Table Of Contents

A Glimpse At Two Approaches To Segment Routing

By Ethan Banks

Old timey network engineers might remember source routing. By default, we turned it off as a security risk, mitigating a potential MitM attack. Best path routing or bust, baby.

Sadly, best path routing isn't sufficient for many networks. Artisanal routing techniques are all the rage, often to meet service chaining needs.

Two emerging artisanal techniques are Segment Routing over MPLS (SR-MPLS) and Segment Routing over IPv6 (SRv6). As Segment Routing is the new hotness, let's squint at these two approaches.

SR-MPLS Vs. SRv6

Both SR technologies encode path information in the packet, but each one does it differently.

1. SR-MPLS uses a label stack to describe the desired path through the network. Label Switch Routers observe the label, pop it, and forward.
2. SRv6 uses a Segment Routing Header (SRH) embedded in an IPv6 packet to describe the desired path through the network. SRH-aware nodes read the header, update a pointer, swap the destination address, and forward.

Both SR approaches require little of the routing infrastructure itself. No SR state is maintained in the core network, as each packet has forwarding instructions embedded in it.

Instead, an edge router, or possibly a controller, creates the paths and encodes the label stack or SRH in the packet before it's sent into the core.

SRv6 Adoption Challenges

SR-MPLS has seen some adoption, particularly among service providers. SRv6, however, lags. Why? In a nutshell, hardware.

SR-MPLS asks nothing special of ASICs when it comes to forwarding. There's a label. Perform a lookup. Pop the label. Forward the packet. Yes, there is control plane software specific to SR-MPLS, but that doesn't impact the ASIC's ability to forward packets. As a result, SR-MPLS has relatively low adoption costs.

By contrast, SRv6 asks something special of ASICs. The SRH is a new IPv6 header to interact with, and it's still an IETF draft. SRv6 nodes must take several actions along an SR path, including reading the SRH, re-writing the IPv6 destination field to the next node in the path, updating a pointer, and performing a node-specific action.

None of these requirements are insurmountable, but as with any new packet-level technology, asking ASIC designers to commit features

requires serious customer demand. Why bake features into silicon that no one will use?

Ameliorating the situation somewhat is that non-SRv6 capable routers forward the packet normally, ignoring the SRH. Not every network node must cope with the SRH.

Hardware & Software Support For SRv6

Thus far, I've found evidence of SRv6 support in Cisco's NCS5500 and Nexus 9300GX, and in Barefoot Networks' programmable Tofino chip. In addition, I found a rumor that Juniper might support SRv6 in its Penta chip. That makes sense, as the Penta is also programmable. There might be more platforms supporting SRv6 I didn't spot.

In a Tech Field Day presentation, Cisco's Jakub Horn stated, "Our first hardware implementation right now is on two platforms. One is the custom silicon. Second is the merchant silicon. Both are shipping right now in [IOS-XR] 6.1 release. Obviously, capability-wise, they are slightly different. Custom silicon is slightly better than the merchant...Both are shipping right now."

In software, the Linux kernel supports SRv6 as of kernel versions 4.10 via the SREXT kernel module. That's an option if you'd like to hack around or roll your own with a whitebox. The open-source FD.io project also supports SRv6.

Those software options sound like reasonable edge computing use cases, as service providers want to do fancy service chaining in central offices. There's enough power in x86 paired with the right projects to handle the forwarding load in an edge scenario.

Who's Driving SRv6?

Cisco seems to be the driving force behind SRv6. Cisco's made a chip investment. Cisco folks have also done much work in the IETF related to SRv6, and seem tied to the FD.io SRv6 implementation.

Why? I'm not sure. In Cisco's SRv6 presentations, the main theme is network simplification. That is, operators can deliver the same apps

and services they did with an MPLS stack, only with a simpler IPv6 fabric. Cisco has all the pieces--hardware and software offerings driving a complete SRv6 solution.

Simplification is an interesting point, but I'm not sure MPLS is that complex of a problem for operators. In addition, any significant change to an SP network is going to be a slow burn. Brownfield is hard, even when new tech is installed as patches of green.

Will "simplification" drive SRv6 adoption? At the moment, I don't see it.

What's Next For SRv6?

I'm sorting out my opinion on where SRv6 might fit into the networking world. If you have an opinion, let me know via ethan.banks@packetpushers.net. I'm writing a hoopla-free technical whitepaper on SRv6 for publication on Ignition, and my research continues.

Sponsor: VIAVI Solutions

When NetOps And SecOps Worlds Collide

Whether managing end-user issues related to network performance or assisting in the identification and remediation of security breaches, one thing is clear – IT requires the right data, and context, to quickly identify and resolve any issue.

The VIAVI Observer platform bridges NetOps and SecOps, providing a single data set that captures and stores the most robust wire data available for issue identification and remediation.

This pure, unaltered packet and flow data is analyzed to present views across the entire IT landscape with simple end-user experience scoring and threat alert features. Use automated workflows to dive

into high-fidelity network evidence and solve the root-cause of the issue – minimizing impact on users and business operations.

Learn more by visiting www.viavisolutions.com/packetpushers. Access free videos, guides and content, and request a free copy of our popular OSI Layers 1-7 Enterprise Protocol Poster.

Tech Blogs: How To

How to Bootstrap Kubernetes the hard way! - Polar Squad

https://medium.com/polarsquad/how-to-bootstrap-kub...

Many “K8s up and running” tutorials exist. There’s even a book. I found this particular tutorial interesting in that it eschews the miniKube approach for a more grown up method. That’s why it’s “the hard way.” You’ll bring up an actual cluster, configure networking, and move all nodes to a ready state. Steps are written with context, so that you’ll know why you’re doing what’s been instructed. - Ethan

Chaos testing with RIFT-Python: an open source implementation of the Routing In Fat Trees (RIFT) - Bruno Rijsman on YouTube

https://www.youtube.com/watch?v=GqebgPmA4Xc

RIFT is an emerging, specialized protocol for routing in large data centers. “In this video we give an overview of the RIFT-Python implementation, and give detailed step-by-step instructions on how to do ‘chaos testing’.” Created for the IETF 104 hackathon. - Ethan

Building the Network Automation Source of Truth - IP Space

https://blog.ipspace.net/2019/02/building-network-...

OK, technically this isn’t a how-to blog post. It’s more of a “Hey, this is a really complicated problem, and here’s why” kind of post. In any case, as Ivan points out, before you can really get into network automation, you’ll need a single source of truth for all the devices and their configurations on the network. That means you’ll need a reliable device inventory. And a reliable data source. In a format that’s useful to you and your automation tools. And a mechanism to update the inventory when configs change and new devices get added. And...you can see where this is going. - Drew

The Advanced VMware vSphere NIC Teaming Guide - vMiss.net

https://vmiss.net/2019/02/22/vmware-nic-teaming-ad...

NIC teaming is configuring two or more network interfaces to act as a single logical interface. The hashing algorithm handles load-balancing across the teamed links. Choosing an appropriate hashing method is important if your goal is to end up with a roughly even distribution across the links. Melissa Palmer talks through configuring IP hashing and source MAC hashing in a vSphere environment. - Ethan

SSH to all of the serial ports - Fragmentation Needed

https://www.fragmentationneeded.net/2019/02/ssh-to...

Chris Marget shares, “a quick-and-dirty script for logging into every serial port on an Opengear box, one in each tab of a MacOS terminal.” If you need to know this, then you need to know. - Ethan

Cross-Data-Center L4-7 Services with Cisco ACI - ipSpace.net

https://blog.ipspace.net/2019/02/cross-data-center...

Craig Weinhold has detailed knowledge about how ACI, especially as of the 4.0 release, is forwarding traffic in a multi-site scenario. He says, “I’m a cynic and a skeptic, but it’s clear that Cisco ACI is making strong progress in delivering an ‘easy button’ multi-tenant, multi-DC

architecture that effortlessly supports clean L3, stretched L2, and stretched L4-7.” - Ethan

Virtual Design Clinic: Join Us March 14th!

The Packet Pushers are hosting a new Virtual Design Clinic on Thursday, March 14th

This is a live, online event that includes technical presentations, Ask Me Anything sessions with industry experts, and a sponsored presentation.

It's free to register and attend. Sign up here.

If you can't make the 14th, we'll post the full event on our membership site, Ignition, which is also free to join. You can check out previous VDCs at Ignition at your leisure.

Tech Blogs: Opinion

Choosing a Cloud Provider Doesn't Have to be a Binary Choice - Gestalt IT

https://gestaltit.com/exclusive/ken/choosing-a-clo...

Ken Nalbone argues that it's worth looking at cloud specialists for certain computing needs as you might save money and prevent lock-in. Ken cites a new partnership between Packet and Wasabi to bolster his point. "Customers accessing Wasabi's object storage will see no egress charges or additional transit fees when leveraging Packet's compute resources." I agree with the idea of considering non-traditional cloud services, but I wonder if developers are driving cloud decisions more than operators or business managers. - Ethan

Contemplating APIs and the WLAN State of Things - wirednot

https://wirednot.wordpress.com/2019/02/22/contempl...

Programming, APIs, and wireless networks. Lee Badman has some questions. My favorite? "If a vendor has historically put out crappy code that is developer-driven versus engineer-driven, how do we trust the developers to get it right when it comes to what data awaits engineers at the end of the APIs?" Tough, but fair. - Ethan

Will transport innovation collapse the Internet? - Christian Huitema

https://huitema.wordpress.com/2019/01/12/will-tran...

Greg spotted this article and then asked Christian to record a podcast with us, which he did. That recording will be released on 19-March-2019 as a Heavy Networking episode. We recommend you read through this article before listening to the show to refresh your knowledge on Internet congestion, TCP variants such a Reno & Cubic, QUIC, and BBR. This will help you get the most this forthcoming, very heavy Heavy Networking episode. - Ethan

Tech Office Sketches - ongoing by Tim Bray

https://www.tbray.org/ongoing/When/201x/2019/02/15...

This piece by Tim Bray, a software nerd working for Amazon, should be interesting to all knowledge workers. Lots of office topics covered, with interesting opinions such as, “All the high-tech companies I’ve worked for have resolutely ignored the research I hear about that seems to say putting expensive engineers out on the floor with no separating walls leads to grievous productivity losses. Isn’t this biz supposed to be data-driven?” - Ethan

Automation: Tools vs Languages - ipEngineer.net

http://ipengineer.net/2019/02/automation-tools-vs-...

David Gee takes some dogma out of the religions that form around workflow engines and programming languages with a grown-up pros &

cons analysis of various approaches. Spoiler alert. There is no one right answer. - Ethan

What Are the Networking Requirements for HCI? - SNIA on Ethernet Storage

http://sniaesfblog.org/what-are-the-networking-req...

A short Q&A follow up to a webinar presentation that SNIA did on hyperconverged infrastructure. Some good questions here focused on architecture and the NVMeoF types likely to be relevant for HCI. - Ethan

The Lulz

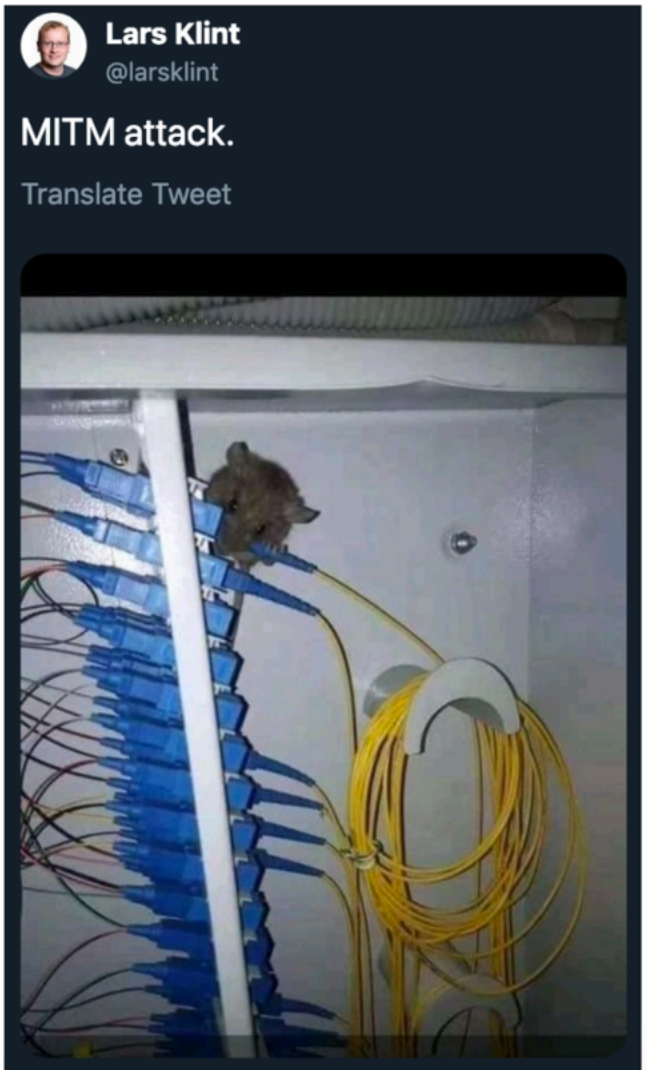

MOUSE IN THE MIDDLE ATTACK - SHARED BY LARS KLINT @LARSKLINT
https://twitter.com/larsklint/status/1098166163828...

Tech News

Changes to APNIC IPv4 Maximum Delegation Size and IPv4 Recovered Pool Waiting List - APNIC Blog

https://blog.apnic.net/2019/02/28/changes-to-apnic...

“The APNIC Executive Council (EC) today announced an interim arrangement for IPv4 delegations from 103/8, as of Thursday, 28 February at 12:00 (UTC+9), to immediately reduce the maximum delegation size from a /22 to a /23.” - Ethan

Suicide instructions spliced into kids’ cartoons on YouTube and YouTube Kids - ArsTechnica

https://arstechnica.com/science/2019/02/youtube-ki...

Who’s to blame? Is YouTube a video streaming platform (like an ISP), or a content provider (like Netflix)? When YouTube is an open platform, is it even possible to safeguard children from seeing inappropriate content pretending to be tyke-friendly? Maybe we’re all to blame. Parents assume YouTubeKids is safely curated. YouTube clearly isn’t curating. Society substitutes screens for substance. We keep opting in. Maybe we need to start opting out. - Ethan

Nike says it’s ‘actively working’ to fix its broken smart sneakers - The Verge

https://www.theverge.com/circuitbreaker/2019/2/21/...

Nike engineers are working on a bug fix for their $350 “smart” sneakers. What makes them smart? These shoes are self-lacing, and use Bluetooth to pair with an app on your smartphone to remember your preferred tightness setting when you put them on--unless the shoes can’t pair with the app and won’t tighten, which is why Nike is working on a patch. Or you could just skip all this nonsense and tie your shoelaces like a grownup. - Drew

128 Technology Is Starting a Revolution for Secure SD-WAN, Networking - SDxCentral

https://www.sdxcentral.com/articles/news/128-techn...

This is a silly clickbait title, but I raise this article because 128T is unique in how they deliver secure WAN. SD-WAN is just one of 128T’s use cases. Worth reviewing, at least for the exercise in thinking through a technology approach you probably haven’t run into before. In September 2016, I recorded a presentation at 128T’s headquarters when they hosted a BOSNOG event. 128T has also appeared at a recent Tech Field Day--scroll to the bottom of this playlist to see those sessions. You might or might not like what 128T is doing, but they have won some converts. - Ethan

Verizon and Cisco to Extend Software-defined Networking to Support Future 5G Mobile Devices for Enterprise - PR Newswire

https://www.prnewswire.com/news-releases/verizon-a...

https://www.sdxcentral.com/articles/news/cisco-and-verizon-build-a-mobile-sd-wan-offering-with-5g/2019/02/

The Verizon and Cisco relationship around SD-WAN goes back a ways, so this announcement is no surprise. The SD-WAN tech employed to add 5G support is Viptela, and apparently leverages “Cisco's intent-based networking capabilities.” Yikes. That’s a lot of stuff to integrate. Glad 5G is still future tech. Maybe they’ve got time to get this 5G/SD-WAN thing right. - Ethan

New Products & Industry Takes

What is the Point of AIOps? - ThousandEyes Blog

https://blog.thousandeyes.com/what-is-the-point-of...

A thoughtful piece that cuts through the growing AIOps hype. “Application disaggregation, the evolution to virtualized and containerized infrastructure, overlay and underlay networks, the ubiquity of cloud and SaaS-based service components communicating over the Internet in nearly every application, network and infrastructure means that correlation is devilishly difficult yet still critically important. That’s where AIOps is defined as the solution: using advanced analytics to consume various streams of monitoring data and do the correlation that humans can’t do via swivel-chair analysis.” - Ethan

VMware Releases NSX-T Data Center 2.4 and NSX Cloud - VMware

https://www.vmware.com/company/news/releases/vmw-n...

You'll have to wade hip-deep through marketing hyperbole and needless adjectives to get there, but you'll find somewhat useful notes starting about halfway down. Highlights include, "a new declarative API model that simplifies network automation using human-readable JSON configuration." - Ethan

Amazon Web Services to integrate its Cloud Services with Sprint's Curiosity IoT Platform to Bring Actionable Intelligence to the Network Edge - Sprint

https://newsroom.sprint.com/amazon-web-services-to...

Sprint is partnering with Amazon AWS around IoT. Sprint's Curiosity IoT platform provides LTE connectivity for IoT devices along with local compute capabilities. In a new partnership, customers can send IoT data from Sprint's Curiosity platform to S3 for storage, to AWS's IoT Analytics service for analysis, and other AWS services. - Drew

CNCF Launches Cloud Native Network Functions (CNF) Testbed - Cloud Native Computing Foundation

https://www.cncf.io/announcement/2019/02/25/cncf-l...

The Cloud-Native Computing Foundation, which is part of the Linux Foundation, has released a new testbed to let organizations such as

carriers and service providers test network functions running in a Kubernetes environment vs. network functions running on OpenStack, on the same hardware. The goal is to “showcase the performance improvements from avoiding virtualization overhead.” OpenStack is an independent open-source project not affiliated with the Linux Foundation. In other words, this reads to me like the CNCF, which backs Kubernetes, throwing a little shade at OpenStack. - Drew

Announcing Updated Professional-Level AWS Certification Exams - AWS Blog

https://aws.amazon.com/about-aws/whats-new/2019/02...

AWS Certified Solutions Architect – Professional exam and AWS Certified DevOps Engineer – Professional exam have been updated. I also caught this lurking at the bottom of the post. “We now only recommend, rather than require, that candidates have an Associate-level certification before taking Professional-level exams.” Save some money if you really know your stuff. - Ethan

AT&T Bridging 5G and SD-WAN with VMware SD-WAN by VeloCloud for a New Layer of Control at the Edge - AT&T

https://about.att.com/story/2019/att_5g_sdwan.html

It’s a buzzword bonanza as AT&T and VMware VeloCloud announce a partnership that will blend SD-WAN with 5G. At present, this announcement is a statement of intent rather than actual shipping product, but AT&T says it will be able to integrate fine-grained application steering within its 5G network. For instance, if a factory site

uses SD-WAN to send regular business traffic over a wired Internet link and the 5G network for manufacturing traffic, it will be able to parcel “slices” of the 5G connection so that time-sensitive robotics can get a low-latency connection, while other manufacturing apps get a best-effort slice of the 5G connection. - Drew

AWS Site-to-Site VPN Now Supports IKEv2 - AWS Blog

https://aws.amazon.com/about-aws/whats-new/2019/02...

AWS is catching up with the times. The title says it all. - Ethan

Now add up to 1000 static routes per VPC route table - AWS Developer Forum

https://forums.aws.amazon.com/ann.jspa?annID=6554

A bit of an obscure announcement, but the forum post states, “You can now add up to 1000 non-propagated/static routes per route table in your Virtual Private Clouds (VPC). Until today, you were limited to a maximum of 125 routes per route table.” I can imagine several scenarios where this might be helpful. - Ethan

The End Bit

Human Infrastructure is bi-weekly newsletter with view, perspectives, and opinions. It is edited and published by Greg Ferro and Drew Conry-Murray from PacketPushers.net. If you'd like to contribute, email Drew at drew.conrymurray@packetpushers.net.

We don't give away your email address or personal details because that would suck.