

Improving WAN Performance with XSP

GEC12 and SC11 SRS demos

GEC12 – Kansas City, MO

Ezra Kissel

University of Delaware

Martin Swany

Indiana University

November 3rd 2011



Overview

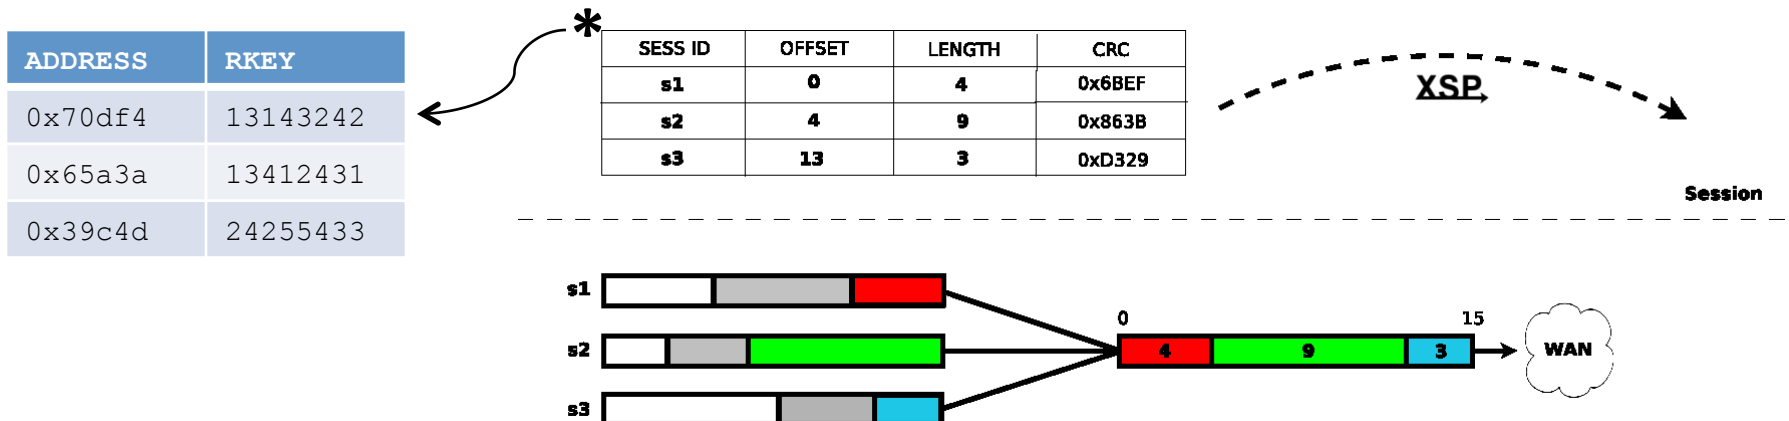
- GENI experiment
 - XSP, Phoebus and SLaBS
- GENI Instrumentation and Measurement project
 - LAMP
- Demos
 - Here, in just a few minutes
 - In two weeks at SC11
 - SC is the conference formerly known as Supercomputing
 - The SCinet Research Sandbox (SRS) will feature an OpenFlow network with switches from HP, IBM, NEC and Pronto (as well as some wide-area 100G experiments)
 - The Disruptive Technologies program will feature Mark Berman

eXtensible Session Protocol

- XSP is a session protocol for IP networks, and provides generalized messaging between gateways, network devices, and services/applications
 - Addresses the necessity for end to end context while allowing innovation in the core
 - Socket compatible shim library for legacy applications
- In-band and out-of-band signaling makes it ideal for dynamic or software defined networks
 - Dynamic path establishment
 - Provides a channel for application data and/or other metadata to be forwarded with session layer data frames
- Other benefits for future Internet
 - Binding E2E context to the session rather than the transport layer enables location / identification separation
 - Explicit signaling enables rich authentication and authorization support with SSH or X.509

Phoebus - SLaBS

- The Phoebus project aims to help bridge the performance gap in modern networks
 - Uses XSP to enable multiple adaptation and buffer points, “gateways”, in the network
 - Achieves 10G+ TCP forwarding performance on commodity hardware
 - WAN acceleration
- SLaBS applies burst switching concepts at Phoebus Gateways (PGs)
 - Use XSP to reframe, “slab”, incoming flows and send relatively large PDUs versus small layer-3, layer-4 PDUs common today, exchange metadata
- Minimize total transfer time of bulk data flows
 - Latency of individual segments is of less concern
 - Use alternative data channel to improve performance (e.g. RDMA, UDT)

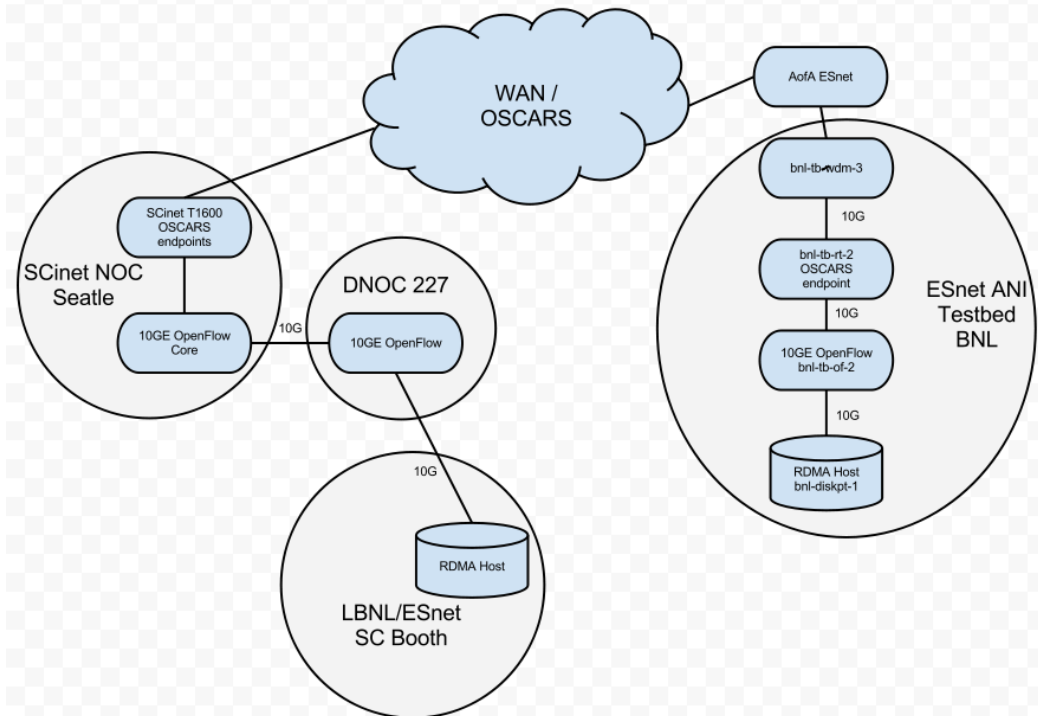


Periscope

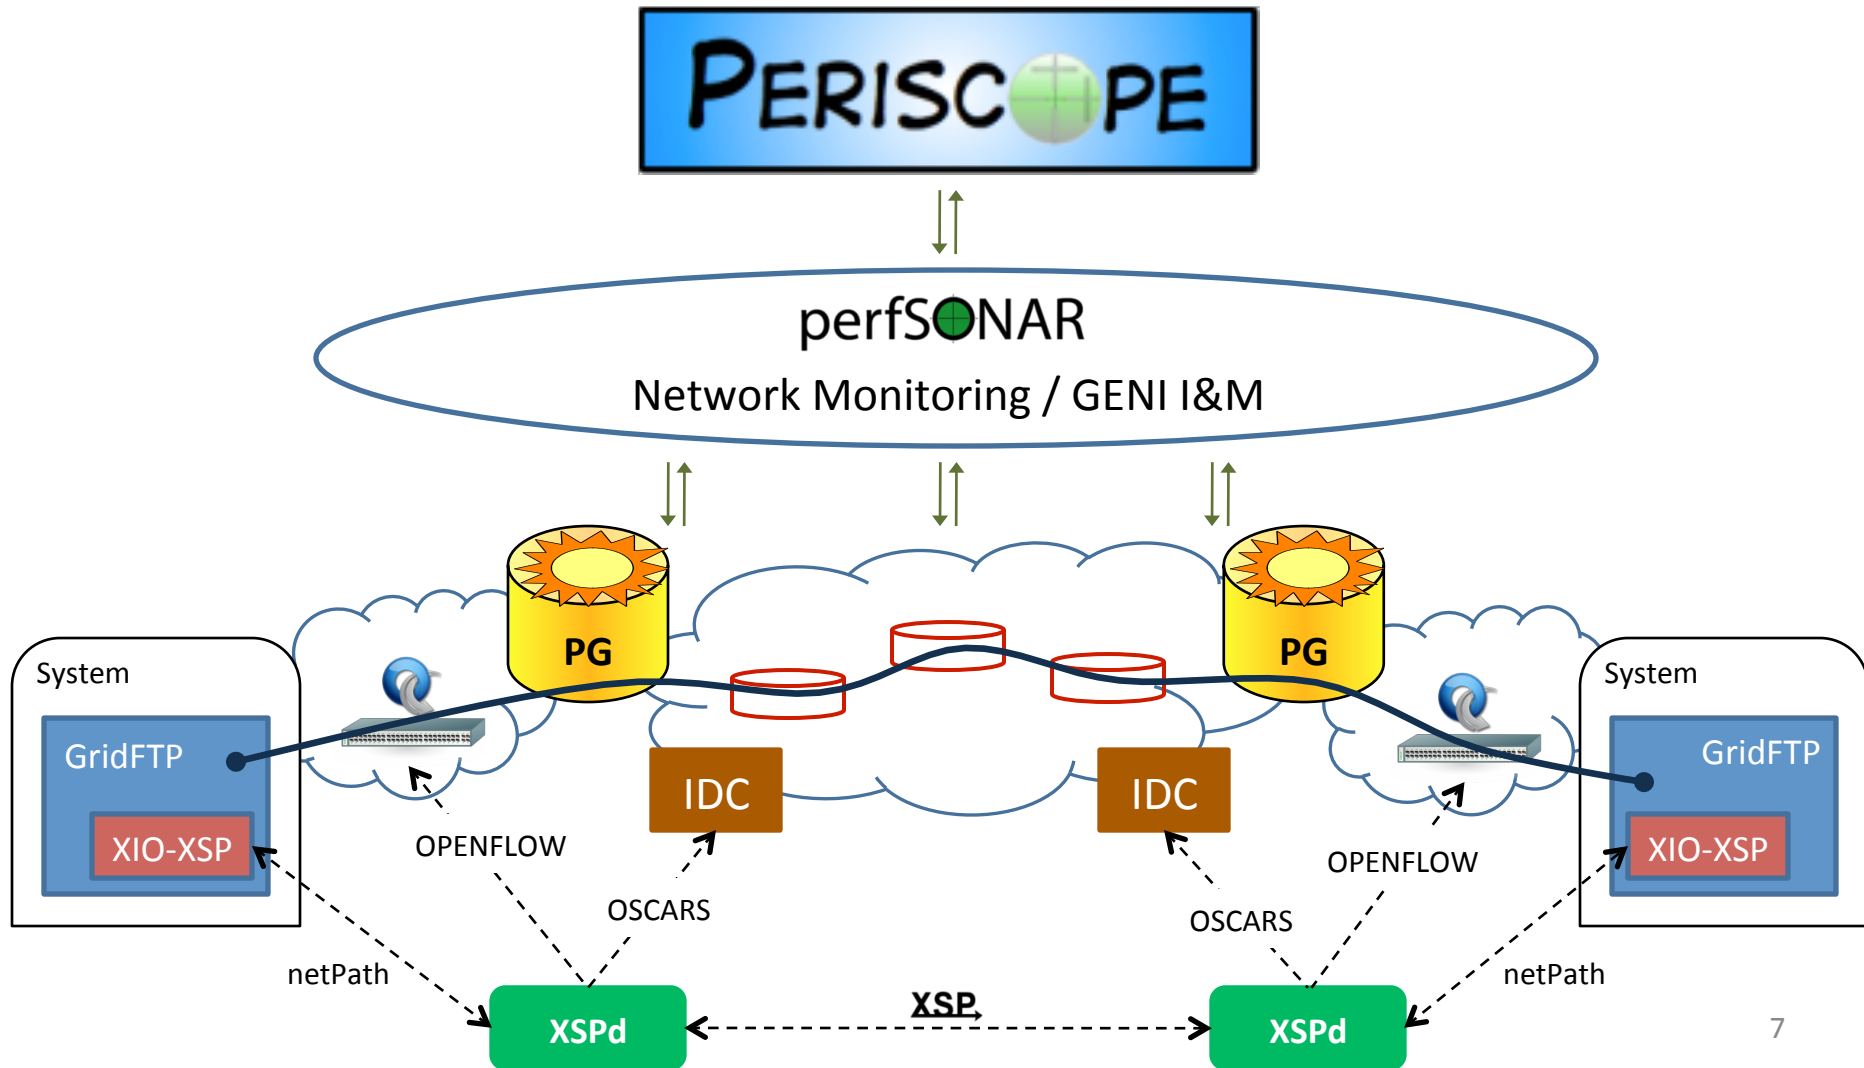
- Periscope was developed in the LAMP project (GENI I&M)
 - Extended into other efforts with various collaborators
- Measurement control and visualization portal
- Extends and leverages the widely-used perfSONAR system
- Provides caching, aggregation, visualization and control of I&M data

SRS Demo

- Collaboration with ESnet
- 10G RDMA end-to-end
 - With and without Phoebus
- Bring circuit to endhosts
 - Openflow L2 path setup
- Include GENI resources used at GEC12



Demo Architecture



GEC12 Demo Resources

