Oct 2015 C b9238

Oct 2015 Cisco Network Architecture Geeks

Oct 2015 Cisco Network Architecture Geeks

Monday, May 22, 2017

7:10 AM

- MLBAM

o MPLS connecting all ball parks to central transmission operations centers in San Fran and NYC

o heavy wireless

o beacons - location based messages, promotions, ticket info based bluetooth sensors

o video — sd versus hd versus 4k not the interesting conversation

o voxel-based 3d imaging array (5k)

- ability to zoom in in 3d mode and look at plays
- decomposed and rendered onsite
- statcast doing real time statistical analysis of plays (how fast ball or player is moving)

o their OTT (over the top) streaming service

- they're now providing OTT services to HBO Now, WWE, PGA Tour as well
- going to do NHL next

o Observations:

- infrastructure used to be an asset, it is now a liability
- there will never again be a one platform world
- consumer choice will drive innovation much faster than Moore's Law
- millennial largest cohort (and largest percentage of workforce); larger than baby boomer

♣ 87% say phone never leaves their side

♣ 60% think everything will be mobile

♣ 14% won't do business that doesn't have a mobile site or app

♣ 67% don’t have credit cards

♣ 30% think a physical tv is important. 55% don’t intend to own one or care.

♣ only 17 out of 100 new households will bother getting traditional pay TV

- question about IPv6 — classic argument about how most of the ecosystem
- Imagine Communications — switching to a fragmented video delivery architecture

o expected that video delivery must be over Internet using Adaptive Bit Rate

o client controls the bit rate switching based on local conditions (cpu, available bandwidth, codec, etc.)

o friendly to do thru both CDN and NAT/firewalls

- Nick McKeown from Stanford — P4

o wants to fix the Internet, but things change too slowly

o Stanford Clean Slate Program - “what would we do if we started again?”

o lead to OpenFlow, NOX, SDN

o basic goal of programmability emerges from this

- note: abstraction (safely push a simple to code change at 1 layer down to something beneath it) is different than layering (change at 1 layer doesn’t change stuff beneath it)

o NOTE: last year, Cavium released a 3.2Tb programmable chip at the same time that Broadcom released their 3.2tb dedicated switched chip. Big deal!

- suspicion that all switching chips will be programmable in 5 years

o PISA — protocol independent switch architecture

- a generic approach to programmable chip development that doesn’t know anything about protocols — could drive much higher adoption of high performance programmable merchant silicon in the next few years

o P4

- high level language for programming these PISA chips
- for operators, demonstration of a p4 script that is directly querying the data forwarding plane to identify what is causing latency in real time
- what does it mean for equipment vendors:

♣ programmable in the field — continuous updates to data plane

♣ 1 box, many uses — switch router, firewall, load-balancer, etc.

♣ your ideas belong to you:

- you no longer need to tell your secrets to switch chip vendors

♣ 1 p4 program, many targets

- competition among chip vendors is to be best target, not best features
- what it does for operators, Input from ATT

♣ uniformity

♣ compliance testing

♣ IPR

♣ etc

- Xilinx, Barefoot, Netronome, Intel all going to be producing P4 compatible hardware

o Working w/ MIT to produce Domino a higher level language (c like) that could produce P4 code

o PERC — fast congestion control by proactive direct calculation of flow rates in the forwarding plane

o Goals

- new features — add new protocols
- red complexity — removed unused protocols
- greater visibility
- modularity
- efficient use of resources
- portability
- own your own IP
- Facebook on programmability (Najam Ahmad) — why we need it?

o lots of traffic, global footprint

o 3.2b cache operations every 30 min ; 5b messages

o 82% of users outside US

o network got so big, that they no longer can afford to troubleshoot network from bottom up

o have to drill down from heat maps down

o Major elements, in 30 days, major components in action:

- FBNet
- NetSonar
- ODS — gets the data
- NetNORAD
- Alert Manager engine
- Megazord == correlates alarms into ~1200 unique master alarms
- Emitter — receives and processes 3.37b network notifications, 0.99% result in an alarm
- FBAR — would run 750k times on networking alarms, automatically resolves 99.6% of them
- Carrier Maintenance — would act on roughly 300 maintenances
- Vendors — not that many escalations
- Drainer — reroutes traffic away from a failed device

o End Result — software mitigates most of the problems, engineers go and fix the failed devices that were taken out of rotation

- only 0.6% of alerts dealt w/ by a human

o Lack of programmability across many areas

- monitoring and correlation at this scale are genuinely difficult
- get ride of CLI as much as possible
- configurations of various devices come in proprietary formats, not leveraging server best practices

o redundancy — primary + backup doesn't cut it

- Facebook runs 16 wide in any given datacenter

♣ detect a failure and drain it, figure it out later

♣ horizontal instead of vertical

o keep everything as simple as possible

- they only run 2 protocols — BGP and ECMP (equal cost multi-path routing)

o don't perpetuate or create legacy

- they run a technical debt week every 6 weeks.

o they are furthest along in completing this automation strategy in their datacenter networks

- backbone not so much

o observation — network engineering in its non-programmable state is mostly systems integration these days

- he had a stated goal of converting his network engineering team into a software engineering team
- software engineers become part of the network team
- 1 network on call guy
- Sky presentation re: IPv6

o broadband provider in UK. 92% coverage across UK. About 20% of households

o they have broadband, bundled cable, and their own OTT service as well

o because they were a late entrant, didn't have a lot of IPv4

o trial runs of carrier-grade NAT demonstrated it was expensive and didn't work great

o decided to convert to IPv6

o IPv6 migration project — cost was less than 5% of capex (even w/ coattails projects)

o took 3 years to migrate

o 92% of users expected to be running IPv6 by April

o only CPE issues w/ their trial efforts were due to IP-enabled baby monitors!

o lots of

- Adam O'Donnell presentation on Security

o came from SourceFire acquisition

o off the shelf AV won't help against targeted attacks

- APT attacks are often custom developed, mutate quickly over time

o enterprise posture

- build best possible defense but assume breach will occur
- minimize amount of time to find out you're attacked
- after discovery, minimize amount of time the attack can propagate

o overview of FireAMP

- threat DB
- centralized behavior analysis
- sandbox testing capabilities
- quarantine actions can be implemented w/ having vendor push custom signatures

o can mitigate via endpoint policies and bump in the wire (snort IPS — either physical or virtual)

- Segment Routing Update from Cisco

o Across portfolio

- ISIS, OSPF, BGP, BGP-TE, BGP-LS, PCEP
- ipv4/6, mols,
- XR, XE, NXOS
- NCS, ASR, CRS, CSR, N9K, N3k
- ODL, WAE

o Assumption

- MPLS dataplane for IPv4 and IPv6
- ASR9K / Tomahawk

o Example of benefit

- doing a classic SP calculation of how to get to a destination w/ a single node failure, calculates a single alternative path
- w/ SR, you can also bring in multiple ECMP paths into the revised path calculation
- Sam Ramji, CEO of Cloud Foundry

o Internet All the Things

o continuous innovation now more important than locking up a market thru 1 time competitive advantage

o Cloud foundry is a joint foundation between Pivotal, IBM, HP

- focused on node.js, cloud, open container
- Deepfield Presentation
- Cisco and RouteViews on BGP

o BMP — BGP Monitoring Protocol

- in Juniper and XR/XE Cisco routers
- going thru IETF

o OpenBMP is a monitoring web interface for pulling in data from BGPMon devices

- OpenDNS

o sign up, point all of your devices at their DNS service

o easy to provide dns-based security anywhere

#learning/conferences