Another week of being a Human in IT Infrastructure. Are you still alive to talk about it?

View this email in your browser

*IN THIS WEEK'S ISSUE: IPv6 In The Enterprise: Why Bother?; Sometimes Vendor Lock-In Isn't So Bad. Please remember to enable the images; the magazine looks a lot better that way!*


**Table of Contents**
**(aka The Project Plan)**



## Issue Number 57

**05/25/2017**

The "Let's get real" issue.

## Thought For The Week:

Finite—to fail, but infinite to Venture - Emily Dickinson

# 1. IPv6 In The Enterprise: Why Bother?

## by Greg Ferro

In 2004, I spent a few hundred hours getting on top of IPv6 on the understanding that it was going to be a critical technology. I thought I needed to be ahead of the curve.

Fast forward to 2017 and ... what a mess.

Your enterprise network does not need or want IPv6. IPv6 does not reduce the cost of equipment, simplify or secure the network, or indeed change anything for the better.

The reality is that IPv6 is only needed on the Internet. You can put IPv6 on your public gateways and services to meet client needs, but keep everything else the same.

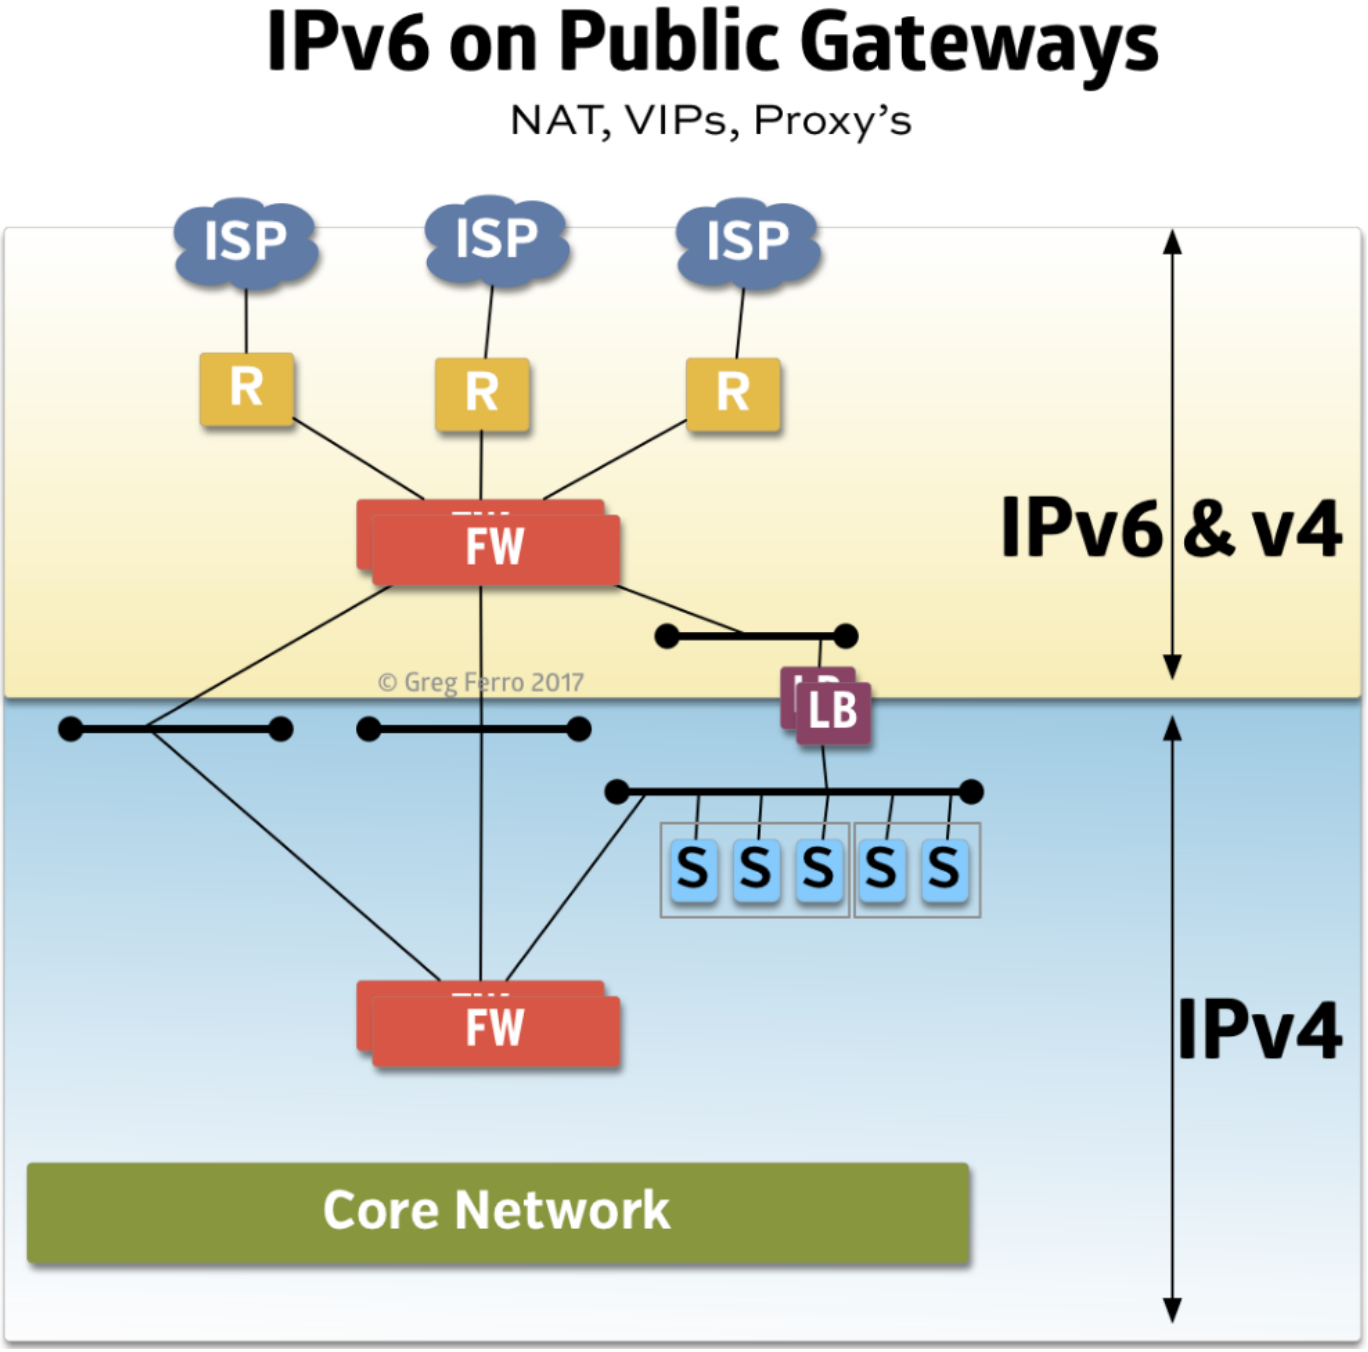


I agree there are long-term benefits to having IPv6 everywhere, but I'm tired of trying to convince IT managers that allocating budget is worthwhile. Because it isn't. There is zero value to the business to upgrade, especially when the quality of vendor implementations is so low.

## Vendor Quality

Vendors have been slow to implement IPv6 because there isn't any money in it for them. This also means that the quality of the IPv6 code has been rubbish because there are no customers to test it. Vendors, especially Cisco, will ship code early so that customer can find and report the bugs (incidentally justifying the cost of expensive SmartNet maintenance). One does wonder if testing is being cut back to maintain those 65% gross profit margins.

Early IPv6 implementations were so bad that I figure they gave the project to their interns or worst programmers.

## Some Examples

What's wrong with IPv6?

1. Too many address types. Link Local, Multicast (four types!), Anycast. Then there are all the address scopes such as interface local, link-local, realm-local, site-local, and so on. Let's pretend that neighbor discovery and MPLS aren't relevant (but they are).

2. Typing an IPv6 address is impressively difficult. Using a semicolon as a delimiter is truly dumb, and the rules for representation are just plain weird.

3. Multicast for ARP. No one likes or wants Multicast and implementing IP multicast in silicon is a serious technical problem.

4. Extension headers. Extension headers are never going to be supported in ASICs or hosts because of memory considerations. They create unnecessary state in the network core, which doesn't scale.

5. False security. Claims that IPv6 is secure are false. Link-local is insecure and RA-Guard has serious flaws.

6. Confusing set of IPv4-to-IPv6 compatibility tools. I don't even know where to start here: 6in4, 6RD, NAT64 (NAT46), 464XLAT,

DS-Lite are the most common.

7. DHCP still isn't finished. Attempts at auto-configuration seem to have failed or lost momentum, and now we are waiting for DHCP to finalize.

I could go on, but there isn't enough space. And I'm bored of the issue already.

The IETF doesn't seem to know when to stop messing around with standards. IPv6 has been a moving target since 1999 with a consistent stream of changes, updates, modifications, extensions, and fussing. Vendors have made various attempts to hijack the process to speed it up, slow it down, derail it, or initiate different things.

Of course, that's the nature of a community/open effort; those who participate get to make the decisions. Unfortunately, enterprise IT doesn't have many people representing their needs. Certainly the vendors aren't.

## Why This Matters

IPv6 is a mess. The only need for IPv6 is when you have customers in Asia on the Internet where IPv6 is widespread in the public WAN/Internet. Simply bang up some external IPv6 and put the full migration on hold for another five years.

For an enterprise IT network, there isn't a business case to upgrade inside the network. If you move to a cloud provider then IPv6 isn't necessary yet.

## Join Apstra and the Packet Pushers at Cisco Live!


If you're going to be at Cisco Live in Las Vegas this June, we'd love to see you at the Apstra Intent to Party Party.

Join the Packet Pushers and Apstra for dinner, drinks, and great conversation. This is an invitation-only event, so if you'd like to participate, please click here to request an invitation.

Request your invitation here and we'll see you in Vegas!

**When:** Monday, June 26, 7:00 - 9:00 pm
**Where:** The Border Grill at Mandalay Bay Resort

# 2. Sometimes Vendor Lock-In Isn't So Bad

by Ethan Banks

In the SDN world, there are some customers who insist on completely "open" technology. That is, they want whatever technology they choose to have a shot at working with whatever else they choose, because all of the tech is open source. Vendor lock-in is anathema to these customers.

The reality of SDN products is that, despite the cry of open-everything, the most useful solutions are vendor-specific. The solutions might or might not have open source components, but that doesn't especially matter to the customer who is simply trying to buy a turnkey solution that works.

Perhaps "turnkey" is the crux of the matter. Most organizations can only make use of a tool that actually does something. They have neither the time nor expertise to assemble Lego blocks into something useful. They need to open the proverbial box, plug it in, and experience success.

In this context, we see the uptake of Cisco ACI, VMware NSX, and SD-WAN solutions such as CloudGenix and Viptela. These tools perform different functions under the broad umbrella of SDN, but make a special point of emphasizing out-of-the-box usefulness—with no assembly required.

That's not to say you don't have to configure these platforms to do something. Certainly, you do. You might even want an integrator to help you get it right. But you don't have to write your own software to stack the Lego bricks together.

## Idealism Vs. Usability

The open-everything vs. turnkey mentality is one I've had to come to grips with as I've watched the networking industry change. As end users, we all think we want interoperability, openness, and the chance

to swap out any individual pieces and parts whenever we want to. Only, that's not actually true. That's idealism talking.

Idealism aside, I believe most people simply want their IT to work with a minimum of fuss. Does that mean vendor lock-in? At this point, probably. Why? Because IT infrastructure is increasingly interdependent, creating ever more complexity.


With the software defined data center, we're taking formerly disparate tech and mashing it together into a unified whole. Systems talk to one another through API calls. Central orchestration software, driven by policy, creates changes in compute, storage, networking, and security.

This makes for a fragile operating environment.

Attending an OpenStack Summit Boston 2017 keynote, it was clear that major OpenStack directives going forward are maturity, stability and usability. There's less emphasis on new features, and more emphasis on polishing what's there.

Attending a session, I saw a frightening OpenStack reference architecture. Imagine Kubernetes (K8s) as the base layer of OpenStack—the foundation that keeps OpenStack running.

And then imagine tenants running their own K8s or Mesosphere on top of THAT to create their own multi-tenant environments. Admittedly, this architecture is only interesting to a small number of cloud operators. Most operators wouldn't need the top layer that completes the OpenStack sandwich. But still…is OpenStack running on K8s a private cloud stability improvement?

What enterprise has the energy to expend on an orchestration service so fragile that it takes yet another orchestration service to keep it running? Every instinct in me as an IT architect screams, "Run away!" when contemplating such a design.

If we, as an industry, can't run OpenStack reliably on bare metal, then I think it's game over. We'll never see mass adoption. Although, in fairness, that was sort of the point I referred to from the keynote earlier. It's time for OpenStack to grow up.

That lengthy illustration returns my thinking to turnkey solutions suitable for consumption by the average enterprise. Suddenly, turnkey doesn't sound so bad, does it? Hmm.

Well, what about the no-no of vendor lock-in? At this point, I argue, "Who cares?" If you're buying on a 3 year cycle, you'll be tossing out what you've got in the not-too-distant future anyway. Therefore, go ahead. Get locked in if what you're buying is easy to use and solves

your business problems.

And by “solves your business problems,” I mean…

1. Enables your product or service to come to market quickly and remain competitive across iterations.
2. Enables excellent customer service such that you retain customers, and grow or at least maintain revenue.
3. Meets your budget requirements, with opex as critical to analyze as capex.

Many of you will benefit from partnering with the right combination of turnkey vendors to take on the technology integration work for you. If they did a good job with both their product and integration with other products, you’ll be well-positioned to make SDN work for you.


Join the Datanauts on their mission to bust silos and explore the latest developments in cloud, convergence, data centers, and more. Sign up free here.


Network Break is a weekly podcast that delivers news & analysis on the networking industry in a fun, fast-paced style. Subscribe here!

# Internets Of Interest

**A collection of pre-loved links that might interest you. "Pre-loved" because I liked them enough to put into this newsletter. It's not *true* love.**

**By Greg Ferro and Drew Conry-Murray**

## Take The Packet Pushers' SD-WAN Survey

The Packet Pushers have talked about SD-WAN a lot over the past couple of years, so we thought it might be a good idea to shut up and give you a chance to tell us what you think about it.

We put together a survey to get an idea about your interest in SD-WAN. Our goal is to get some vendor-neutral data on questions like current and potential adoption, WAN challenges, and relevant features and functions. If you have a few minutes, we'd appreciate your responses.

LINK

## On the Myth of the 10X Engineer and the Reality of the Distinguished Engineers

This post is a few months old, but Fintan Ryan offers some compelling observations about what actually makes for an excellent engineer. And while he focuses on software developers, the qualities of excellence are such cross-disciplinary attributes as humility, customer focus, and responsibility.

*"Rather than looking at what allows someone to crank out code that ultimately someone else ends up maintaining, let's look at some of the areas...that help define what makes a makes a distinguished engineer."*

---

## VXLAN: BGP EVPN with Cumulus Quagga

Vincent Bernat provides a detailed introduction to using BGP EVPN with the Cumulus Linux version of Quagga, the open-source router. The post is chock full of diagrams and code configuration samples, including three different software options for setting up route reflectors: Cumulus Quagga, GoBGP, and JunOS. If you want to get your hands dirty with BGP EVPN, check out the post.

---

## Network Engineer Would Be Systems Manager If He Could Do It All Over Again

The timing of this mock news piece from The Onion is pretty good given all the network industry soul-searching going on over programming, DevOps, and other silo-eroding trends. It also nicely captures the empty jargon that pervades the tech industry.

*"Who knows where life might have taken me if I hadn't spent the past two decades devising, configuring, and supporting communications networks? What if, instead, I were managing, planning, and coordinating IT systems?"*


# WEEKLY SHOW

Where Too Much Networking Would **NEVER** Be Enough

The Weekly Show channel is our one-hour deep dive on networking technology. Subscribe today!


# Priority Queue

Where Too Much Networking Would **NEVER** Be Enough

Priority Queue tackles niche and nerdy tech topics and cutting-edge research projects. Subscribe here!

# Product News

**Find out about interesting new products, or get essential information about things you might already be using.**

## Versa Networks Joins The Battle To Box Up Your Branch

SD-WAN is quickly transforming from a standalone product to a table-stakes feature within a broader set of capabilities at branch and remote networks. Startups and incumbents alike are adding new capabilities to—and supporting third-party virtualized network functions (VNFs) within—their branch gateways.

One of the companies driving this “branch in a box” idea is Versa Networks, which has announced that it’s adding new security functions, Wi-Fi, and third-party VNF support, to its SD-WAN devices.

LINK

---

## Pluribus Networks Gets Back On My Radar With Adaptive Cloud Fabric Announcement

Some companies I sort of lose track of. It’s the way of the technology business. Those vendors making noise tend to stay on my radar. Companies who stop engaging tend to fall off the edges of my scope.

Pluribus Networks was one of those companies that had fallen off my radar. Recently, the Packet Pushers were briefed by Pluribus, an excellent way to rekindle the relationship.

As we chatted with these folks, I was reminded of many of their product aspects that appealed to me the first time around. Here are a few Adaptive Cloud Fabric highlights if you need a bit of information to put Pluribus back on your own radar.

LINK

# Recent Podcasts

**The last five podcasts published on Packet Pushers**

## PacketPushers.net - The Last Five

Datanauts 085: Understanding In-Memory Databases

[Network Break 136: Extreme Debuts New Switches; Google Jams Cisco Smart Board]

[Show 340: OpenFlow, Fabrics & Network Virtualization]

[PQ Show 116: Practical YANG For Network Automation]

[Datanauts 084: A Fireside Chat: PowerShell & DevOps Global Summit]

# Watch This!

Where we collect some videos that make us reflect, think about our inner lives, or just entertain us.

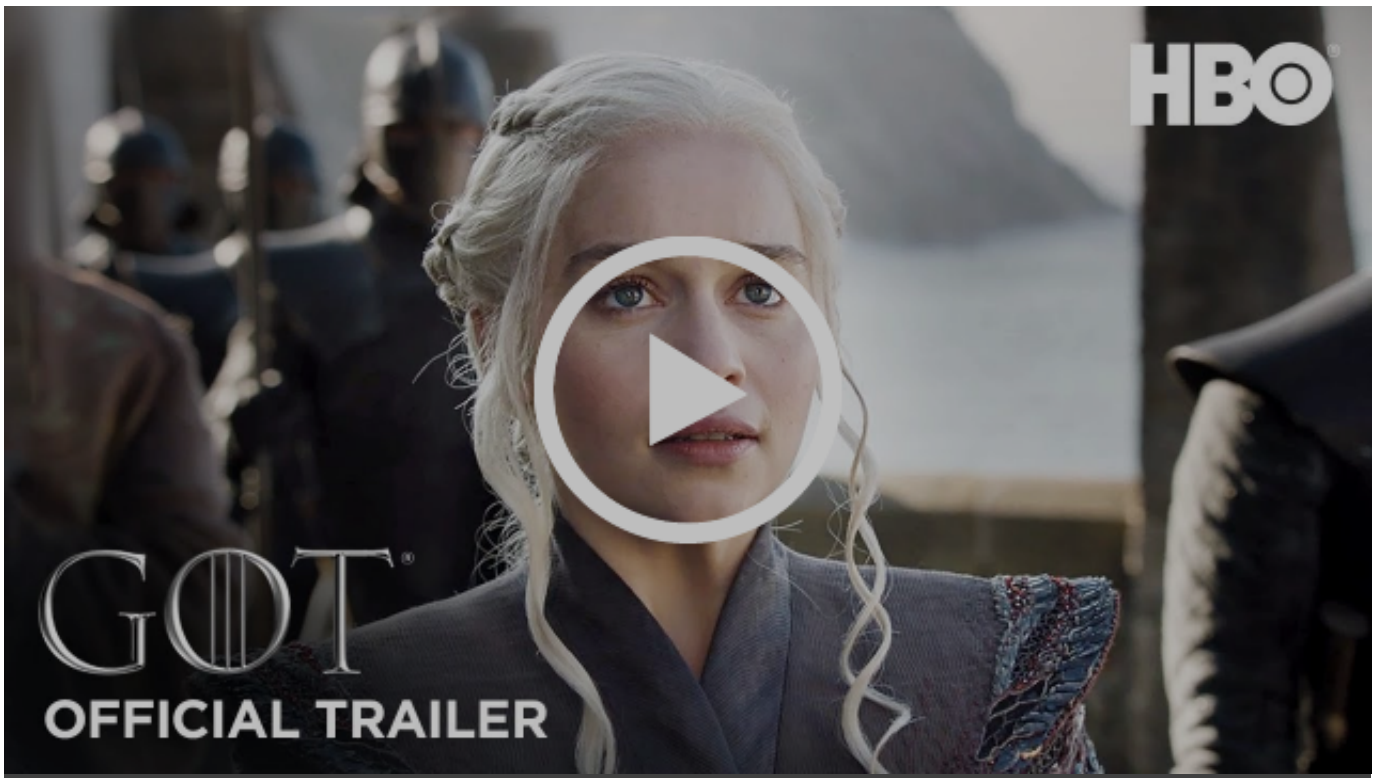


YESSSSSSS!!!!!!!


# Link Propagation Newsletter

**Our weekly newsletter delivering essential headlines, announcements, and useful news to your inbox**

Can't get enough newsletters? Check out Link Propagation, our newest publication. We send you a free weekly digest with tech news, interesting blogs, and industry announcements, all curated by the Packet Pushers. It's an easy way to keep up and stay informed. Subscribe at packetpushers.net/link-propagation.

# Quick Survey: IPv6 Adoption

Are you supporting IPv6 in your enterprise?

A. We've fully migrated to IPv6
B. We run dual-stack v4 and v6
C. We have limited IPv6 support
D. IPv6 is on our roadmap
E. No plans at this time

# Last Issue's Survey Results

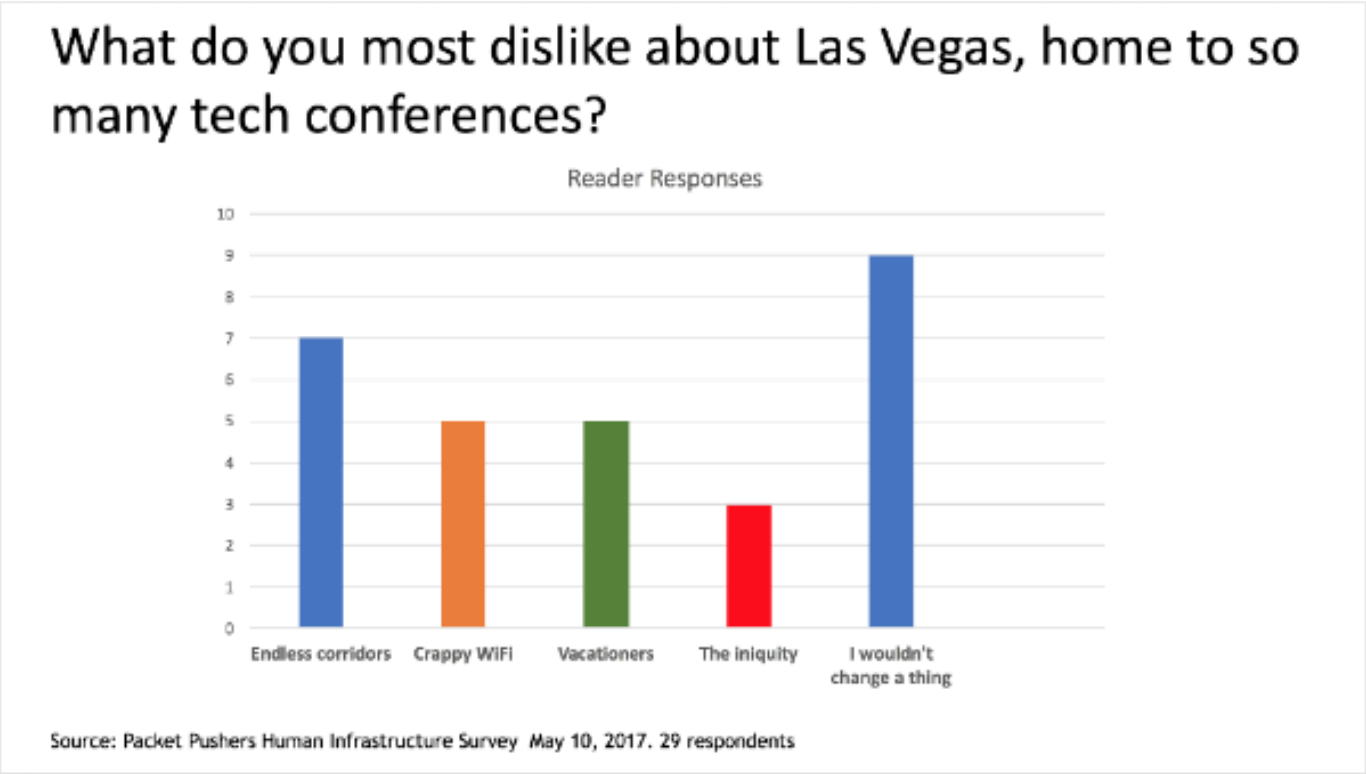


# Did We Miss Something?

Got an link or an article to share? Email it to humaninfrastructure@packetpushers.net

## The End Bit

*Sponsorship and Advertising - Send an email to humaninfrastructure@packetpushers.net for more information. You could reach 5,013 people.*

**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.*