



Ph☀ebus

Network Middleware for Next-Generation Network Computing

Martin Swany

University of Delaware

Visiting Professor, CERN



Ph☀ebus



Phoebus Overview

- ☀ *Broad project goal:* Investigate next-generation network environments for high-performance distributed computing with novel middleware and protocols
- ☀ *Tenet:* Opaque, homogeneous view of the network is limiting and only becoming more so with the advent of nets like USN
- ☀ *Long-term question:* How can we augment current network models for the future?





Network Processor-Based Model

☀ Our inquiry focuses on using general-purpose Network Processors as a platform

- * *A network processor is a special-purpose, programmable hardware device that combines the low cost and flexibility of a RISC processor with the speed and scalability of custom silicon*
- * Growing in popularity, provide a lower bound on functionality

☀ The Intel IXP can sustain 10Gb/sec with a cycle budget of around ~80 operations per packet

- * We have more headroom for slower links and the technology continues to improve

☀ While this effort does not aim to provide a production-quality implementation, we are working with real hardware to stay grounded in reality

Session Layer Model

- ☀ One key enabler is the use of a Session Layer (above the Transport layer)
- ☀ End to end transport layer connections are not sufficient in many cases
 - ✱ Why do we continue to be constrained by this?
- ☀ By embracing network heterogeneity and performing explicit negotiation, we can optimize the end to end flow and functionality
- ☀ Our model assumes “session processors” in the network
 - ✱ May or may not be in the actual data path



Session Authentication

- ☀ Strong, explicit authentication of a “session” enables significant functionality
- ☀ The Grid uses X.509-style certificates
- ☀ A problem with many network protocols is the lack of authentication model
 - ✱ So we layer it on top with ACLs, holes poked in firewalls
- ☀ Simple session layer interaction can use certificates to establish an ACL
 - ✱ This works today
- ☀ Authentication can cascade with proxy certificates or the user can authenticate with each session processor



Authorization and Policy

- ☀ Another key component for next-generation network systems is the application of authorization and policy
- ☀ The creation of an explicit session provides a locus for the application of policy
 - ✱ This pervades the other aspects of the session processing
- ☀ Two simple examples
 - ✱ Policy-based signaling from a trusted location
 - ✱ Don't have to worry about ACL rot
 - ✱ We have designed a link-state routing protocol for flows that addresses the case where administrative domains touch in multiple places



Signaling and Traffic Engineering

☀ Having authenticated the user at the edge, the session processor can initiate or modify signaling requests

☀ Three Examples:

- ✱ Insertion (or validation) of ER options in RSVP-TE Path messages
- ✱ Mapping between provider-specific DSCPs at line speed based on policy
- ✱ Fusion and Fission of LSP-based traffic for inter-domain connectivity
 - ✱ Esnet and Internet2 for instance

Multiple Transports

- ☀ The part of the our session protocol that has been perhaps the most controversial
 - ✱ Or the “terrible idea that will never work” (but does) and which continues to be reinvented with increasing frequency
- ☀ Short-circuit TCP based on an established session
 - ✱ Possible on dedicated hardware
- ☀ Convert from TCP to a rate-based transport protocol in the network
- ☀ Analytical work on signaling granularity



Conclusions

- ☀ Continued exploration of a promising new paradigm
- ☀ Single paradigm and protocol that addresses many of the problems faced by advanced DOE networking in the future
 - ✳ This reinforces its promise in my opinion
- ☀ Currently supporting Ph.D student Aaron Brown
 - ✳ 2nd Student beginning in Spring

