Human Infrastructure is a bi-weekly magazine-style newsletter from EtherealMind.com with News, Views and Opinions on being a human in IT Infrastructure. Hit the signup page to subscribe and join 5,013 subscribers.

View this email in your browser

ETHEREALMIND

Human Infrastructure Magazine

A Newsletter on a Life in Networking

Issue Number 12
06/30/2015

The "Three Years is an IT Lifetime" Edition

Thought For The Week:

A good project is one that you walk away from.

- **State And Configuration**
- **Open Source Tools To Build An Open Source Lab**
- **44% Reduction In Page Load By Block Tracking Requests**
- **On My Mind: Virtual Doughnuts**
- **PacketPushers.net - The Last Five**
- **Still Here?**

This Week: A Lifetime In IT Is Three Real Years

Inspiration: How long should your IT infrastructure last?

In 1996 I was attempting to sell customers on the benefits of installing Category 6 cable with the reassurance that copper cabling would be good for 20 years. It seemed logical when companies were signing 20-year leases on their buildings to have the cabling last the lifetime of a tenancy.

This was a different time in networking. Token Ring, FDDI and ATM were still around and Ethernet was increasing from 10MB to 100MB. Category 5 hadn't been invented and no one knew which Layer 2 protocol would dominate the future.

In 2015, I wonder how many businesses actually survive 20 years? How many business stay in the same building for 20 years? How many business keep the same floor layout, decoration or office fittest for 20 years? Five years ago, building leases were 5 years with first option on 5 more. Today, it

is common for building leases to be 3 + 3.

What about routers, firewalls, servers and arrays? Bandwidth & storage is doubling every 2 to 3 years. Firewalls are now "Next Generation" and the next generation of firewalls after that will have no way to break encryption.

Today, budgets and business plans are measured in quarters with a long-term horizon of "financial year".

Yet I talk to IT architects and their strategies are planned in decades, with asset lifetimes of 5 to 10 years. In networking, enterprises buy backbone switches in the data centre and campus with expectations of 10 years of ownership.

That's a HUGE gap in expectations and planning between business and IT.

Cost Implications

If a business operates in quarterly cycles and has a financial horizon of 12 months, then at best, a 36-month strategy can be justified as "long term."

There is a substantial cost implication in purchasing a product for 10 years of ownership. Like that 20-year guarantee I sold on cabling plant in 1996, there was a significant cost to producing high quality products that would last. Plus the "up market" pitch generated far higher profit margins for manufacturer and reseller alike. My rule of thumb is that high quality costs 200% more CapEx than "good enough for now." My rule of thumb for "10 year survivability" is 300% in CapEx and 400% for OpEx.

Is that really worth it?

Buying What You Need Now

Logically, IT assets should have a 3-year lifecycle. Selection criteria must include the ability to cheaply, rapidly and easily replace the physical equipment. For software, selection criteria should include data portability so that your word processor, banking application or human resources platform can be replaced at any time.

My Second Law of Infrastructure (Law of Design Choices)
2. Good, fast, cheap: choose only two.
2a. Anything free is worth what you pay for it.

I'm coming to the view that the last twenty years of IT has been dominated by Good & Fast. Our attempts to make things cheap was achieved by reducing our value to the business.

The next 10 years is going to be "Fast and Cheap". We will spend the next ten years attempting to deliver value to the business with rapid asset rotation and fast-moving software and there's not much good about that.

Postscript: and Virtual Appliances

Virtual Appliances are products that meet the need for short, fast lifecycles. They don't need dedicated hardware, cabling or power. They are easier to replace and upgrade. For some people, subscription licensing means you can walk away from them at any time.

Just make sure that you avoid fancy functions or vendor-specific features when you deploy them.

Augury: SFP Modules, DRM And Per-Port Licensing

I've been tracking a number of rumours that I haven't been able to prove.

First, I heard that Extreme Networks had implemented a DRM scheme on SFP modules in certain models of its switches. If you purchase and install an SFP

unit from Extreme then operation is normal. Using third-party or OEM SFP will cause a licensing message to appear on console indicating that a per-port license is needed to continue using the interfaces at full speed. A license must be purchased and configured within 90 days or the port speed will reduce to 25%. That is, a 10GbE will operate at 2.5Gbps, 40GbE will operate at 10Gbe and so on.

Second, I heard that Arista Networks has implemented full DRM on all SFP ports in certain switch models. Any use of OEM SFPs in the switch causes the port to completely disable itself and throw an error message. I understand that there is no workaround; you must buy Arista SFP modules.

Should I Be Worried

Again, I haven't been able to confirm these stories but it presents a worrying trend.

1. Buying SFP modules from a vendor at substantial markup is a per-port licensing fee. I doubt that most customers have realised this yet.
2. In speaking with third-party OEM companies, there is no technical reason to use vendor-certified modules.
3. When purchasing a modern switch, more than 50% of the total deployment cost is in the SFP modules when all or most ports on switch are in use. This cost is sometimes hidden when adding just a a few ports at a time. (How many people actually price the full cost of a switch including SFPs, cables, power etc)
4. Tracking the spare parts of certified SFP modules is a complete nightmare.

If you are forced, through DRM, to use vendor-certified SFP modules what are your thoughts ?

Guest Writer: The Push Button Data Center

I am freshly returned from an event where I heard a number of data center related vendors discuss their wares. Nearly all of them were startups. Startup tech companies are interesting in the data center space, as they are often begun by founders who have lived through IT pain, then try to assuage that pain. Pain relief is a risky business. The pain has to be poignant enough to drive demand for the product they create. The pain also has to be long-standing enough that it will still exist by the time their product comes to market.

A common thread I noticed in presentations was the idea that, "We make your life easier." In the case of the modern data center, I don't think startups are taking big risks by trying to reduce operational complexity. Data center management is so challenging that products making the data center easier to consume will gain customers.

The companies I heard from included Rubrik, VMTurbo, Data Gravity, Pernix Data, and Scale Computing. As those are storage, backup, and virtualization related companies, you might wonder why I bring them up in a networking-centric magazine. Consider this. The data center is becoming highly automated. Converged and hyper-converged platforms have taken a small but significant share of the infrastructure market. New data center management tools are consumed by APIs. Vendors are partnering with one another to deliver unified solutions. There are several product options to manage storage, compute, and virtualization in tandem. The network is the last silo to fall. In speaking with what I'm loosely calling "automation vendors" about their approach to the network, the general answer is that the network is difficult to integrate, but on their minds

Networking is still hard to automate — still hard to consume. There's a lot of nuanced detail to networking, and it doesn't help that most of us build snowflakes. But snowflake network construction is not going to last forever. Conversations around NFV frameworks and policy are moving networking to a

place where it, too, will be able to be consumed in a predictable way by automation systems. Networking will have to become predictable, with fewer options and nerd knobs required to deliver an effective transport.

The network engineer of the future will not merely have deep knowledge of networking. They will also know a lot about automation. What network elements are being consumed? By what systems? In what way? And why isn't it working as expected? Knowledge of automation and comprehension of other IT silos coupled with deep networking knowledge will be the calling cards of the next generation network engineer. That will allow us to function effectively in the emerging push button data center.

Contributed by: Ethan Banks,
Packet Pushers Co-founder

Web: **ethancbanks.com**
Twitter: **@ethancbanks**

Drew-Conry Murray Contributor

Guest Writer: To Impress At Work, Try Talking

Most business communication happens via the written word, whether it's email, texts, documents, tweets, or messaging tools such as Slack or HipChat.

Well-crafted sentences and good spelling demonstrate a degree of intelligence, but **recent research** shows that if you really want to be regarded as thoughtful or intelligent, the human voice trumps written text.

Researchers at the University of Chicago **conducted studies** that measured how people reacted to pitches from candidates applying for hypothetical jobs.

The pitches were either spoken aloud, or read by the evaluators.

According to an article published in the journal Psychological Science "… evaluators rated a candidate as more competent, thoughtful, and intelligent when they heard a pitch rather than read it and, as a result, had a more favorable impression of the candidate and were more interested in hiring the candidate."

The study noted that the results were the same whether the pitches were given by trained actors or just regular people.

I prefer text when communicating, both as a sender and a receiver. As a sender, I can compose the message when it's convenient for me, and if I have time I can revise and rewrite.

As a receiver, I can quickly scan a piece of text and often get the gist much faster than waiting for someone to talk it out.

But it's also true the human voice carries nuances that don't always come across in writing, including passion, energy, sarcasm, and subtle emotional cues.

And as someone who creates and listens to podcasts, I know from my own experience that the human voice can resonate in ways that text alone can't.

So don't forget your voice as a business tool. Maybe the next time you need to pitch a new idea or train customers, don't send an email or a document.

Talk to people. Look them in the eye. Take advantage of the power of your voice. If this study is right, you might just come out looking smarter.

Drew Conry-Murray
Packet Pushers

Web: **andrewconry-murray.com**
Twitter: **@Drew_CM**

Got something to say? Send your contribution to ***humaninfrastructure@packetpushers.net*** *and, if we publish it, we will pay you a pittance for it.*

You Need Business Continuity, Not Business Connectivity

Most software-defined wide area network (SD-WAN) vendors today provide an extended LAN solution to protect and connect enterprise locations. We call this business connectivity. What an enterprise really needs is an SD-WAN solution that optimizes and ensures WAN availability. We call this *business continuity*.

For organizations with customer-facing systems (Web site, customer care, sales, etc.), any application downtime has an immediate impact on revenue and customer satisfaction. From an internal IT perspective, when critical applications and data are not available, employee productivity suffers a significant reduction. Customer and employee expectations mean that critical systems now require 100 percent uptime. The key to limiting the impact of system failure is to fix the problem fast by providing rapid (sub-second) failover to alternative data center resources. How organizations achieve that is the challenge.

Access to the continuously available data center architecture relies on a high-performance, highly reliable WAN. Advances in software-defined networking

(SDN) enable the network to support continuous availability via geographically distributed data centers. **This hands-on demonstration** illustrates how Sonus' business continuity solution is helping State Street Bank and other customers keep business up and running around the clock, even in the event of disaster. Ensure you know the differences between SD WAN business connectivity and SD WAN business continuity— and why *business continuity* is a must for your network. Learn more at **www.sonus.net**.

Q&A: Outspoken, Public And Career Damage

Question:

You seem like the no-holds barred kind of guy that is always willing to speak your mind even if your opinion is not popular. Do you ever feel like that has hurt you in your career?

Answer:

The person you hear on the podcasts is not the same person in real life. I think it is important to keep the show entertaining so that you listen, to keep it interesting so it isn't boring, and to challenge the audience on certain ideas in networking. Being a little crazy helps to keep it interesting and, maybe, fun.

When I'm am working for clients or companies, I am about listening,

researching, advising and covering the different options. Speaking your mind has a time and place, but it must be backed up by knowledge, research and references.

To prepare the podcasts and write blogs, I spend many hours researching technology, briefing with vendors, reading RFCs, manuals etc etc before finalising the format and discussion topics.

Has this hurt my career?

I work freelance so I don't speak for my employer. I've never worked for a vendor. I never mention my employers anywhere.

I would guess that I missed a few freelance gigs for those people who actually look you up on the Internet and get worried about that. But then, I have gained so much from being in public that its probably a fair trade.

> Got questions? Sure you do! Send them to **humaninfrastructure@packetpushers.net** and get answers (with no guarantee you'll like them).

Internets of Interest : Pages To See

Webpages that have caught my attention in the last couple of weeks.

Toxiproxy - CI Framework For Network Simulation

Stream your TCP requests through this proxy to simulate latency, delay and connection loss in your CI framework.

> Toxiproxy is a framework for simulating network conditions. It's made specifically to work in testing, CI and development environments, supporting deterministic tampering with connections, but with support for randomized chaos and customization. Toxiproxy is the tool you need to prove with tests that your application doesn't have single points of failure. We've been successfully using it in all development and test environments at Shopify since October, 2014.

LINK: **https://github.com/shopify/toxiproxy**

How Facebook Is Eating The $140 Billion Hardware Market

Whenever I look at the OpenCompute Project I get the feeling that custom server hardware for the data center, especially blade servers, is history.

> Facebook’s extraordinary Open Compute Project is doing for hardware what Linux, Android, and many other popular products did for software: making it free and "open source."

But revenues from blade servers and converged systems continues to grow and Enterprises continue to spend (waste?) big dollars on custom servers that runs commodity silicon....

> “I wrote a short paper, circulated it to Zuck and the rest of team,” he remembers, referring to Facebook CEO Mark Zuckerberg. Heiliger argued that the technology, particularly the hardware, “is not our competitive advantage." and that "open source should be a core tenet at Facebook.

For an Enterprise, IT hardware isn't a core function or competitive advantage. The real advantage is the software that runs and the design and operation of

that harware.

> Now that OCP has become a phenomenon, Google's top hardware-infrastructure guy (a legend in his world), Urs Hölzle, offers a begrudging respect for the project.

Google considers its hardware (network, storage and compute) as part of its magic powers and core business advantage. But Facebook is now gaining lower hardware costs by encouraging wider participation ..... which leads to volume manufacturing which leads to lower risks for setting up the production line and investing in producing Open Compute hardware.

> Flash forward to March 2015: Over 2,500 people came to US conference held at the San Jose Convention Center, OCP's new full-time CEO Corey Bell tells us.

Open Compute hardware is growing in popularity. I have spoken with many companies who are conducting early stage tests on building private clouds with the hardware. The universal fact is that unique or vendor branded hardware is expensive and not needed for a true cloud architecture. It doesn't necessarily repalce existing software but future applications are planned to run on full cloud-style systems.

Link: **How Facebook is eating the $140 billion hardware market**

How To Receive A Million Packets Per Second

This article talks about configuring Linux to handle 1 million UDP packets per second:

> That got me thinking. While I agree that 50kpps per core is probably the limit for any practical application, what is the Linux networking stack capable of? Let's rephrase that to make it more fun:

Most network engineers would immediately consider implementing a load balancer and tens of servers to handle the load. But with some Linux skills and serious Linux know-how (also, google-fu) you can handle a vast amount of packets....

First, IP receive queue is linked to single CPU core. Which is silly, there are always 4 or more CPU cores in a modern server (bare metal). ....

> To utilize multicore systems, NICs began to support multiple RX queues. The design is simple: each RX queue is pinned to a separate CPU, therefore, by delivering packets to all the RX queues a NIC can utilize all CPUs. But it raises a question: given a packet, how does the NIC decide to which RX queue to push it?
>
> Since the hashing algorithm on our NIC is pretty limited, the only way to distribute the packets across RX queues is to use many IP addresses. Here's how to send packets to different destination IPs:

OK, getting harder. Having multiple IP addresses needs some networking skills to understand what an IP address is and how to have more than one on a single server.

> While we had shown that it is technically possible to receive 1Mpps on a Linux machine, the application was not doing any actual processing of received packets - it didn't even look at the content of the traffic. Don't expect performance like that for any practical application without a lot more work.

OK, caveat emptor. Except it comes for free. Anything free is worth what you paid for it.

LINK: **How to receive a million packets per second**

State And Configuration

One of my favourite pastimes is watching ex-Cisco employees discover a world that doesn't have all-cisco-all-the-time-being-the-only-answer-to-everything

> Having spent my entire career with Cisco (working with IOS and NX-OS), until recently, I never realized there was a difference between the state of a system and the configuration applied to it.

ORLY?

> It’s important to understand that state and configuration are different ideas. In software like IOS the configuration is always the state. In software like IOS-XR or Linux you can have configuration somewhat independent from state allowing for syntax or dependency checking first, before applying state. Finally Linux can take it a step further and allow for state to be applied without configuration.

Congratulations, welcome to five years ago. Linux makes sense, but not for the reasons that you think.

LINK: **State and Configuration - Pete Lumbix - Excessive Redundancy**

Open Source Tools To Build An Open Source Lab

Mierdin (Matt Oswalt) provides a deep dive into using Open Source Tools to Build An Open Source Lab.

> In this post, I’ll be using VirtualBox, and also Ansible and Vagrant. For this purpose, I’m assuming you’re at least somewhat familiar with these tools.

I'm trying to be. I just can't set aside enough time but this post is a big help.....

> I hope you enjoyed this lab that I set up to help get your feet wet with open source routing! There’s much more where this came from - please don’t stop here, take this as far as it will let you. There are a lot of options involved that I didn’t have space to explore

Please do. In the meantime, go and read this post!!!

LINK: **Open Source Routing: Practical Lab**

44% Reduction In Page Load By Block Tracking Requests

This research paper from the Mozilla Corporation shows just how much bandwidth is consumed by cookies, tracking gifs, and other "personalisation" tools so that you can see "better ads".

> We present Tracking Protection in the Mozilla Firefox web browser. Tracking Protection is a new privacy technology to mitigate invasive tracking of users' online activity by blocking requests to tracking domains. We evaluate our approach and demonstrate a 67.5% reduction in the number of HTTP cookies set during a crawl of the Alexa top 200 news sites. Since Firefox does not download and render content from tracking domains, Tracking Protection also enjoys performance benefits of a 44% median reduction in page load time and 39% reduction in data usage in the Alexa top 200 news sites.

A 44% reduction in load time is truly significant. I have removed most of the tracking/analytics on EtherealMind for this reason.

LINK:**Tracking Protection in Firefox For Privacy and Performance**

Things On My Mind

On My Mind: Virtual Doughnuts

By the hairy balls of **Crom**, I love doughnuts. These are home made donuts that look more delicious than usual.

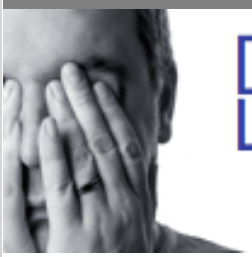

ETHEREAL MIND

Recent Articles

Here are the last five articles from the EtherealMind and Packet Pushers blog.

EtherealMind.com Latest

Logical Razors Can Take on Corporate Babble - **Link**
Canned Response to BGP Networking Questions – Reddit - **Link**
IETF RFC 8374 BGPsec Design Choices and Summary of Supporting

Discussions - **Link**
Net Neutrality Hasn't Ended, We Don't Know When - **Link**
Next Market Transition ? Cheaper Buying, Less Selling - **Link**

PacketPushers.net - The Last Five

Network Break 182: BGP Hijacked For Cryptocurrency Heist; Juniper, Big Switch Unveil New Products - **Link**
Show 387: AWS Networking – A View From The Inside - **Link**
PQ 147: Connecting Security And GDPR Compliance (Sponsored) - **Link**
Datanauts 131: Masters And Mentorship - **Link**
Network Break 181: Russia Accused Of Infrastructure Attacks; US Targets ZTE - **Link**

I welcome your feedback, questions, and corrections. Send an email to **humaninfrastructure@packetpushers.net** and we'll get things sorted out.

Still Here?

2CELLOS Luka Sulic and Stjepan Hauser playing their arrangement of The Trooper by Iron Maiden & William Tell Overture by Gioachino Rossini mashup.

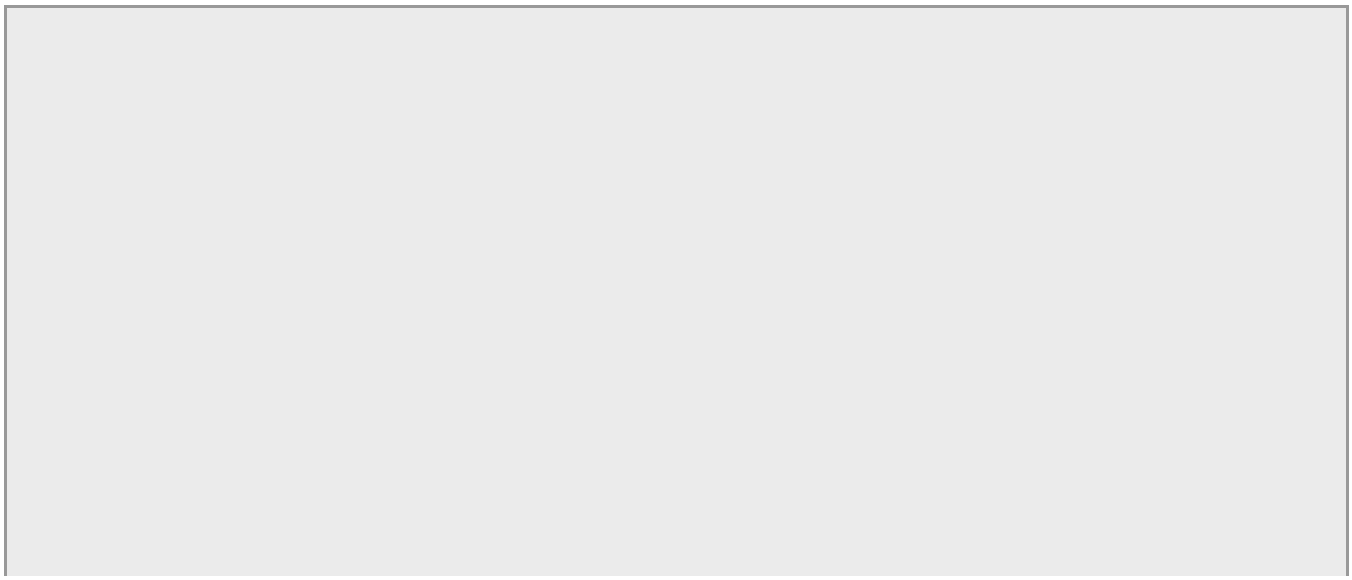

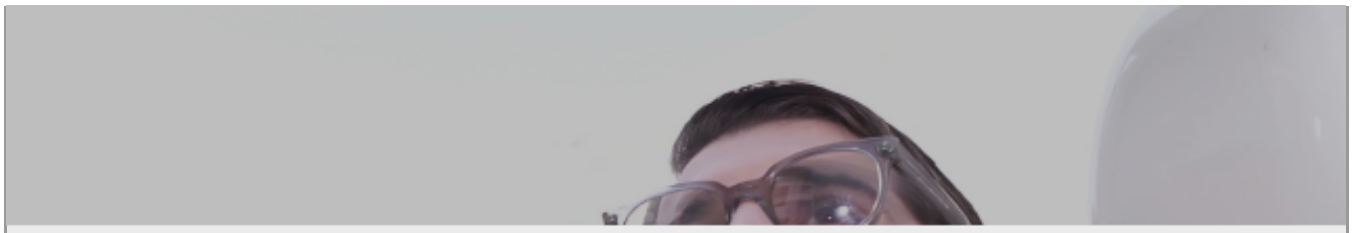

Iron Maiden and William Tell Overture. On Cellos? ..... HELL YES

Sponsorship and Advertising - Send an email to humaninfrastructure@packetpushers.net *for more information.*

About Piece of A Human Infrastructure

A strongly curated newsletter produced by Greg Ferro that contain observations and thoughts on IT Infrastructure with a networking focus that he has seen, done, discussed, reviewed or just simply found on the Internet.

The format is structured but flexible (like any magazine) and will change over time as I settle into a routine of capturing ideas, topics that seem relevant and ultimately finding patterns that seem to be interesting to you. Your feedback will drives changes so don't hesitate to email with feedback or ask followup questions for the next edition.

About Greg Ferro

Greg is a co-host of the Packet Pushers Podcast a weekly podcast on Data Networking which has over 8000 subscribers. He blogs regularly at EtherealMind.com for the last eight years and is pretty well known these days. He also write as an analyst for Network Computing and Gigaom Research. He speaks at major events on Data Centre Design, SDN and life in technology. He moderates panels, advises customers and technology companies.

He works as a part-time network engineer in the UK on a freelance basis. Because real work configuring routers and switches remain not only a passion but important to keeping touch with the industry.