-paper awards (2008 and 2010) on his work on wireless ad hoc and sensor networks. He has served as chair or co-chair of many technical program committees, such as the IEEE Online Conference on Green Communications (2012), IEEE Computer Communications Workshop (2011), and International Conference on Network Protocols (2005). He has served on many organizing committees, including as general chair of the 4th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (2006). He is a senior member of the ACM and IEEE. He is a member of the Networks Research Group in the Computer Science Department at BU.

**Thomas A. DeFanti, Ph.D.**

**Research Scientist/UCSD's Calit2/Qualcom Institute**

**Professor Emeritus of Computer Science – University of Illinois at Chicago**

Thomas A. DeFanti, PhD, is a research scientist at UCSD's Calit2/Qualcomm Institute, and a distinguished professor emeritus of Computer Science at the University of Illinois at Chicago. He received the 1988 ACM Outstanding Contribution Award and became an ACM Fellow in 1994, and the Distinguished Alumni Award from Ohio State in 2018. He co-founded StarLight in Chicago with Joe Mambretti and Maxine Brown. He and Dan Sandin conceived the CAVE virtual reality theater in 1991. Since coming to UCSD in 2004, he and Greg Dawe built the StarCAVE, NexCAVE, TourCAVE, WAVE, SunCAVE, and CAVEkiosk/4KAVE VR systems. With Larry Smarr, Dima Mishin, John Graham, Phil Papadopoulos and teams of colleagues, Tom designs and builds 40/100Gbs state-of-the-art wide-area computer networked computing, storage, machine learning/graphics processing and visualization facilities to serve advanced scientific and artistic computing communities. He is co-PI of the NSF Pacific Research Platform project ([prp.ucsd.edu](http://prp.ucsd.edu), Larry Smarr, PI), of the CHASE-CI NSF community infrastructure project for machine learning, and is project manager for the NSF-funded Toward the National Research Platform project. He is profiled in Wikipedia, at <http://www.evl.uic.edu/tom/>, and [http://www.calit2.net/people/staff\\_detail.php?id=67](http://www.calit2.net/people/staff_detail.php?id=67)

## **About the Presenters (Continued)**


### **Paul Earsy Senior Sales Engineer at Confluent**

Paul is a Senior Solutions Engineer at Confluent. Previously he has worked at Splunk via the acquisition of Rocana. Paul has over 30 years working in technology, most recently focusing on distributed computing and security. Paul has a BA in Business from the University of Massachusetts at Amherst. Paul lives in the Boston area and enjoys spending time with his adult children. <https://www.linkedin.com/in/paulearsy/>

### **Christopher Misra Interim Vice Chancellor IT and CIO, UMass Amherst Welcoming/Breakout Sessions/Closing Remarks**

Christopher Misra is the Interim Vice Chancellor IT and CIO at the University of Massachusetts Amherst where he has worked for many years. His responsibilities include management of overall campus technology coordination, networking, data centers, information security program, identity management, and enterprise architecture. Chris has been active for many years with various regional and national information security organizations including the Security Task Force, Internet2 Salsa, and REN-ISAC, serving on program committees, chairing working groups, and presenting at conferences. Chris is also an instructor at UMass where he has taught undergraduate courses on Network Security for many years.


## About our Sponsors



Intel invents at the boundaries of technology to make amazing experiences possible for society. Harnessing the capability of the cloud, the ubiquity of the Internet of Things, the latest advances in AI, and the promise of always-on 5G connectivity, Intel's [scalable and secure solutions help](#) solve global challenges. Leading on policy, diversity, inclusion, education and sustainability, we create value for our customers and society. Find out more at [www.intel.com/government](http://www.intel.com/government)



Red River brings together the ideal combination of talent, partners and products to disrupt the status quo in technology and drive success for business and government in ways previously unattainable. Red River serves organizations well beyond traditional technology integration, bringing more than 20 years of experience and mission-critical expertise in security, networking, analytics, collaboration, mobility and cloud solutions. Red River delivers ongoing IT consulting, maintenance and support to ensure optimal success through an array of unique service offerings, including integrated Managed, Professional, Technical and Support Services. Whether you need help answering simple end-user inquiries with our 24x7x365 Network Operations Center, need to develop a path to the cloud, or implement new network security protocols, our certified experts are available to help. Learn more at [redriver.com](http://redriver.com)



**Listen. Learn. Deliver. That's what we're about.**

Dell empowers countries, communities, customers and people everywhere to use technology to realize their dreams. Customers trust us to deliver technology solutions that help them do and achieve more, whether they're at home, work, school or anywhere in their world. Learn more about our story, purpose and people behind our customer-centric approach. <https://www.dell.com/en-us>



Mainline represents over 120 Technology Companies along with an Information Management/Analytics Practice, Network Practice, Security Practice, Managed Services, and Staff Augmentation Services. We work with clients as a technology advisor to assist in selecting best fit technologies and solutions for their organizations. Steve Frodey, Account Executive, Mainline Information Systems, 518-441-4933, [steve.frodey@mainline.com](mailto:steve.frodey@mainline.com)



Juniper offers high-performance network infrastructure built on simplicity, security, openness, and scale. We are a market leader in routing, switching, and security for Higher Education Campus and Wide Area R&E Networking. JJ Jamison, CTO, Research & Education Markets, Juniper Networks, 703-347-1622, [jj@juniper.net](mailto:jj@juniper.net) and Allen Scalise, Account Executive, Juniper Networks, 508-683-9757, [ascalise@juniper.net](mailto:ascalise@juniper.net)

# COHESITY

Cohesity delivers a hyper-converged, cloud-native, massively scalable storage platform for secondary data. By consolidating all secondary data – including backups, files, objects, test/dev copies, and analytics data - and associated management functions on a unified solution, customers save as much as 70% when they modernize with Cohesity's Data Platform. Contact: Tim Snell, Account Executive, [tsnell@cohesity.com](mailto:tsnell@cohesity.com), (781) 258-1867.