Meet HPC Innovator
Chin Guok

Helping to run the world's fastest science network wasn't originally on Guok's career horizon

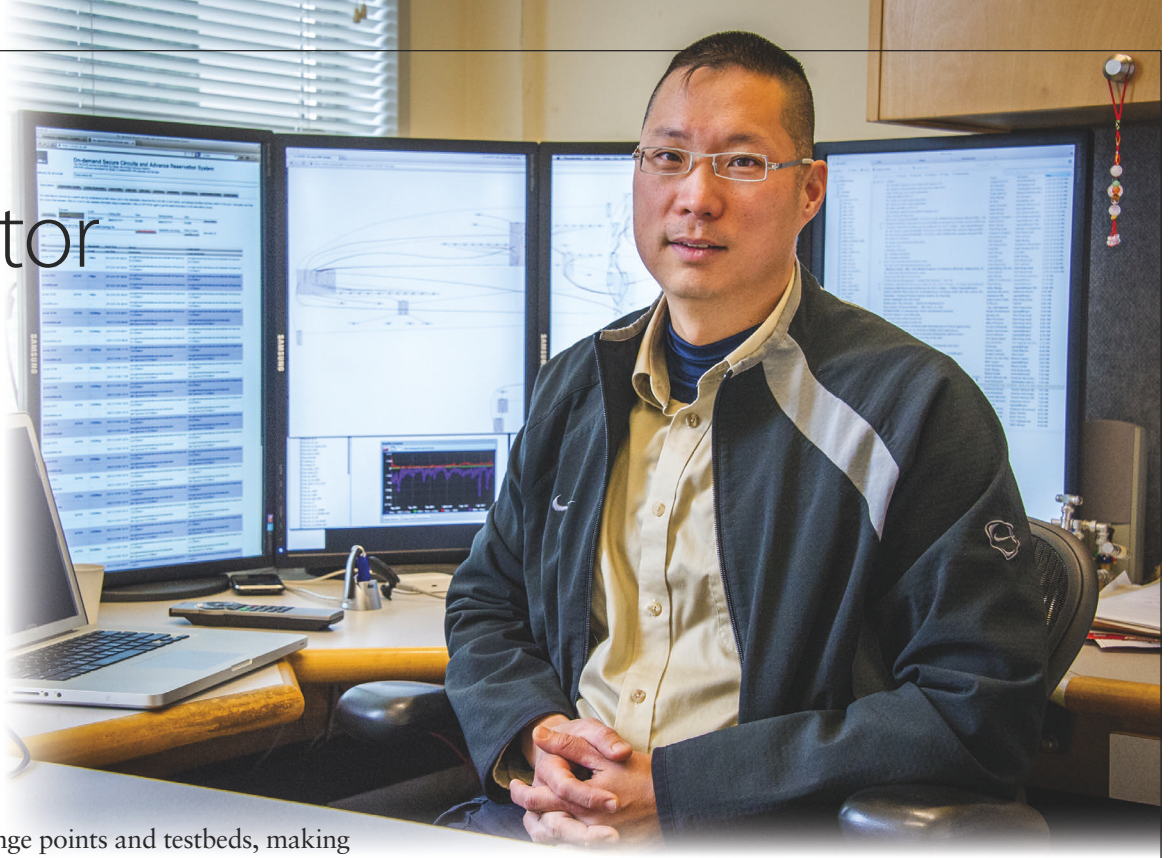

Photo: Roy Kaltschmidt, LBNL

Jon Bashor

Although the subject of software defined networking, or SDN, is now a hot topic in the networking community, Chin Guok of the Department of Energy's ESnet (the Energy Sciences Network) has been helping lay the foundation for SDN since 2004. Guok is the technical lead for OSCARS, ESnet's On-Demand Secure Circuits and Advance Reservation System.

OSCARS allows users to quickly configure multi-domain, high-bandwidth virtual circuits that guarantee end-to-end network data transfer performance. Although the networking industry has fielded competing SDN technologies as companies seek leadership in the field, ESnet has been developing and deploying OSCARS since soon after it was first proposed as a research project in August 2004.

OSCARS has grown into a robust production service used by more than 40 other networks, exchange points and testbeds, making it the most widely adopted inter-domain dynamic circuit services application within the global research and networking community.

Broadly defined, SDN makes it easier for software applications to automatically configure and control the various layers of the network. ESnet has been conducting localized tests to innovate, experiment and demonstrate the value of SDN when applied toward end-to-end support of data-intensive science collaborations and applications.

SDN gives users a measure of predictability by giving them more control over their dataflows. Without this ability, the larger science data flows could compete unpredictably with other data transfers for bandwidth making it difficult for scientists to get critical data when it is needed, especially if the data is stored across multiple sites, which is increasingly common.

"When we proposed OSCARS, it was intended to be just within our network and not communicate with any outside networks. We were only operating at Layer 3 of the network providing IP circuits," Guok said. "But, as we worked with more collaborators, we also gained new requirements. We are now running it on Layer 2 as Ether-

net VLAN circuits and are extending the capabilities across multiple domains. The focus is now how to grow this service out beyond the 40 networks as well as up to the software applications."

That extensibility is important, as 80 percent of the traffic currently carried by ESnet has only one end in the ESnet domain and the other end somewhere else. Not only is OSCARS getting bigger, it's also getting faster and easier.

"For SC10, we ran the numbers on what it took to build a virtual circuit to Europe — it required 10 hours of phone calls — about 100 e-mails over three months," Guok said. "Now, with OSCARS, it takes one person about five minutes and one e-mail."

In order to work, OSCARS has to be "technology agnostic" and work independent of the underlying transport resources available. But to succeed, it has to overcome other hurdles, such as the fact that some nations' networks have a pay-per-use model, while DOE researchers are not charged for ESnet. The result has been a number of international agreements working out trades-in-kind.

DOWN TO THE NEXT LEVEL

In November 2012, ESnet partnered with network equipment vendor Infinera on a test to see if OSCARS could be used at Layer 1, the optical transport upon which the higher network layers are built. Using ESnet's testbed in New York, they successfully moved data between Manhattan and Brookhaven National Laboratory on Long Island. What the demo showed is that Infinera's optical transport switch could be dynamically configured by software at that level. This would allow data transfers to reroute themselves based on network demand, finding unused bandwidth and using the resource more efficiently.

The demonstration marked the first time an open architecture with SDN was used to provide traffic-engineered paths at the optical layer and was accomplished through extensions to the OpenFlow protocol. The open source OpenFlow application was developed to work on Layer 2. However, in this demo, it was configured to run on Layer 1 as well.

"Conceptually, this was a very big deal, as it makes the management of network devices much easier," Guok said. "When you allow a Layer 1 device to speak the same 'language' as the rest of the network, it's a very powerful tool for controlling and managing the network."

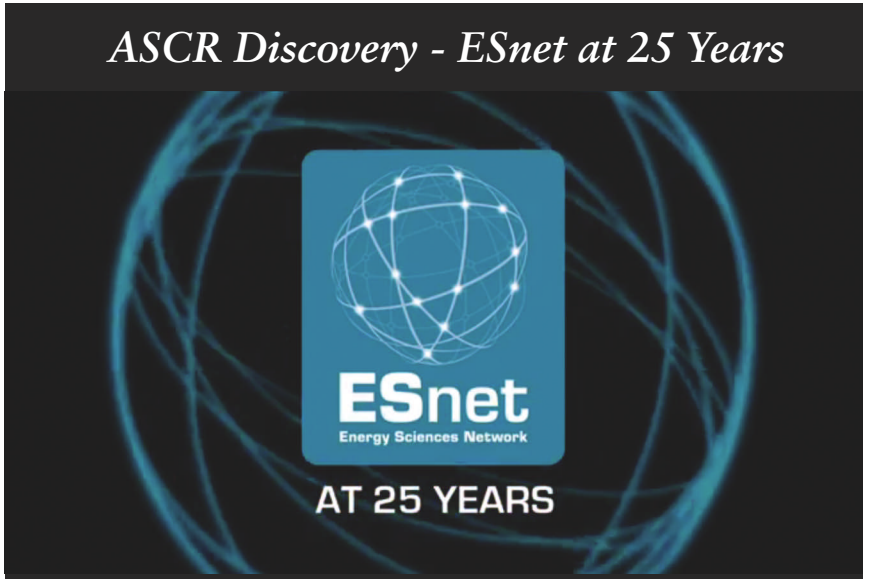

Computational science administrators and users reflect on 25 years of ESnet, the data network that ties 25,000 scientists to Department of Energy laboratories, computers and instruments. During its anniversary year, ESnet celebrated an upgrade to carry even more data even faster.

There is also an economic factor at play. At each higher layer, the equipment is more expensive as features and capabilities are added. For example, a 100 Gb interface for the optical layer may cost about $50,000, while a 100 Gbps-capable router for Layer 3 would cost about $250,000. This means the cost of sending a bit across the network goes up by 1.5 to 5 times as the traffic moves up to the next layer.

"You really want to get down to the lowest layer you need," Guok said. "One of the prime features of OSCARS is its 'intelligence' — it doesn't matter

OSCARS Timeline

Aug 2004: OSCARS proposal submitted.

Sep 2004: Initial network protocols testing completed (for Juniper M-series platform), first Layer-3 VC manually configured between BNL and FNAL, policed at 50Mb/s.

Feb 2005: Started collaboration with I2's BRUW project.

Apr 2005: First production use of OSCARS VC. Transatlantic LHC 10GE connection between CERN and Chicago severed by fishing boat, causing LHC Service Challenge data to reroute through NY and flow over FNAL production OC12. OSCARS VC set up to carry LHC Service Challenge traffic from NY to Chicago onto FNAL non-production connection at Starlight.

Jun 2005: End-user beta testing with Les Cottrell (SLAC), Sean Flanagan (GA), and Dantong Yu (BNL); jitter measurement testing with SLAC's Datagrid Wide Area Monitoring Infrastructure (DWMI).

Jan 2006: More jitter measurement testing with SLAC.

Mar 2006: Adoption of subset of GÉANT AMPS WSDL service descriptions; Formation of DICE (Dante, Internet2, Canarie, ESnet) Control Plane WG; Venue UCLPv2 meeting in Ontario.

Apr 2006: First Layer-3 interdomain VC dynamically negotiated between I2 (BRUW) and ESnet (OSCARS). Unidirectional 25Mb/s VC from I2 test host in Indianapolis, IN, to ESnet host in Sunnyvale, CA.

Jul 2006: Start of rewrite of OSCARS/BRUW software into Java.

Aug 2006: Successful Layer-3 interdomain reservation between BNL TeraPaths and ESnet OSCARS.

May 2007: Adoption of OGF NMWG topology schema in consensus with DICE Control Plane WG; Collaborative measurements of Hybrid Multilayer Network Data Plane testing with Internet2, USN, and ESnet (using OSCARS VCs)

Aug 2007: First ESnet Layer-2 VC (Ethernet VLAN) configured by OSCARS

Sep 2007: First interdomain topology exchange between GÉANT2, Internet2 and ESnet

Oct 2007: First Layer-2 interdomain VC dynamically negotiated between I2 (HOPI) and ESnet (OSCARS)

Nov 2007: Successful Layer-2 interdomain reservation between BNL TeraPaths and ESnet OSCARS ; FNAL LambdaStation and ESnet OSCARS; GÉANT2 AutoBAHN and ESnet OSCARS; Nortel DRAC and ESnet OSCARS; demonstrated token based authorization concept with OSCARS VC setup.

Feb 2008: Formation of the GLIF GNI-API Task Force of which OSCARS is one of the core representatives.

Apr 2008: Active collaboration with OGF NML-WG to combine work of NMWG and NDL.

May 2008: OSCARS VC operational change to remark over-subscribed packets to scavenger service instead of discarding; OSCARS operational change to support site coordinator role; DICE IDCP v1.0 specification completed.

Dec 2008: Successful control plane interdomain interoperability between ESnet OSCARS and g-Lambda using GLIF GNI-API GUSI.

Jan 2009: Formation of the GLIF NSI-WG. OSCARS is a core contributor in the writing of the OGF NSI-Architecture document; Draft architecture designs for OSCARS v0.6.

Oct 2009: Successful control plane interdomain interoperability between IDC (ESnet OSCARS, g-Lambda and Harmony using GLIF GNI-API Fenius (GUSI (GLIF Unified Service Interface))).

Nov 2009: Successful control and data plane interdomain interoperability between IDC (ESnet OSCARS, g-Lambda and Harmony using GLIF GNI-API Fenius); OSCARS used by SC09 SCinet to manage bandwidth challenges.

Feb 2010: DICE IDCP v1.1 released to support brokered notification.

Aug 2010: OSCARS is selected for NSF-funded Dynamic Network System (DYNES) project.

Nov 2010: OSCARS used by SC10 SCinet to manage demo and bandwidth challenges.

Jan 2011: OSCARS v0.6 PCE SDK released allowing developers to test OSCARS v0.6 flexible PCE framework.

Sep 2011: OSCARS interoperates with OGF NSI protocol v1 using an adapter at NSI Plugfest at GLIF in Rio.

Nov 2011: PSS for ALUs is developed and tested to support ESnet ANI 100G prototype network and SC11 SCinet.

Dec 2011: OSCARS v0.6 RC1 is released.

Jun 2012: OSCARS v0.6 Final is released.

Nov 2012: Successful demonstration of interdomain control-plane signalling using OGF NSI CS v2.0 (between OSCARS, g-Lambda, OpenDRAC, BoD, and AutoBHAN) at SC12.

if it's on Layer 1, 2 or 3."

Helping to run the world's fastest science network — ESnet just upgraded its national network to 100 Gbps — wasn't originally on Guok's career horizon. Growing up in Singapore, he was unsure if he could master Mandarin well enough to pass the compulsory exam to enter the local university. So, he came to the United States and earned his undergraduate degree in computer science from the University of the Pacific in California.

While there, he met Joe Burrescia, who was working at ESnet. After graduating, Guok returned to Singapore to fulfill his military service obligation. He subsequently returned to the United States, and was working as a teaching assistant at the University of Arizona while earning his master's in computer science when Burrescia found him via a Web search, asked "Are you the Chin who attended UOP?" and then encouraged him to apply for a network engineer job at ESnet.

"I was really thinking of going into operating systems, but Joe convinced me to give networking a try," Guok said. "My whole career at ESnet is based on a Web search, which I guess is fitting now that I'm a network engineer." HPC

Jon Bashor is a communications manager at Lawrence Berkeley National Laboratory. He may be reached at editor@ScientificComputing.com.