

Large Scale Science, The Science DMZ, SDN/OpenFlow, Security and Cyberinfrastructure Architectures

Joe St Sauver, Ph.D.

(joe@internet2.edu or joe@uoregon.edu)

Internet2 Nationwide Security Programs Manager
and InCommon Certificate Program Manager

Internet2 Member Meeting, Philadelphia PA

<http://pages.uoregon.edu/joe/sdn-science-dmz/>

Note: These slides are provided in a detailed format to ensure accessibility for the deaf and hard of hearing, and for ease of search engine indexing.

Disclaimer: all opinions expressed are those of the author and do not necessarily represent the opinion of any other organization or entity.

This Isn't My First Architectural "Rodeo"

- At the 9/2005 Member Meeting, I delivered "**Thinking About Lambda-Based Network Architectures and Your Applications,**" (see pages.uoregon.edu/joe/lambdas/). In considering some of the wave-based offerings that were coming online then, I suggested that at many sites, waves weren't really needed, and application requirements could be easily satisfied by simple 10Gbps packet connectivity.
- In 4/2008, I did another architectural talk entitled "**Cyberinfrastructure Architectures, Security and Advanced Applications**" for the Spring Member Meeting (pages.uoregon.edu/joe/architectures/architecture.pdf). In that talk, I looked at some of the ways that firewalls can be helpful, and some of the ways that firewalls can end up hurting us, noting that: *"I believe the time has come for us to move beyond the traditional firewall..."*

Fast Forward: Architectural Issues Today

- Today, as was true years ago, too, researchers doing large-scale science still rely on high throughput networks to share data between collaborating sites.
- Many such sites, under pressure from ongoing cyber attacks, have reacted by deploying perimeter firewalls. These devices break the classic end-to-end transparent-network model (ala RFC2775). Suddenly, instead of just forwarding packets or passively monitoring traffic, some appliances become an "active party" to the conversation.
- When that happens, some scientists may find that security appliances interfere with their ability to actually do their work, and obviously, that's not what we want.

An Example: High Throughput Flows

- Firewalls may act as a **choke point** for high throughput flows. Most firewalls are neither designed nor built to handle sustained traffic at even 10Gbps. As a result, behind some firewalls, throughput may be disappointing.
- In fact, if firewalls are unable to keep up and **drop packets** (even if only rarely), it may be difficult or impossible to sustain reasonably fast TCP throughput.
- When that happens, flows may take longer than they otherwise might -- sometimes many hours or days -- and that poor performance can trigger still other firewall-related issues.

Long Duration Flows

- Enterprise-oriented stateful firewalls, optimized for enterprise-class loads, may not cleanly handle **long duration flows** characteristic of data-driven science.
- That is, most stateful firewalls are meant for short or "chatty" (highly interactive) traffic environments. In that world, firewalls can readily distinguish between (a) connections which are still active (and which thus must be left alone), and (b) other connections which may have been summarily abandoned (and which thus can be safely reaped during periodic state-table housekeeping).
- Unfortunately, some normal long duration science data connections may appear abandoned even when they're not

Clawing Our Way Back Toward Transparency

- Given the reality that it may be impossible to completely fight firewall encroachment and related issues, one incremental approach that some have tried is the "**Science DMZ.**" See fasterdata.es.net/science-dmz/
- From a researcher's point of view, an important feature of the Science DMZ is that **systems in the Science DMZ are NOT behind a firewall, although typically they MAY still be sheltered behind router access control lists (ACLs).** Those ACLs can be configured by network engineers to block problematic ports and addresses that aren't needed by the legitimate users of that enclave; other traffic will just transparently flow across the wire unimpeded.

BUT... We Haven't Fixed The Rest of Ye Olde Overly-Firewalled Campus

- While deploying a Science DMZ helps eliminate firewall-related performance issues **for systems in that enclave**, it is strictly a "point solution." It only fixes firewall issues for the systems that are located there.
- It does *not* address, nor does it *claim* to address, the more general-case needs of scientists working from regular network connections in their labs or offices. Those scientists will typically continue to be "protected" by campus perimeter firewalls, and many may continue to struggle with firewall-related issues as a result.
- We must continue to press on this important problem

"So What About SDN/OpenFlow?"

- Another part of today's network "recipe" is OpenFlow/SDN. To understand its potential relevance, as I mentioned in my earlier 9/2005 "Lambdas" talk, there are two sorts of "network research":
 - research conducted OVER or VIA the network, and
 - research ABOUT networking
- If you're doing research ABOUT networking, SDN/OpenFlow has tremendous potential to establish a new experimental environment in which you can work.
- If you're doing research OVER or VIA the network, you likely don't care how your bits get carried -- as long as they get where they're going, fast and affordably.

The SDN/OpenFlow Value Proposition for Scientists Doing Work Over The Network

- Because SDN/OpenFlow is still new & a work in progress, we don't yet fully know what technical innovations it may bring to those doing work OVER the network. For now, from the perspective of a scientist, the network is still TCP/IP, whether the transport is OpenFlow/SDN or not.
- Thus, if I'm an experimental scientist, and I don't care about how you carry my traffic (as long as it gets where it is supposed to go fast and affordably), the basic value proposition for SDN/OpenFlow may be that using SDN/OpenFlow will allow us to **go faster/get more bandwidth in an affordable sort of way today.**
- That's an important accomplishment, and one consistent with my September 2005 call urging sites to deploy fast, simple packet connectivity rather than focusing on waves.

"What About SDN/OpenFlow and Security?"

- Security people, including me, normally worry about three fundamental issues, the so-called "C-I-A" objectives of

confidentiality,
integrity, and
availability.

- For instance, "Can I eavesdrop on what's being sent?" "Can I potentially modify your traffic without you noticing?" and "Can I deny you the ability to use your systems or network altogether?"
- I think it is likely too soon for us to know if there are problematic C-I-A issues associated with SDN/OpenFlow, although there certainly are a number of open questions...

Examples of Some Open Questions

- Does use of a **controller-based architecture** increase a network's vulnerability to denial of service attacks? Or does use of that architecture actually make it **easier** to filter such attacks on distributed infrastructural devices?
- Is control plane traffic **always** adequately protected against eavesdropping and tampering? For example, the OpenFlow 1.2 spec at section 6^[1] mentions that "The OpenFlow channel is **usually** encrypted using TLS, but **may** be run directly over TCP." [emphasis added]
- If I can force an OpenFlow switch into "fail secure mode" or "fail standalone mode" [per OpenFlow 1.2 Specification 6.4], can I undermine the integrity of controller-based security processes? Can I live with that potentiality?

[1] <https://www.opennetworking.org/images/stories/downloads/specification/openflow-spec-v1.2.pdf>

Other Critical Details Are Still "Imprecise"

- As someone who's been increasingly focused on TLS as a result of my work with the InCommon Certificate Service, I find insufficient details in the OpenFlow 1.2 spec about security critical details, such as how TLS should be used.
- In **"SSL/TLS Certificates: Giving Your Use of Server Certificates a Hard Look"** from the 10/2011 Member Mtg, (see pages.uoregon.edu/joe/hardlook/) I explained a number of the ways that TLS can fail, including some as simple as continued reliance on old/insecure versions of the SSL/TLS protocol, or allowing weak cipher choices.
- The OpenFlow 1.2 spec, like many SSL/TLS-relying applications, takes a lot for granted and really doesn't "drill down" as precisely as it should in this important area. As a result, weak crypto may end up being used.

The Other Side of The Coin

- OpenFlow/SDN is not just about potential new issues. It may also help us **fix some of the gaps** that many of our networks have when it comes to network operations.
- In particular, we know that when it comes to **intra-subnet traffic** -- traffic that normally never passes beyond a local ethernet switch -- that traffic tends to have **very limited visibility**. This is an example of a network instrumentation "blind spot" that OpenFlow/SDN may potentially allow us to scalably fix for the first time (at least for selected "interesting" intra-subnet traffic).
- OpenFlow/SDN may also improve our ability to scalably and efficiently **mitigate** potentially compromised hosts.
- Wide area OpenFlow/SDN transport, since it is working at layer two, neatly sidesteps some persistently troublesome **routing security issues** that arise at layer three.

Bottom Line: Security Researchers, We Need Your Help

- If you are a security researcher, or know someone who is, I have a personal request I'd like to ask of you: **please add SDN/OpenFlow to your security research agenda!**
 - Are there **opportunities** to leverage SDN/OpenFlow capabilities to **improve** the security of our networks?
 - Are there **protocol-level vulnerabilities** in the SDN/OpenFlow protocol that need attention?
 - Are there security-specific **implementation flaws** that exist in an individual vendor's SDN/OpenFlow code? We need those to be found and corrected, too.
- We need your help to identify & tackle these issues!