

The IETF Network

... an overview




I'm going on an adventure!


IETF Meeting Router Statistics

To view these stats go to <http://dashboard.meeting.ietf.org>


Downstream Bandwidth



Total Downstream Data




Total Upstream Data




Upstream Bandwidth



MetroOptic



Telus



The Scout™

- A Ubiquiti router
- Shipped to site / installed during site visit
- Starts announcing our address space
 - Allows testing of the circuits
 - Validation of the BGP peering, etc.
 - Provides an anchor for geo-location data
 - Gets the **ietf-hotel** SSID up for NOC

Circuits

- At least 2, but up to 5 circuits
- Almost always donated by local providers
 - Try for redundant:
 - providers
 - fiber
 - entrances
- 1Gbps -> 10Gbps
- Dual stack (IPv4 / IPv6)
- BGP

Routers

- 2 Juniper routers
 - Were MX80s, upgraded to MX204s (this meeting)
 - Convergence: ~25 minutes -> ~1.5 minutes
- Core routers for network
 - BGP (eBGP, iBGP)
 - RPKI
 - OSPF / OSPFv3
 - DHCP relay / RA
 - BCP38
 - Passive ARP learning (more if we have time...)

ARP ARP ARP... ARP ARP...




ARP ARP ARP... ARP ARP...

```
aggregate {  
    inactive: route 130.129.0.0/16;  
    route 31.133.128.0/18;  
    route 31.130.224.0/20;  
}
```


ARP ARP ARP... ARP ARP...

```
aggregate {  
    inactive: route 130.129.0.0/16;  
    route 31.133.128.0/18;  
    route 31.130.224.0/20;  
}
```



RPKI

```
routing-options {  
    validation {  
        group rpki-servers {  
            session 31.130.229.4 { # Dragon Research Labs RPKI Toolkit  
                preference 100;  
                port 323;  
            }  
        }  
    }  
}  
  
policy-statement RPKI {  
    term whitelist {...}  
    term invalid {  
        from {  
            protocol bgp;  
            validation-database invalid;  
        }  
        then {  
            validation-state invalid;  
            community add RPKI_Invalid;  
            reject;  
        }  
    }  
}
```

```
policy-statement RPKI {
  term whitelist {
    from {
      protocol bgp;
      prefix-list RPKI_Whitelist;
    }
    then {
      validation-state valid;
      community add RPKI_Whitelist;
      next policy;
    }
  }
  term invalid {
    from {
      protocol bgp;
      validation-database invalid;
    }
    then {
      validation-state invalid;
      community add RPKI_Invalid;
      reject;
    }
  }
  term valid {
    from {
      protocol bgp;
      validation-database valid;
    }
    then {
      validation-state valid;
      community add RPKI_Valid;
      next policy;
    }
  }
}
```

```
term unknown {
    from {
        protocol bgp;
        validation-database unknown;
    }
    then {
        validation-state unknown;
        community add RPKI_Unknown;
        next policy;
    }
}
/* This should not happen -- things should be valid,
invalid or unknown */
term failed {
    from protocol bgp;
    then {
        community add RPKI_Failure;
        next policy;
    }
}
}
```

Switches

- 2 x Cisco Catalyst 4500X Core stacked
- 10 x Cisco IDF switches
- 40 x Cisco 12 port switches
- "Joe's magic..."
 - Y'all keep plugging in DHCP servers :-(
 - A new switch to a fully provisioned switch in ~15 minutes (including a software upgrade).
 - Rooms are dynamic - this means we need to reconfigure things often and quickly

Switch Automation

- Feature-wise, the switch automation includes:
 - Initialize new switch with desired config and software image
 - Validation of config and image (checksum)
 - Auto-generation of SSH host key
 - Call-home for when a switch should re-ZTP
 - Auto-detection of connected device type (switch, AP, probe)
 - Port auto-config and auto-doc update
 - Detection of lost device and port description update

NEW: Device sw-122 made a request to bootstrap

Serial Number

FOC2129Y3X5

Current Version

15.2(6)E2

DHCP IP Address

31.130.224.239

Platform ID

WS-C3560CX-12PD-S

Current Image File

c3560cx-universalk9-mz.152-6.E2.bin

NEW: Device sw-123 made a request to bootstrap

Serial Number

FOC2129Y3X4

Current Version

15.2(6)E2

DHCP IP Address

31.130.224.240

Platform ID

WS-C3560CX-12PD-S

Current Image File

c3560cx-universalk9-mz.152-6.E2.bin



ztp APP 1:55 PM



VERIFY SUCCESS: Device sw-120 has been successfully bootstrapped

Serial Number

FOC2129Y3X8

VERIFY SUCCESS: Device sw-121 has been successfully bootstrapped

Serial Number

































































FOC2129Y3X6

IETF Switch Registration Tool: Physical Switches

Show Only:

Switch Search:

[Add Physical Switch](#) | [Export To Ansible](#) | [Logical Switches](#)

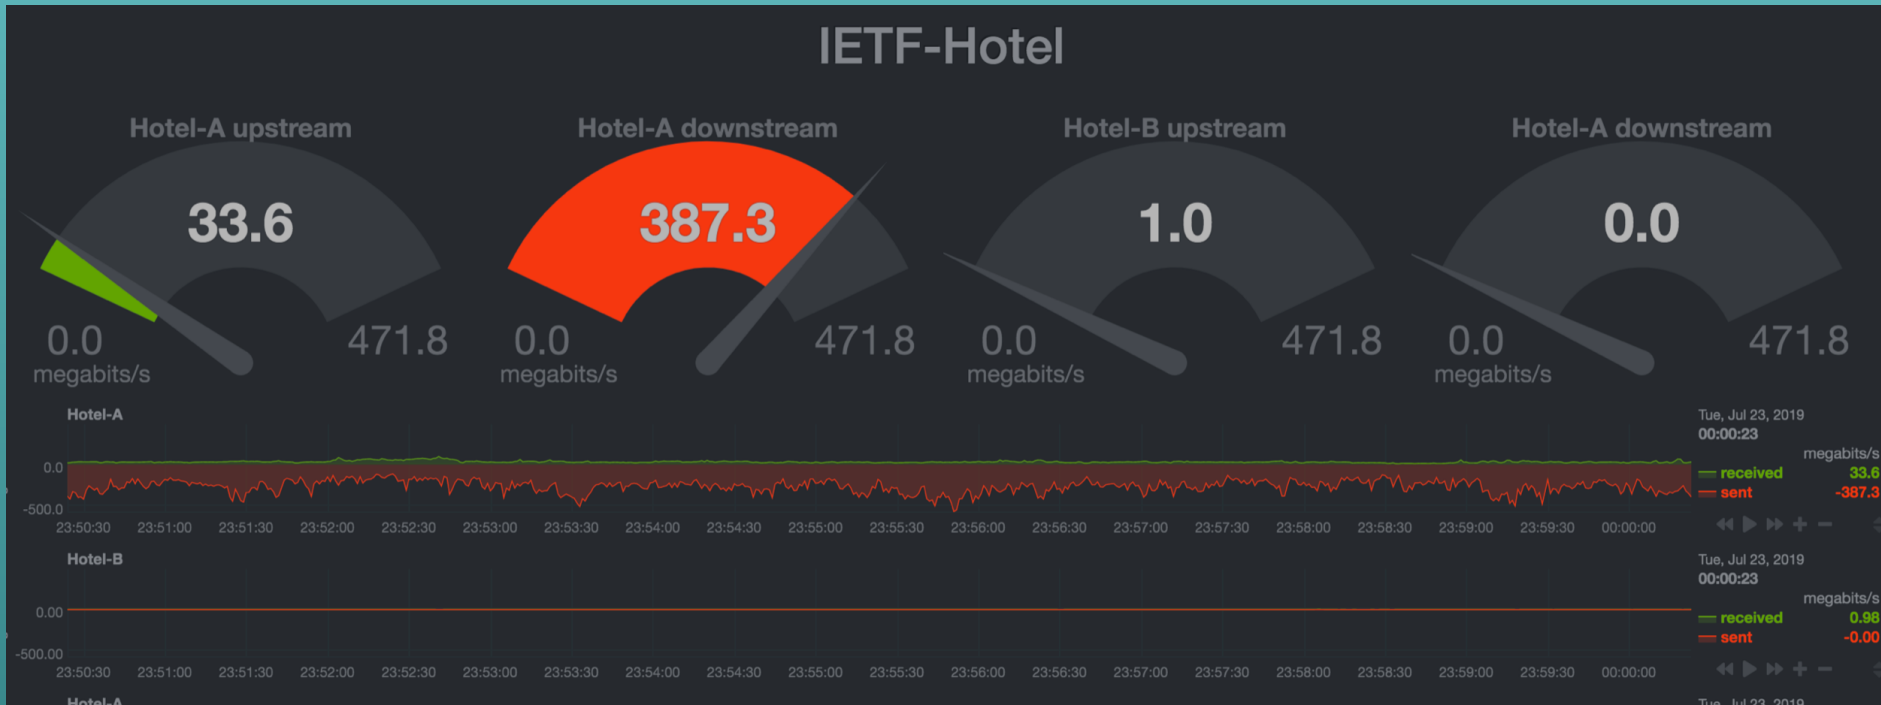
Search: <input type="text"/>									
Row No. ▲	Delete? ▼	Assigned? ▼	Re-ZTP? <input type="checkbox"/> ▼	Serial Number ▼	Product ID ▼	Max Ports ▼	Assigned Logical Switch ▼	Provision Status ▼	Reachability ▼
1.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	FCW2132C063	WS-C3850-24U-L	24	sw-101   		
2.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	FCW2132C08E	WS-C3850-24U-L	24	sw-102   		
3.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	FCW2132C08H	WS-C3850-24U-L	24	sw-103   		
4.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	FCW2132C08N	WS-C3850-24U-L	24	sw-104   		
5.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	FCW2132D02S	WS-C3850-24U-L	24	sw-105   		
6.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	FCW2132D07Y	WS-C3850-24U-L	24	sw-106   		
7.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	FOC2129Y3VL	WS-C3560CX-12PD-S	12	sw-111   		
8.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	FOC2129Y3VN	WS-C3560CX-12PD-S	12	sw-112   		
9.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	FOC2129Y3WL	WS-C3560CX-12PD-S	12	sw-113   		
10.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	FOC2129Y3WP	WS-C3560CX-12PD-S	12	sw-114   		
11.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	FOC2129Y3WZ	WS-C3560CX-12PD-S	12	sw-115   		
12.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	FOC2129Y3X1	WS-C3560CX-12PD-S	12	sw-125   		
13.	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	FOC2129Y3X2	WS-C3560CX-12PD-S	12	sw-124   		

Showing 1 to 50 of 50 entries

Wireless

- 2 x Cisco WLC 5520 in an HA pair
 - Cisco WLC 2504 for ISOC & testing
- Somewhere between 50 and 70 Access Points
 - [TODO] 55 this time
 - We do both 5Ghz and 2.4Ghz, prefer 5Ghz
- This has largely solved much of the ARP problem
 - Does your phone battery now last >3/4 day?
 - Thank Panda...!
- Multiple **encrypted** SSIDs
 - "ietf-legacy, ietf, ietf-2.4only, ietf-nat64, ietf-v6only, ietf-nat64-unencrypted, eduroam, isoc, ..."

Guestroom / "hotel"



Guestroom Network

IETF participants are "weird"..
... no, really weird...

- Guest networks are built for *normal* people
 - Captive portal
 - Intercept / rewrite DNS
 - HTTP munging...
 - NAT
 - Drop no-good, bad, dangerous ports (like 22!)
 - Assumptions:
 - Limited devices
 - Limited bandwidth
 - Limited sessions
- IPv6? Ain't nobody got time for that...

From recent stay

```
wkumari$ git push  
ssh: connect to host  
git.kumari.net port 22:  
Connection refused  
fatal: Could not read from remote  
repository.
```

Please make sure you have the
correct access rights
and the repository exists.

```
wkumari$
```

~~Guestroom network~~

- Bypass guestroom gateway with Ubiquiti routers, open SSID
 - "Free Internets for all!"
- Some hotels have ~~truly bizarre~~ inventive architectures...
 - Really bad channelizing
 - Mac Mini in "Internet Sharing Mode"
 - Access Points on elevators... much hilarity...
- Too few access points in guest rooms (getting better)
- Ethernet over Coax / DOCSIS / DSL / Cat3
- Integrated PoS, TV, mini-bar, signs, thermostats, ...

Servers / Services

- 3+3 Physical servers
 - Ganeti, Docker
-
- DNS / DNSSEC, DPRIVE
 - DNS64
 - DHCP / DHCPv6
 - NTP
 - Tickets
 - RPKI server
 - TACACS+ / RADIUS
 - ZTP server
 - Etherpad
 - Ansible for automation
(Yay! DevOps!)
 - SMTP
- Git repo
 - VMs for Meetecho
 - Backups
 - Syslog
 - Monitoring:
 - Prometheus
 - Deadman
 - Intermapper
 - Smokeping
 - Rancid
 - Netdisco
 - Observium, ...

Scrubbing PII....

Remote Participation

- Live streaming gets their own VLANs
- ... and VMs
- ~60 Mbps BW from VMs to Internet
- The network we build makes remote participation possible
- Meetecho / Kaskadian have done events on venue networks
 - but only streaming (not remote participants)
 - Meetecho remote participation depends on "but the limited bandwidth, NATs, firewalls, lack of IPv6, would likely prevent us from providing good remote participation."
 - Kaskadian: "Hotel network won't work!.... :-P"




You deserve a kitten now...

Experiments...


DPRIVE

- Ran one of the early DNS-over-TLS services
- Now it is a "standard service"

IPv4



IPv6



V6ONLY – no, really.

- Turned off IPv4 on all radios near V6OPS, 6MAN
 - Hilarity ensues... :-P

NAT64 Testing

	Meetecho	Jabber	Etherpad	Skype	Signal	Spotify	Outlook	Dropbox	Air Display
MacOS	Adium								
iOS							TBT	TBT	
Android									
Windows		TBT		TBT	TBT		TBT		
Web based									

Ref: <https://datatracker.ietf.org/meeting/100/materials/slides-100-v6ops-sessa-readout-from-ietf100-hackathon-01>


MAC Randomization



Universidad
Carlos III de Madrid
www.uc3m.es

DHCP Logs

- ◆ **144 local MACs seen during the week (IETF92)**
- ◆ **97 IP addresses were assigned to local MAC addresses. Out of them:**
 - ❖ **76 IP addresses were assigned to one local MAC address, e.g., because no DHCP client identifier was used by the client**
 - ❖ **21 IP addresses were assigned to multiple local MAC address**



Questions?

