NETWORK RESEARCH CENTER, TSINGHUA UNIVERSITY

Design, Build and Operate IETF79 Meeting Network

An internal report

Xiaoliang ZHAO
Zhiyan ZHENG
Jilong WANG
Xing LI
{*xzhao, zhzhy, wjl, xing@cernet.edu.cn*}

12/01/2010

This report is to document how IETF79 meeting network was setup from scratch and how it was operated. The report is intended for internal use, but there are no restrictions to distribute the material for information sharing purpose.

Revision History:

Date	Version	Author	Note
2010.12.21	1.0	xzhao	Initial draft
2010.12.31	2.0	xzhao	First readable version

Table of Contents

1 INTRODUCTION 3

2 NETWORK REQUIREMENTS 3

3 DESIGN 4

3.1 ADDRESS SPACE 5
3.2 EXTERNAL CONNECTIVITY 6
3.3 EQUIPMENT 9
3.4 TEAM 10
3.5 DEPLOYMENT BLUEPRINTS 11
3.6 REDUNDANCY CONSIDERATIONS 14

4 BUILD 15

4.1 CIRCUITS 15
4.2 MDF ROOM 15
4.3 IDF ROOM 19
4.4 MEETING ROOM 20
4.5 GUEST ROOM 22

5 OPERATION 24

5.1 NETWORK MONITORING 24
5.2 ISSUE REPORTING AND HANDLING 25

6 TEARING DOWN 26

7 APPENDIX 27

APPENDIX A: IETF MEETING NETWORK REQUIREMENTS 27

1 Introduction

IETF 79th meeting, *IETF79* for short, was held at Shangri-la hotel in Beijing from Nov. 7 to Nov. 12, 2010. More than 1000 people from all over the world attended the meeting, and as usual, most of attendees were actively using the Internet during this period. It is the meeting host's responsibility to build a capable network, therein referred as *IETF meeting network*, to provide network coverage for all IETF attendees. This report documents our experiences to design, to build and to operate the IETF79 meeting network.

The document starts with the network requirements mandated by IETF, which guides the network design. Then we describe in detail on how the network was built and operated. One important note is the whole process was rather dynamic and fluid, while the report only recorded it in a static picture.

Below is a quick summary of some key points in this report:

1. The 1-2-3-4 of IETF79 meeting network:
   a) ***One*** network.
   b) ***Two*** design goals: working and reliable.
   c) ***Three*** essentials: circuit, equipment and people.
   d) ***Four*** phases: design, build, operation and tearing-down.
2. 50 seats per AP seems a good rule of thumb to estimate the total number of wireless APs.
3. Information availability and information sharing are very important when different groups work together during busy hours. Spreadsheet seems a good tool. For example, below spreadsheet can be used to track where a particular AP is located and connected:

AP #	Room #	Patch Panel #	Switch Port #
AP101	Pearl	D315	vw1-sw 0/20

4. Guest room coverage is almost as important as meeting room coverage regarding to user experiences.
5. Be open and flexible, changes will happen, surprises may come.

2 Network Requirements

The IETF Meeting Network Requirements document is the guideline as well as the inputs to the network design and the network operation. The latest version can be found online at: http://iaoc.ietf.org/network_requirements.html. For this report to be self-contained, a copy is also included in Appendix A.

For a quick summary, a "minimal" IETF meeting network must have:

1. Two redundant circuits providing external connectivity with minimal 45Mbps bandwidth, or 100Mbps recommended bandwidth.
2. BGP peering with network providers to advertise IETF's PI (Provider Independent) address space.
3. IPv6 support.
4. No content filtering, no limiting firewall, no NATs.
5. Full 802.11b coverage for all meeting rooms and common areas.
6. A terminal room with about 100-150 Ethernet drops.
7. A manned helpdesk.
8. Redundant DHCPv4 servers.
9. Redundant DNS servers.
10. A SMTP relay server.
11. Network monitoring and data collection system.
12. Two network-connected enterprise-class printers.
13. Enough power strips.

3 Design

There are two key inputs to network design. One is the IETF network requirement document, the other is the good understanding on the meeting venue, including how many meeting rooms, how many Ethernet drops and power outlets in each room, how each room covered by wired line and wireless, how the guest rooms covered, and so on. The more understanding on the meeting venue facility and its network infrastructure, the less last-minute changes (or surprises) of the design.

We decided building a *working* network is our first design goal. It means the network will be designed to provide necessary functionalities only, instead of a cutting-edge or feature-rich or showing-off network. To achieve such goal, probably the best and safest approach is the "copy-n-paste" approach to replicate previous IETF meeting network settings. We collected and studied available information on previous meetings, and also actively participated IETF78 meeting network setup and operation at Maastricht which provided us a great opportunity to learn and to practice.

Our second design goal is to build a *reliable* network. To achieve this goal, the design emphasized on redundancy and reduction of complexity. We follow the Simplicity Principle as stated in RFC1958: "*Keep it simple. When in doubt during design, choose the simplest solution.*"

The following subsections will detail the design considerations on address space, circuits for external connectivity, network equipment, team building. The outcome of the design phase is a detailed deployment blueprint.

3.1 Address Space

IETF meetings usually use dedicated address space, i.e., 130.129.0.0/16 for IPv4, 2001:df8::/32 for IPv6, and AS290 for BGP. The address space is belonging to Interop Show Network, which may use the same address for different show events. It is a good idea to double check with IETF before announcing the address to the Internet.

Internally, the IPv4 and IPv6 address space is divided into several smaller blocks for different uses. The below spreadsheet was used to manage the address allocation for IETF79.

Network	IPv4 Network	IPv6 Network	VLAN ID	SSID
Loopbacks	130.129.0.x/32	2001:df8::x/128	N/A	N/A
Management	130.129.1/24	2001:df8:0:1::/64	1	manage
Media	130.129.3/24	2001:df8:0:3::/64	3	N/A
Peering	130.129.4/24	2001:df8:0:4::/64	N/A	N/A
Server	130.129.5/24	2001:df8:0:5::/64	5	N/A
NAT64-downlink		2001:df8:0:7::/64	7	ietf-nat64
Wireless-v6ONLY	N/A	2001:df8:0:8::/64	8	ietf-v6ONLY
Wireless-Secure	130.129.32/21	2001:df8:0:32::/64	32	ietf.1x
Wired	130.129.48/21	2001:df8:0:48::/64	48	N/A
Wireless-Guest-Room-Inside	130.129.64/21	2001:df8:0:64::/64	1164	ietf-hotel
Wireless-Guest-Room-Outside	N/A	N/A	4001	N/A
Wireless-Guest-Room-Secure	130.129.72/21	2001:df8:0:72::/64	72	ietf-hotel.1x
Wireless-A-Secure	130.129.96/21	2001:df8:0:96::/64	96	ietf-a.1x
Wireless-Portal-Inside	130.129.112/21	2001:df8:0:112::/64	1112	ietf-portal
Wireless-Portal-Outside	N/A	N/A	112	N/A
Wireless-A-Portal-Inside	130.129.128/21	2001:df8:0:128::/64	1128	ietf-a-portal
Wireless-A-Portal-Outside	N/A	N/A	128	N/A
IVI-uplink	130.129.46/24	2001:df8:0:46::/48	46	N/A

The most of items in the spreadsheet are self-explanatory. But it may need some explanation on Portal related items. IETF79 authenticates users through web portal system and 802.1x. Web portal system uses two VLANs. One ("outside VLAN") is used in the direction from user to web portal system, and the other ("inside VLAN") is used in the direction from web portal system to the border router. After a user is authenticated, the portal system will do a VLAN ID swap to change the "outside VLAN ID" to "Inside VLAN ID". In other words, the traffic will be logically taken out from "outside VLAN" then put into "inside VLAN" where the border router is logically connected. If authentication fails, there is no such VLAN ID swap, and the traffic won't be able to reach the border router.

3.2 External Connectivity

There is an exchange point located inside Tsinghua University which connects several Research and Education networks, including CERNET (China Education and Research Network) and CSTNET (China Science and Technology Network). CERNET and CSTNET are quite large networks connecting almost all Chinese universities and research institutes, as well as the rest of the Internet through other ISPs. We found a local dark-fiber rental company who can provide layer-1 connectivity between Shangri-La hotel and Tsinghua University. At Tsinghua's exchange point, we then connected to CERNET and CSTNET as our upstream providers. Both connections were designed to run at 1Gbps rate.

Besides physical connectivity, the traffic engineering and routing policy at border needed to be sorted out. Because we had plenty bandwidth, to keep the design simple, we chose 1) using prefix deaggregation as primary traffic control mechanism for incoming traffic; 2) using local preference to control outgoing traffic; 3) no load balancing. In other words, the CERNET circuit would act as the primary link for both IPv4 and IPv6 traffic, while the CSTNET circuit would remain idle unless CERNET circuit went down. Figure 1 illustrates such design. The drawback for this design is the recovering time might be high. During our switchover test between two circuits, we found it took about 3 minutes for both v4 and v6 traffic to recover after circuit being switched.

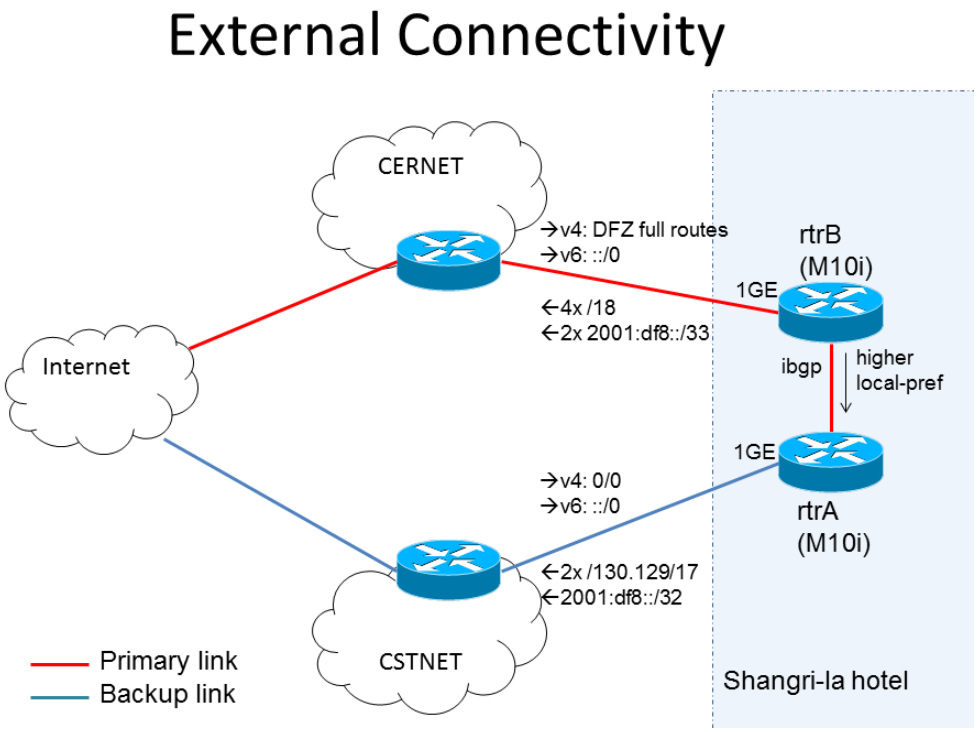

Figure 1: External Connectivity

Based on our design, BGP configuration at border routers are quite straightforward. We accept full routes or default route from CERNET and CSTNET, and we announce more specifics to CERNET but aggregated route to CSTNET. Juniper Networks Inc. generously loaned us two M10i to be border routers. One connected to CERNET was named RTR-B and the other connected to CSTNET was named RTR-A. Two border routers ran iBGP between themselves and RTR-B was configured with a higher local preference. The partial BGP configuration on RTR-B is shown below.

```
protocol {
    bgp {
        group ebgpv4-cernet {
            type external;
            import ebgpv4-cernet-im-policy;
            export ebgpv4-cernet-ex-policy;
            peer-as <CERNET ASN>;
            neighbor <CERNET v4 address>;
        }
        group ebgpv6-cernet {
            type external;
            import ebgpv6-cernet-im-policy;
            family inet6 {
                unicast;
            }
            export ebgpv6-cernet-ex-policy;
            peer-as <CERNET ASN>;
            neighbor <CERNET v6 address>;
        }
    }
}
policy-options {
    policy-statement ebgpv4-cernet-ex-policy {
        term self-routes {
            from {
                route-filter 130.129.0.0/16 prefix-length-range /16-/18;
            }
            then accept;
        }
        term others {
            then reject;
        }
    }
    policy-statement ebgpv6-cernet-ex-policy {
```

```
        term Self {
            from {
                route-filter 2001:df8::/32 prefix-length-range /32-/34;
            }
            then accept;
        }
        term Others {
            then reject;
        }
    }
    policy-statement ebgpv4-cernet-im-policy {
        term deny-self {
            from {
                route-filter 130.129.0.0/16 orlonger;
                prefix-list UNALLOCATED;
                prefix-list MARTIANS;
            }
            then reject;
        }
        term set-local-pref {
            from {
                protocol bgp;
                neighbor <CERNET v4 address>;
            }
            then {
                local-preference 120;
                accept;
            }
        }
    }
    policy-statement ebgpv6-cernet-im-policy {
        term Self {
            from {
                route-filter ::/0 exact;
            }
            then {
                local-preference 120;
                accept;
            }
        }
        term Others {
```

```
            then reject;
        }
    }
}
```

3.3 Equipment

Mainly there are five different kinds of equipment to run an IETF meeting network, including

1. layer-3 routers;
2. layer-2 switches;
3. wireless Access Points (AP)
4. servers
5. printers.

We collected equipment from different sources. Routers were loaned from Juniper. Switches and APs were provided by IETF. Servers and printers were bought from local market. Below is a detailed list of all equipment used for IETF79 network

Type	Model	Specifications	Quantity	Notes
Layer-3 Router	Juniper M10i	- Dual RE - Dual PS - 1x 4-port GE PIC	2	Used as border router.
Layer-2 Switch	Cisco WS-C3750G-48PS-S	- 48 Ethernet 10/100/1000 ports - 4 SFP-based Gigabit Ethernet ports	4	Used for core switch, and terminal room switch
Layer-2 Switch	Cisco WS-C3560E-24PD-S	- 24 10/100/1000 PoE ports + 2 X2-based 10 Gigabit Ethernet ports - allowing > 15.4W to all 24 ports	8	Powering up wireless APs
Layer-2 Switch	Cisco WS-C3560-8PC-S	- 8 Ethernet 10/100 ports and 1 dual-purpose 10/100/1000 and SFP port	12	For NOC, registration desk, and other small subnets.
Wireless AP	Cisco AIR-AP1252G-A-K9	- a/b/g/n support - 2.4G and 5G	74	
Server	Dell R410	- Intel 2.4G CPU - 8G RAM - 300G SAS HDD	6	To run network services and monitoring

		- 4x GE NIC		system.
Printer	HP ColorCP2025dn HP 2035n	- Network printer - Color & Black	2	

For software systems, IETF provided us with two FreeBSD VM images. One was built with DNS, DHCP Radius and other network services, and the other was built with Cacti, Netdisco, Email spam control and other network management applications. The VM images are properly configured and maintained by IETF volunteers and they have been used by many previous IETF meetings. Once the network is up, it is quite straightforward to run the VM images to provide needed software services. One exception is the web portal system, which is the Captivator software from U. of Wisconsin. It is not part of VM images and we had to build and configure it separately.

3.4 Team

We divided all engineers and volunteers into 4 small groups, namely, wireline, wireless, system, and helpdesk group. Each group included one or two dedicated engineers and a number of volunteers.

Group	Responsibility	Head counts
Wireline	- Configure and maintain all routers and switches - Lay down cables	- 2 Tsinghua engineers - 4 student volunteers
Wireless	- Configure and maintain all wireless Aps - Lay down cables connecting to APs	- 3 Tsinghua engineers - 10 student volunteers - 5 Cisco volunteers
System	- Install and maintain VM images - Configure and maintain DNS / DHCP / RADIUS / SMTP / NTP services - Configure and maintain user authentication system - Configure and operate Cacti / Netdisco network management system - Configure and operate network security system	- 6+ IETF volunteers - 2 Tsinghua engineers
Helpdesk	- Run helpdesk from 8am to 8pm - Manage ticketing system	- 1 Tsinghua staff - 1 IETF volunteer - 5 student volunteers

Although most of our engineers had adequate experiences and training, we were still lack of experiences of running an IETF meeting network. To gain as much experiences as possible, we decided to build a small but fully functional testbed with the same equipment for meeting. IETF agreed to ship all its equipment (switches and APs) a month earlier to us. The shipment arrived on Sep. 29 and then our team worked on the testbed for weeks. There were at least three things

achieved by this mini-IETF network. First, we found a routing policy issue with service providers, which was eventually resolved but it took quite a while. The testbed provided us enough time to handle such issue. Second, all services such as DNS and DHCP were installed and tested, which also saved us some time. Third, our engineers gained valuable experiences.

Most student volunteers came to help when we set up the network on Nov. 6. During network operation, they took rotations mainly helping with the helpdesk and monitoring network status at night.

3.5 Deployment Blueprints

To build IETF meeting network inside Shangri-La hotel, we need understand its facility and network infrastructure better. The hotel's IT department was quite collaborative and helped us with a lot of information and assistance. Shangri-La hotel has two main buildings, Garden Wing and Valley Wing, connecting with each other. Valley Wing was built around 2007 and it is quite network-friendly with adequate Ethernet drops and power outlets in most of rooms. Comparatively, Garden Wing, which was initially built in 1987 and renovated significantly in 2004, has a rather old network infrastructure. The difference between two Wings mainly affected where we put PoE switches. For Valley Wing, PoE switches were put in the IDF (Intermediate Distribution Frame) room to co-locate with the patch panel, which means a lot less work of re-wiring. Garden Wing had one specific problem: from its IDF room, some Ethernet cables were more than 100 meters long to the wall ports. It is then impossible to put PoE switch in the IDF room and use it to power up APs, because of too much power loss and poor signal quality over such long distance. But the good news was only few meeting rooms were in Garden Wing and there were optical fibers available at major locations. We took the advantage of those fibers and put PoE switches in the Ball room and terminal room. The network deployment blueprints were shown below:

Network Topology

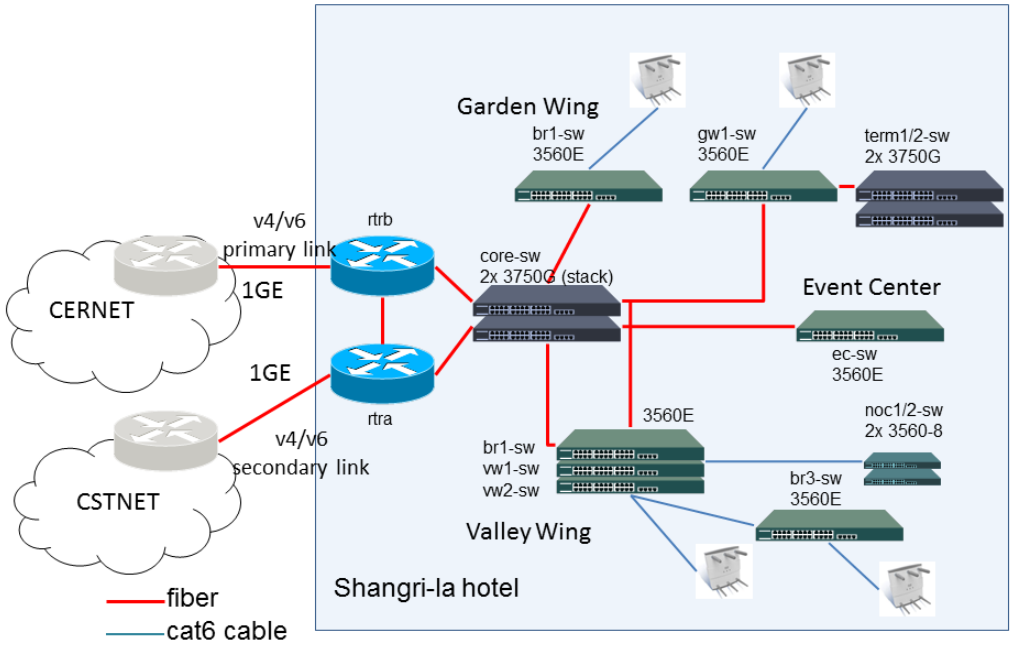

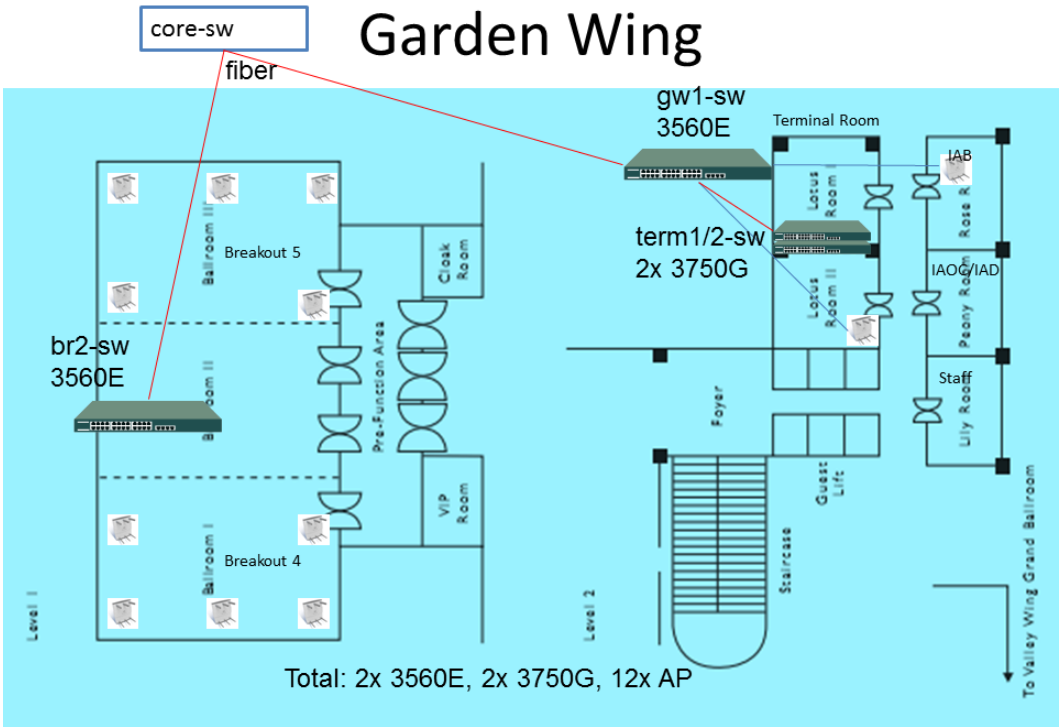

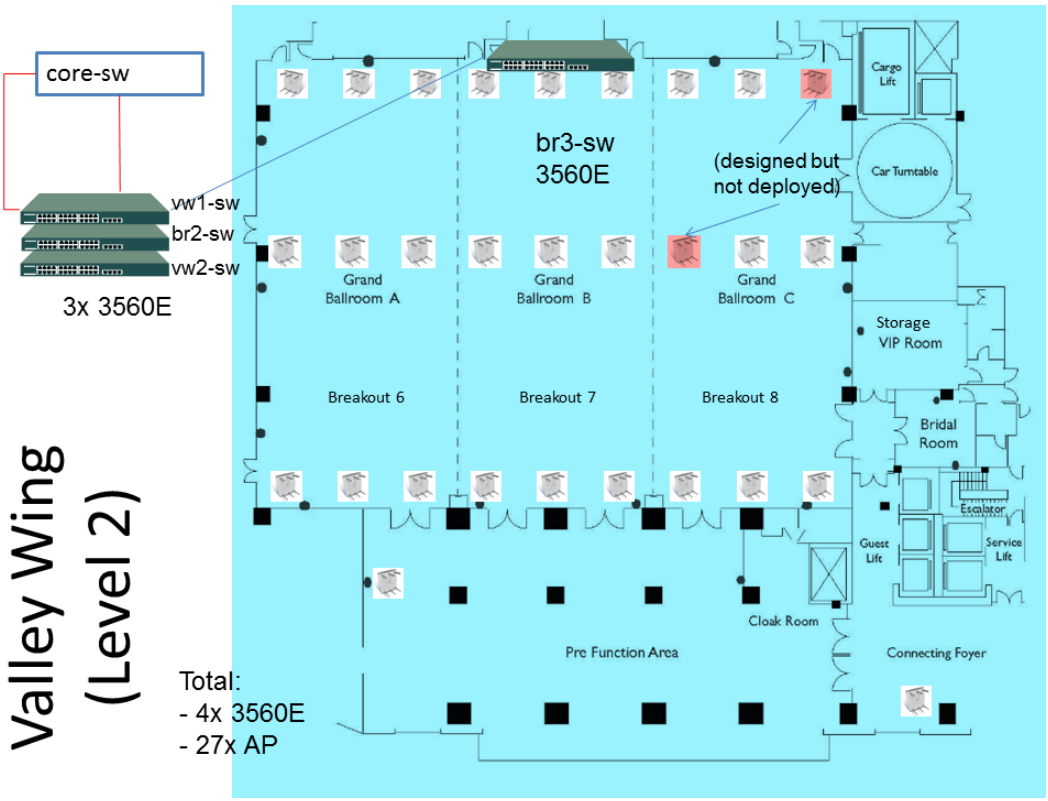
core-sw
vw1-sw
br2-sw
vw2-sw
3x 3560E
br3-sw
3560E
(designed but
not deployed)
Cargo Lift
Car Turntable
Grand Ballroom A
Grand Ballroom B
Grand Ballroom C
Storage VIP Room
Breakout 6
Breakout 7
Breakout 8
Bridal Room
Escalator
Guest Lift
Service Lift
Cloak Room
Pre Function Area
Connecting Foyer
Valley Wing (Level 2)
Total:
- 4x 3560E
- 27x AP

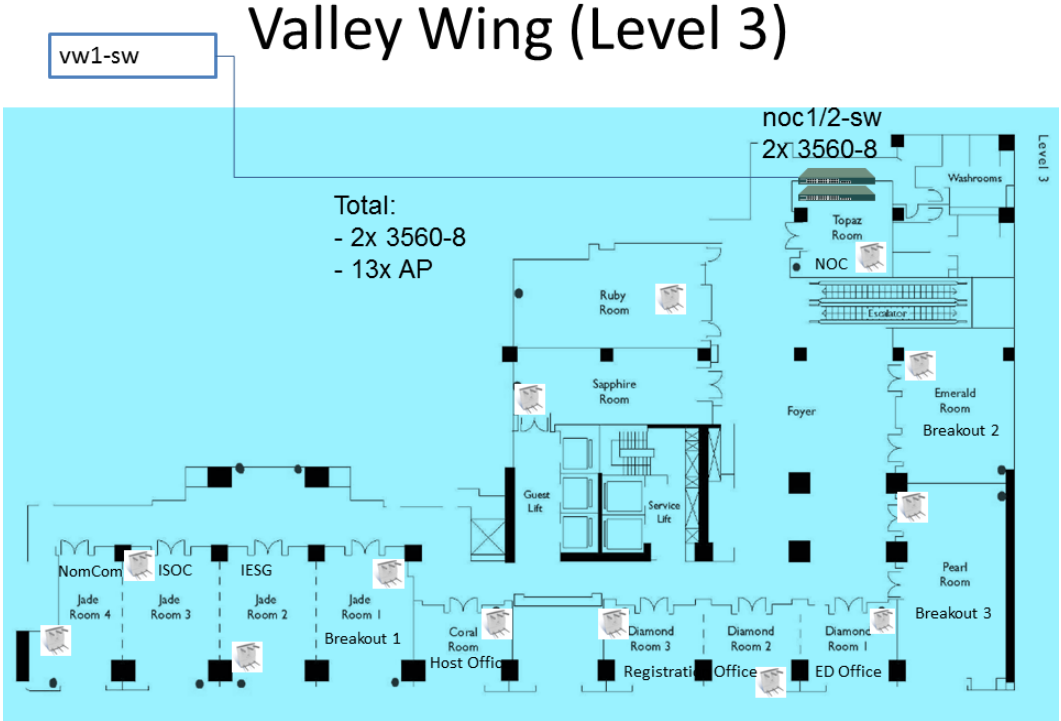
Valley Wing (Level 3)
vw1-sw
noc1/2-sw
2x 3560-8
Total:
- 2x 3560-8
- 13x AP
Washrooms
Level 3
Topaz Room
NOC
Escalator
Ruby Room
Sapphire Room
Foyer
Emerald Room
Breakout 2
Guest Lift
Service Lift
NomCom
ISOC
IESG
Jade Room 4
Jade Room 3
Jade Room 2
Jade Room 1
Breakout 1
Coral Room
Host Office
Diamond Room 3
Diamond Room 2
Diamond Room 1
Registration Office
ED Office
Pearl Room
Breakout 3

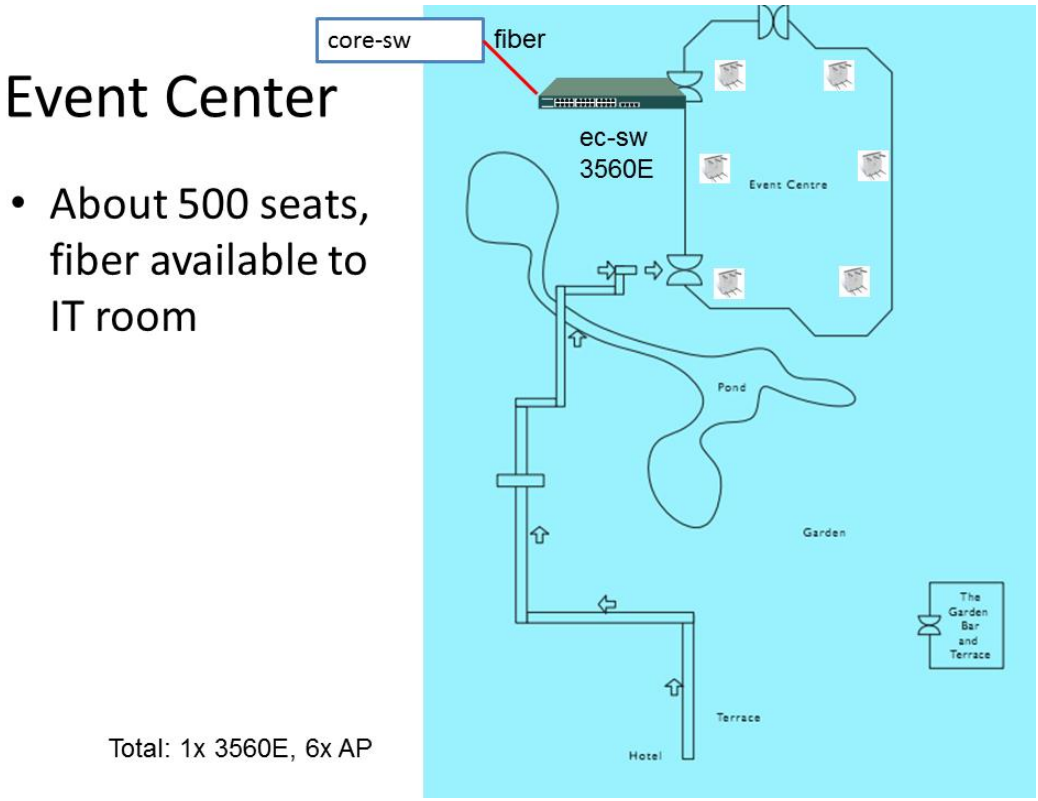

3.6 Redundancy Considerations

The external connectivity was designed with as much redundancy as possible. First, the design had two circuits for link redundancy. We rented two dark fibers and they were on non-overlapping routes to minimize the impacts of fiber cut. Second, two border routers could back up each other. Each router had two power supplies and they were fed by UPS-protected power sources.

The critical services such as DNS and DHCP were also completely redundant. First, they were part of the same VM image and two copies of the image were running simultaneously on two different physical servers. All the servers again were protected by UPS. The two DNS services were configured with different IP addresses and both were available to end users. The two DHCP services were protected by the failover feature from ISC dhcpd.

The rest of the network, mainly switches and APs, was only protected by closely monitoring the network so failure hardware could be quickly identified and replaced.

4 Build

4.1 Circuits

We started working with on circuits back to June. The fiber rental company had a telecom well outside the hotel, and the problem was how to get fiber from the well into the IT room, *i.e.*, our MDF (Main Distribution Frame) room. The Hotel was very cooperative and only asking the work should be good looking. We laid the fiber around the wall and lifted it up to cross a street. Fiber work was completed three week earlier than IETF meeting.

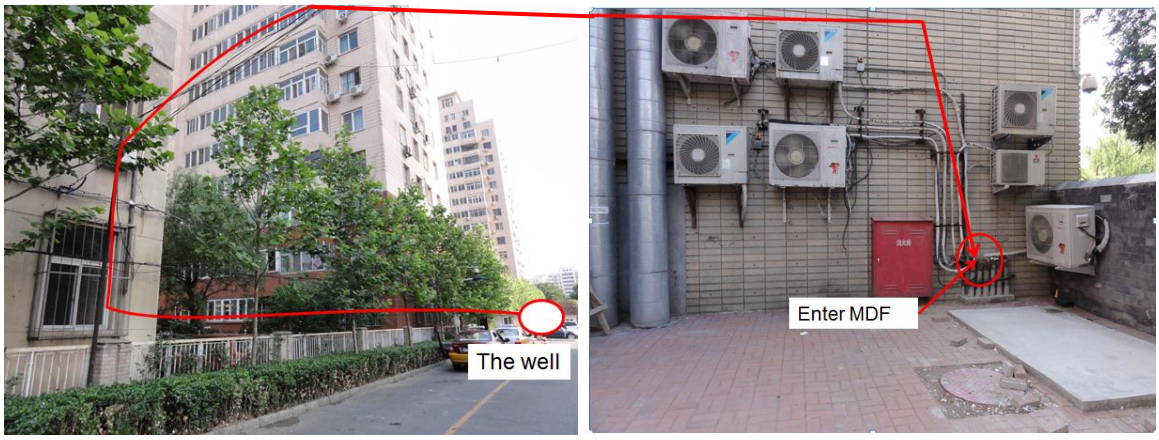

Figure 2: fiber from telecom well to MDF

4.2 MDF Room

Two weeks before the meeting, from Monday, Oct. 25, we started working full time to build the network at hotel. At that time, none of the meeting rooms available to us, so we started with MDF room, *i.e.*, the IT room. Two M10i was delivered to the hotel and the interface went up/up without any problems. With the help from CERNET engineers and others, BGP session also came up quickly. We started announcing IETF address blocks to the world on Oct. 28. BGP session with CSTNET was also established a week later.

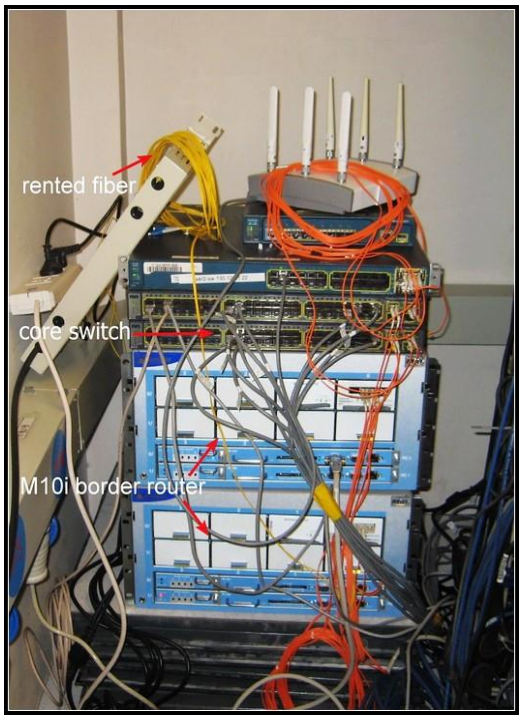

Figure 3: Border Router in MDF room

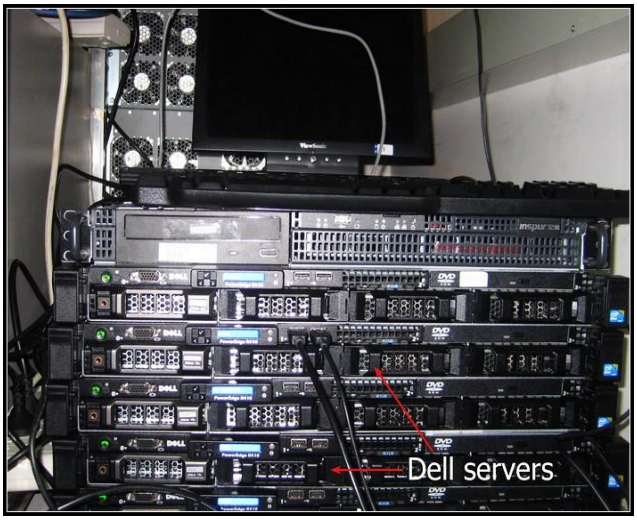

Figure 4: Servers in MDF room

All servers were moved from the staged testbed to here. The VM images were already installed and tested during the testing. So we only need turn them and connect them to the core switch. Cables were probably labeled, which was a minor detail and it took a while to finish but later it turned out to be quite helpful for troubleshooting. An excel file was used to print all cable labels.

A end	Z end	label	
core switch g1/0/47	service VM eth0	CS-g1/0/47 Service-eth0	CS-g1/0/47 Service-eth0

The core switch was configured as a layer-2 central point to connect border routers, downstream switches and servers. Below configuration exerts show a typical setting on core switch. First, the interface configuration for uplinks connecting to border routers:

```
! uplink configuration
!
interface GigabitEthernet1/0/49
 description "Link to RtrA"
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 4050
 switchport mode trunk
 snmp trap mac-notification added
```

```
 snmp trap mac-notification removed
 spanning-tree portfast trunk
!
interface GigabitEthernet2/0/52
 description "Link to ge-0/1/0.RtrB"
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 4000
 switchport mode trunk
 snmp trap mac-notification added
 snmp trap mac-notification removed
 spanning-tree portfast trunk
 !
```

Second, the interface configuration for downlinks connecting to other switches:

```
!
! Connecting to downstream switches with port channel
!
interface Port-channel2
 description "to ValleyWing1-Switch"
 switchport trunk encapsulation dot1q
 switchport mode trunk
 logging event bundle-status
 snmp trap mac-notification added
 snmp trap mac-notification removed
 spanning-tree portfast trunk
!
interface GigabitEthernet1/0/51
 description "to vw1-sw g0/25"
 switchport trunk encapsulation dot1q
 switchport mode trunk
 ip access-group killDHCPservers in
 snmp trap mac-notification added
 snmp trap mac-notification removed
 channel-protocol lacp
 channel-group 2 mode active
 spanning-tree portfast trunk
!
interface GigabitEthernet2/0/51
 description "to vw1-sw g0/26"
 switchport trunk encapsulation dot1q
```

```
 switchport mode trunk
 ip access-group killDHCPservers in
 snmp trap mac-notification added
 snmp trap mac-notification removed
 channel-protocol lacp
 channel-group 2 mode active
 spanning-tree portfast trunk
!
! Connecting to downstream switches without port channel
!
interface GigabitEthernet1/0/50
 description "to gw1-sw g 0/25"
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 4000
 switchport mode trunk
 snmp trap mac-notification added
 snmp trap mac-notification removed
!
```

Last, the interface configuration for downlinks connecting to servers:

```
! Connecting to servers or wired end hosts
!
interface GigabitEthernet1/0/47
 description vmhost-1
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 4000
 switchport trunk allowed vlan 1,5
 switchport mode trunk
 snmp trap mac-notification added
 snmp trap mac-notification removed
 spanning-tree portfast trunk
!
interface GigabitEthernet2/0/6
 description "registration"
 switchport access vlan 48
 switchport mode access
 ip access-group killDHCPservers in
 snmp trap mac-notification added
 snmp trap mac-notification removed
```

```
 spanning-tree portfast
!
```

4.3 IDF Room

From Tuesday, Nov. 2, we started laying down cables and deploying more switches and wireless APs. Since the external connectivity and all servers were already up, it made easy to check if a newly deployed switch or AP worked or not. The IDF room is the place where all the cables to meeting rooms are physically aggregated here on the patch cable. To take over a meeting room to make it connect to IETF network, we pulled the hotel's cable from the patch panel then plugged in our cables connecting to IETF switches. The below pictures illustrate such process:

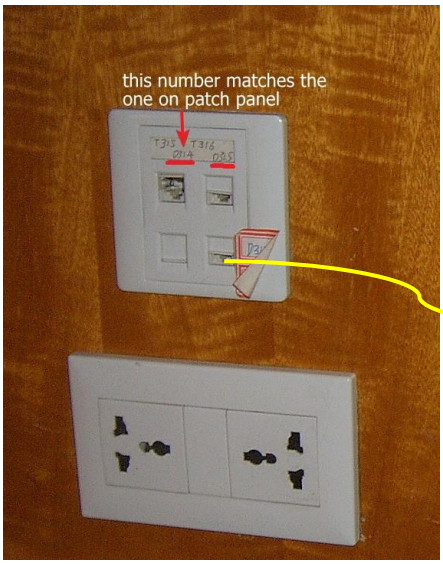

Figure 5: Ethernet ports in a meeting room

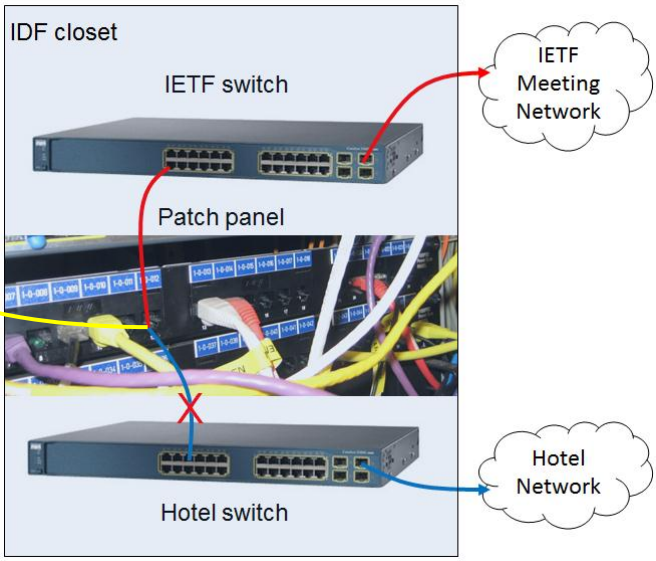

Figure 6: re-wiring work at IDF room

With the help from network engineers from the hotel, the re-cabling work at IDF room took just couple of hours. On late afternoon, Nov. 5, all meeting rooms were wired to the IETF network. Below is a typical interface configuration on a Cisco 3560E switch for wireless APs:

```
/* A distribution layer switch */

! a typical config for AP trunk
!
interface GigabitEthernet0/2
 description "AP Trunk"
 power inline port maximum 20000
 switchport trunk encapsulation dot1q
```

```
 switchport trunk allowed vlan 1,7,8,16,32,80,96,112,128,144
 switchport mode trunk
 ip access-group killDHCPservers in
 ipv6 traffic-filter v6_Access_IN in
 snmp trap mac-notification change added
 snmp trap mac-notification change removed
 spanning-tree portfast trunk
!
! firewall filter to drop any DHCP server replies from clients
!
ip access-list extended killDHCPservers
 deny udp any eq bootps any log
 permit ip any any
!
! firewall filter to drop any IPv6 RA advertisements from clients
!
ipv6 access-list v6_Access_IN
 deny icmp any any router-advertisement
 permit ipv6 any any
!
```

In above configuration, there are two firewall filters configured for network security purpose. The filter *killDHCPservers* will drop any incoming DHCP server reply messages (UDP source port 67) from wireless clients. It makes sure no wireless clients sending out false DHCP replies which may damage the whole network. Similarly, the filter *v6_Access-IN* drops any incoming IPv6 Router Advertisement (RA) messages from wireless clients to make sure no wireless clients falsely acting as IPv6 routers. All switch ports connecting to APs as well as the switches in the terminal room were configured with the two filters to prevent unintentional errors.

4.4 Meeting Room

Two days before IETF meeting, on Nov. 6, we were granted accesses to all IETF meeting rooms. The major work left was deploying wireless APs, wiring the terminal room and the event center. Many student volunteers came to help so we could finish all the work in one day.

We adopted the standalone APs, instead of controller-based approach. The only reason was because IETF'78 used standalone APs and the result was pretty good. The controller-based approach could be a good choice too because it becomes more and more popular now. Deployment of standalone APs were straightforward. When an AP was deployed at the spot, it was manually reset so that the AP would load its configuration from the DHCP server. In this way, the configuration management of 60 more APs was much easier.

In a big room, all APs will start broadcasting signals and try to tune them up automatically so they work on different channels with least interferences with each other.

A typical DHCP configuration for an AP is shown below:

```
host ap101 {
    next-server 130.129.5.12;
    hardware ethernet 00:1E:7A:81:3B:76;
    fixed-address 130.129.1.104;
    option bootfile-name 1252startup.ini;
    option host-name "ap101";
}
```

AP configuration example for SSID ietf.1x:

```
dot11 vlan-name Wireless-Secure vlan 32
dot11 ssid ietf.1x
   vlan 32
   max-associations 50
   authentication open eap eap_methods
   authentication key-management wpa
   mbssid guest-mode
!
! IMPORTANT BITS BELOW
power inline negotiation prestandard source
!
interface Dot11Radio0
 encryption vlan 32 mode ciphers aes-ccm tkip
 ssid ietf.1x
 speed  basic-11.0 6.0 9.0 12.0 18.0 24.0 36.0 48.0 54.0
 channel least-congested 2412 2437 2462
!
interface Dot11Radio0.32
 encapsulation dot1Q 32
 no ip route-cache
 bridge-group 32
 bridge-group 32 subscriber-loop-control
 bridge-group 32 block-unknown-source
 no bridge-group 32 source-learning
 no bridge-group 32 unicast-flooding
```

```
 bridge-group 32 spanning-disabled
!
```

For power strips, we hired a local company which had proper licenses to install all power strips for us.

4.5 Guest Room

The IETF network coverage to guest rooms was not explicitly required, but it was very important for attendees. The hotel had wireless coverage to all guest rooms but the traffic would go through hotel's network instead of ours. We worked closely with hotel's IT department to add two SSIDs into their wireless controllers. User traffic was then switched to the IETF network through VLANs. Two Ethernet cables connected hotel's switch with IETF switch. One was used for users' traffic, and the other was used to connect the wireless controller to our Radius server directly. We found if the wireless controller and the radius server were located in different address spaces, the radius traffic actually went out to the public Internet first then came back to the IETF network, even both devices were co-located in the same room. A direct cable made data transfer more efficient and secure for Radius traffic.

The below Visio drawing reflected the final network layout in more detail.

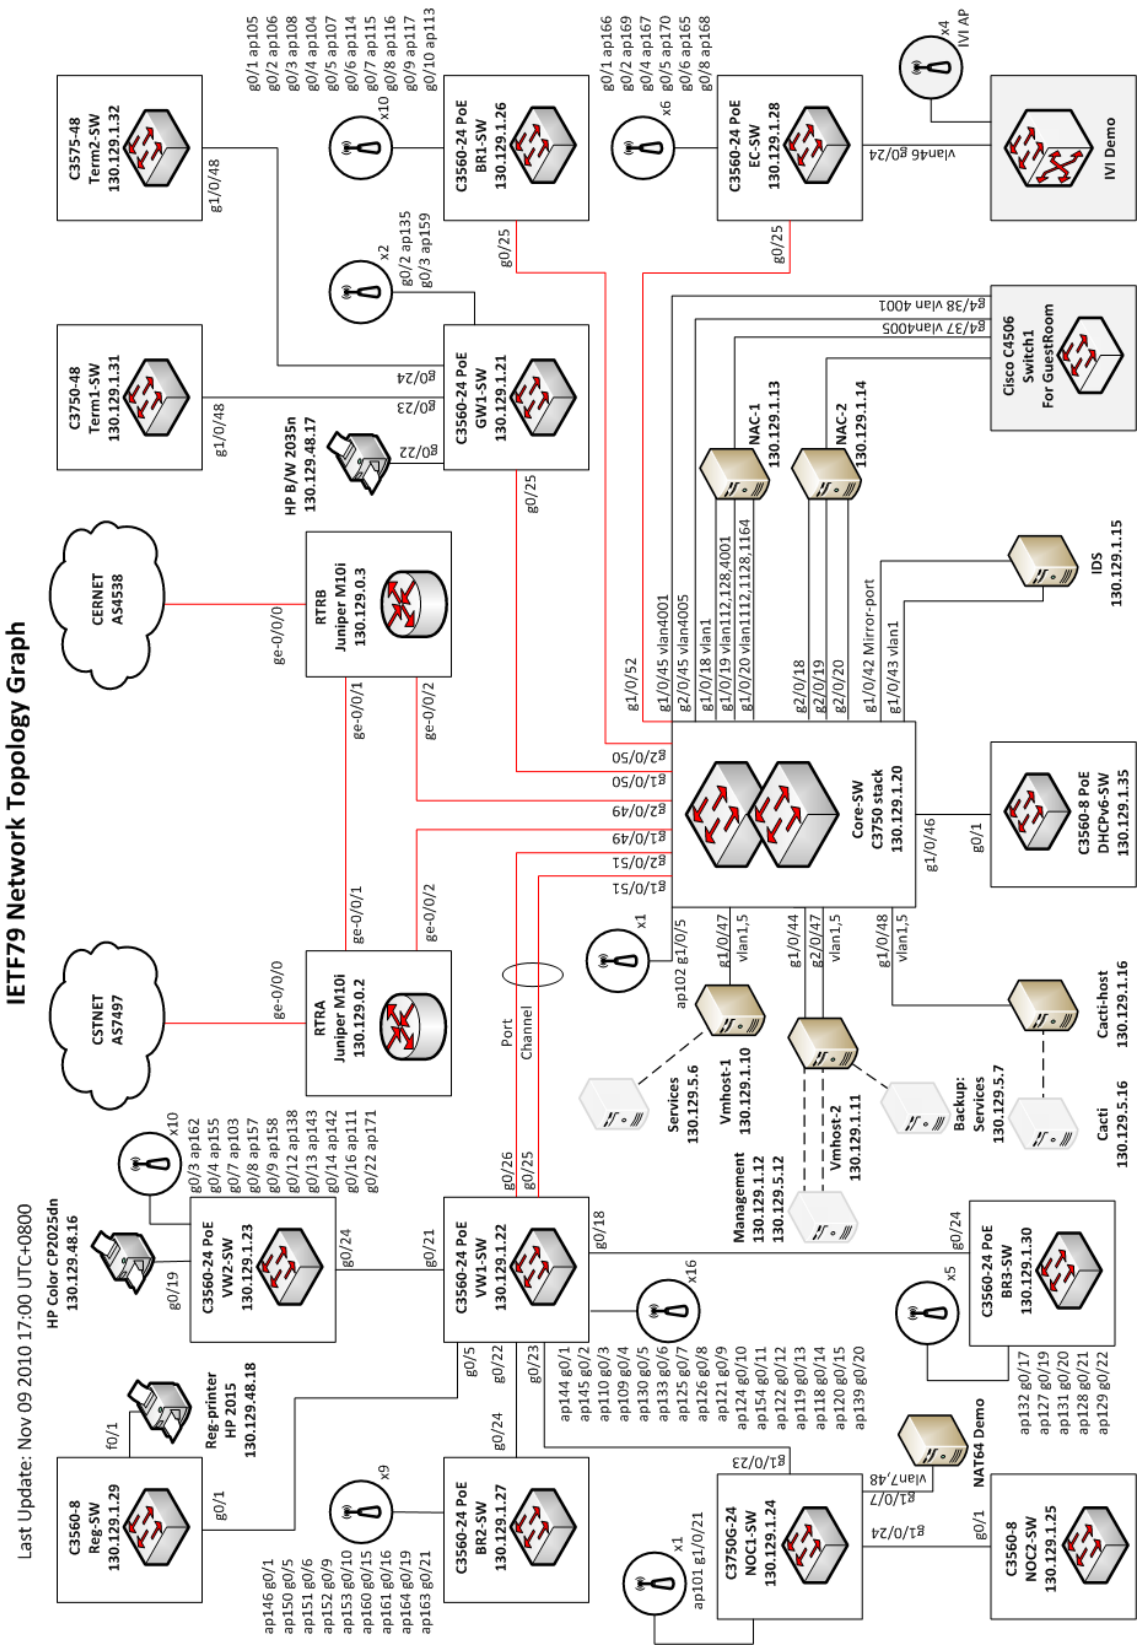
IETF79 Network Topology Graph
Last Update: Nov 09 2010 17:00 UTC+0800
CERNET AS4538
CSTNET AS7497
RTRB Juniper M10i 130.129.0.3
RTRA Juniper M10i 130.129.0.2
HP B/W 2035n 130.129.48.17
HP Color CP2025dn 130.129.48.16
Reg-printer HP 2015 130.129.48.18
C3575-48 Term2-SW 130.129.1.32
C3750-48 Term1-SW 130.129.1.31
C3560-24 PoE BR1-SW 130.129.1.26
C3560-24 PoE GW1-SW 130.129.1.21
C3560-24 PoE EC-SW 130.129.1.28
Cisco C4506 Switch1 For GuestRoom
IVI Demo
NAC-1 130.129.1.13
NAC-2 130.129.1.14
IDS 130.129.1.15
Core-SW C3750 stack 130.129.1.20
C3560-8 PoE DHCPv6-SW 130.129.1.35
Vmhost-1 130.129.1.10
Vmhost-2 130.129.1.11
Services 130.129.5.6
Backup: Services 130.129.5.7
Cacti-host 130.129.1.16
Cacti 130.129.5.16
Management 130.129.1.12 130.129.5.12
C3560-24 PoE VW2-SW 130.129.1.23
C3560-24 PoE VW1-SW 130.129.1.22
C3560-24 PoE BR3-SW 130.129.1.30
C3560-24 PoE BR2-SW 130.129.1.27
C3560-8 Reg-SW 130.129.1.29
C3750G-24 NOC1-SW 130.129.1.24
C3560-8 NOC2-SW 130.129.1.25
NAT64 Demo
IVI AP
Port Channel

5 Operation

Starting from Sunday morning, Nov. 7, IETF79 meeting network entered production mode[1]. To keep the network stable, major changes to the network were not allowed unless very necessary. The majority work shifted to the network monitoring and the helpdesk.

5.1 Network Monitoring

We kept ping every network device to make sure their aliveness and we kept eyes on the traffic pattern changes. Both were made visible through graphs created by Cacti tool[2]. Some screenshots were shown below.

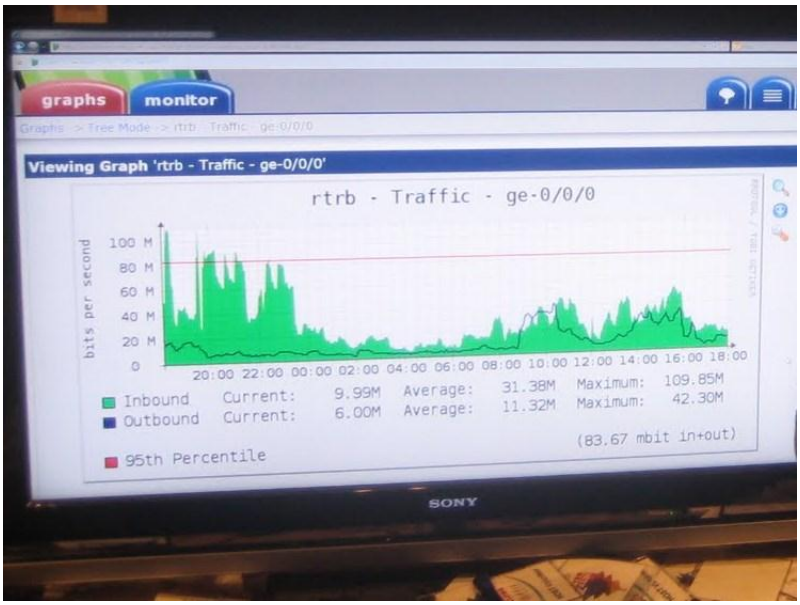

Figure 7: Cacti graph showing border router traffic on a 40“ Sony HDTV.

[1] The terminal room went to production one day earlier because of code sprint activity.
[2] Cacti: http://www.cacti.net/

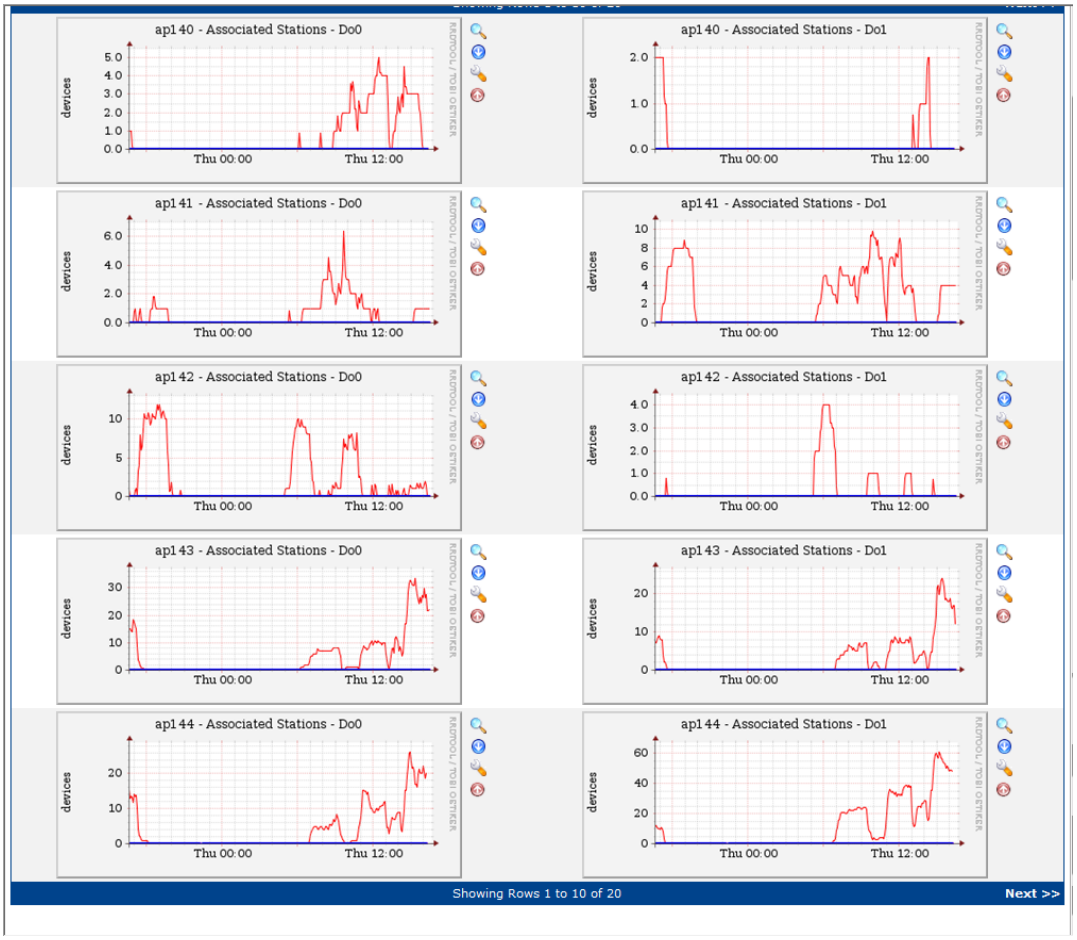

Figure 8: Cacti graph showing wireless AP association counts

We also used RANCID tool[3] to monitor and track configuration changes on network devices. On servers, RCS was used to avoid multiple users modifying same configuration file at same time as well as the version control.

5.2 Issue Reporting and Handling

There were three ways for an IETF attendee to report a network issue: 1) go to the helpdesk; 2) open a ticket; and 3) chat with NOC engineers through Jabber.

1. The helpdesk was open from 8am to 8pm during the meeting week. NOC engineers and student volunteers took rotations to man the desk. On Nov. 7, there was a wrongly installed power strip which burned about 10 power adapters. The helpdesk became the go-to place for people to report and replace the bad adapters.
2. We also had a Trac-based ticketing system, monitored by a dedicated ticket master. Once a new ticket was opened, it would be dispatched by the ticket master to a NOC engineer to resolve it.

[3] RANCID: http://www.shrubbery.net/rancid/

3. Jabber chat room was monitored together by the helpdesk and the NOC. It was convenient especially when a user sat in a meeting and experienced network problems.

One lesson we learned is we should plan duty shifts better so the helpdesk would always have someone there during working hours. Moreover, we should also provide more training to our engineers so they could be more familiar with the ticket system and Jabber tools.

Based on the user survey data, IETF79 meeting network was reasonably successful (56% users rated the network as Excellent). But there are always rooms for improvements. User complaints came from three major areas:

1. *Wireless*
    a) Some users complained about 802.1x authentication, they couldn't get it work. It was known that some Windows operating system defaulted to its own authentication mechanism, so certain configuration changes need to be applied. For other operating systems like Linux, we suspected a compatibility issue with drivers.
    b) Some complained about general wireless problems like intermittent associations, slowness, and so on. We were not quite sure about the root cause and we hoped we could have had more time to investigate it.
2. *Printer* was reported not working until Tuesday night in the terminal room. It was not a mere technical problem. It was a bad call to move printers around and people got confused. A lesson learned here is we should fully test printers before use. In addition, a printing station may be a good idea by setting up a dedicated computer directly connecting to printers, so people can use it without need downloading drivers.
3. *Power strip* problem was caused by a wrong installation. It was our fault so we replaced all damaged adapters as quickly as possible. It was recommended to check all power strips after the installations.

6 Tearing Down

From late afternoon Friday, Nov. 12, we started to tear down the IETF network. All hardware was collected to a storage room and a spreadsheet was used to track them to make sure we didn't miss any. Cables were removed and the hotel network was restored. IETF equipment was shipped back on Nov. 15. M10i was returned to Juniper on Nov. 16. The rest was shipped back to Tsinghua.

7 Appendix

Appendix A: IETF Meeting Network Requirements

Date: June 9, 2009
Editor(s): Karen O'Donoghue, Jim Martin, Chris Elliott, Joel Jaeggli
(Source: http://iaoc.ietf.org/network_requirements.html)

Abstract
The IETF Meeting Network has become integral to the success of any meeting. Building such a network, which provides service to thousands of heavy users, spread throughout the event venue, with very little time for setup and testing is a dramatic challenge. This document provides a set of requirements, derived from hard won experience, as an aid to anyone involved in designing and deploying future networks.

1. Requirements Notation
The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2. External Connectivity

- A primary and backup link MUST be provided for redundancy. If technically feasible, these links SHOULD be aggregated or load balanced.
- The primary link MUST provide at least 45Mb bandwidth in both directions, and SHOULD have at least 100Mb bidirectionally. (Note: Historically, bandwidth utilization peaked at 80Mb and averaged 35Mb.
- Recent events have peaked at 50 Mbs and averaged in the 25Mb range.) The backup link MUST have 10 Mb bandwidth in both directions and SHOULD have at least 45 Mb bandwidth in both directions.
- The backup link SHOULD be supplied by a different Internet service provider from the primary link.
- The primary and backup links SHOULD have physical and logical path diversity.
- IPv6 MUST be provided (possibly via a tunnel).
- The transit provided in support of the IETF MUST be capable of providing access to the IPv4 and IPv6 default free zones without the imposition of content filtering (e.g., URL, Site, application, port, or DPI based filtering).
- The primary link MUST support BGP peering, and the backup link SHOULD support BGP peering.
- Routing MAY be configured to allow the simultaneous usage of the bandwidth of both the primary and backup links.

- Access to research networks, like those that are part of Internet 2, MAY be provided on one of the external links.
- AS Numbers MAY be supplied by the network provider. If not, the network provider MUST use the AS Numbers provided by IETF.
- The network provider MUST provide at least a /19 of provider independent public IPv4 space or allow the IETF to advertise their own space.
- The network provider MUST provide at least a /32 of LIR public IPv6 space or allow the IETF to advertise their own space.
- If providing access space, the network provider MUST provide proper IP address delegation for DNS reverse lookups.

3. Meeting Facility

- The facility SHOULD have as much physical separation as possible in the meeting room area to improve the RF environment. In addition, the facility SHOULD avoid using airwalls and other partitions with low RF attenuation in the 2.4Mhz spectrum between meeting rooms.
- The facility SHOULD provide a RF environment in all meeting rooms (as identified by the Secretariat), common gathering spaces around the meeting rooms, the registration area, and the terminal room that has a reasonable noise floor in the 2.4Mhz spectrum.
- The meeting facility SHOULD have installed network cabling that can be used to deploy the network infrastructure.
- The meeting facility SHOULD provide the network installation team with 24 hour access to key telecom spaces. The meeting facility MUST provide the network installation team with access to key telecom spaces from one hour prior to the beginning of sessions to one hour after the end of sessions and 9am to 5pm daily during the setup period.
- All locations for network gear, with the exception of wireless APs, MUST be secure.
- If wireless will be used for an external link then access to the roof or installed location MUST be provided.
- The meeting facility MUST have adequate ventilation to support the equipment rooms and the terminal room.
- The meeting facility MUST have adequate power available to support the equipment required to support the network infrastructure and its users. This may include 110v/220v requirements in technical closets, roof locations, and various public and back-of-the-house areas.
- The meeting facility The meeting facility SHOULD have UPS power available to support key network infrastructure components, including at least the core routers, core switches, and hardware to maintain the external links.
- The meeting facility MUST provide sufficient power in all meeting rooms to handle the projected load from users' laptops, using 100% congruency between the projected number of attendees in each meeting room and the number of laptop users and projecting 70 watts of power usage per laptop.

4. Internal Network

- Wired Ethernet connections (network drops) MUST be provided in all the locations used for meeting room audio distribution for the purposes of audio recording and transmission.

- Wired network drops MUST be provided to the registration desk.
- The network SHOULD have separate VLANS for wired (primarily terminal room and audio) and wireless traffic.
- The network MUST NOT prohibit end-to-end and external connectivity for any traffic (no limiting firewalls or NATs).
- The network SHOULD have mechanisms for detecting and silencing rogue servers (DHCP, IPv6 RA's, etc)

5. Terminal Room or equivalent

- Terminal Room or equivalent A terminal room MUST be provided. This terminal room MAY be a single room or distributed sites in reasonable proximity to the meeting rooms.
- The terminal room MUST provide Ethernet 10/100 connectivity with RJ-45 connectors (approximately 100-150 drops required). (note: this number should be revised based on terminal room usage statistics)
- The terminal room SHOULD provide a small number of desktop or laptop computers for emergency use by attendees (minimum application requirements are web browsing, word processing, presentation production, and printing capability).
- The terminal room SHOULD have 24 hour access. This access SHOULD include security, but it MAY not include a 24 hour staffed help desk.
- The IETF users MUST have access to the terminal room from one hour prior to the beginning of sessions to one hour after the end of sessions.
- The terminal room MUST provide at least two network connected enterprise class printers. These printers SHOULD have duplex capability.
- A color printer MAY be provided.
- The terminal room MUST have a manned help desk from one hour prior to the beginning of sessions to one hour after the end of sessions. The help desk provides technical assistance to attendees, provides one potential interface to the trouble ticket system (see next requirement), and maintains the printers.
- The network supplier SHOULD provide a trouble ticket system to track attendee network issues. This trouble ticket system SHOULD be accessible to the help desk staff in addition to NOC staff.
- Power strips MUST be provided in the terminal room.
- Power strips MAY be provided in common gathering areas (desirable).
- The terminal room MUST have physical security (guards) during operating hours.

6. Wireless

- The network MUST provide 802.11b coverage in all meeting rooms (as identified by the Secretariat), common gathering spaces around the meeting rooms, the registration area, and the terminal room.
- The network SHOULD provide 802.11b coverage in additional common spaces in the meeting venue. The lobby, bar, restaurant, and most commonly used hallways of the primary meeting hotels, SHOULD also be provide 802.11b access.
- The network SHOULD provide 802.11g in all the spaces identified above.

- The network SHOULD provide 802.11a coverage in as many of the above identified spaces as possible focusing on the spaces with the highest user density first (e.g. plenary meeting room).
- The network design MUST anticipate 100% congruency between the projected number of attendees in each meeting room and the number of wireless network users (historical utilization in excess of 1000 simultaneous wireless users has been observed during a plenary session).
- The network SHOULD provide separate SSIDs for 802.11b and 802.11a networks.
- The network MUST provide fully open (unsecured) wireless access.
- The network MAY provide additional secured (WEP, 802.11i, WPA) wireless access.
- There SHOULD be mechanisms for identifying and silencing rogue Wireless Access Points.

7. Services

- The network MUST provide redundant DHCPv4 servers.
- The network SHOULD provide DHCPv6 service.
- The network MUST provide local redundant DNS servers.
- The network SHOULD provide NTP.
- Printers MUST support IPP and SHOULD support LPR and Windows printing.
- The network MUST provide a SMTP server providing relay services for the IETF network.
- The network SHOULD provide a full on-site mirror of the RFC and I-D directories.

8. Network Monitoring

- The network MUST provide sufficient monitoring to ensure adequate network availability and to detect faults before they impact the user experience.
- The network SHOULD provide some visibility into the state of the network for attendees (e.g. public graphs of network utilization, number of wireless associations, etc.).
- The network MUST collect data for future use in scaling IETF meeting network requirements. Minimum required metrics include bandwidth utilization (average and peak) for each external connection and user density per AP and radio.
- The network provider SHOULD provide SNMP read-only access to the network devices to individuals as identified by the Secretariat for network management and planning purposes.

9. Miscellaneous Requirements

- The network provider SHOULD maintain spares of critical network components on-site.
- Attendees SHOULD be notified of power connector requirements well in advance of the meeting via both the IETF meeting web page and the IETF- announce mailing list.
- A document MUST be provided to attendees detailing on-site network configuration information, including wireless configuration details, services available (e.g. printing, SMTP), instructions on how to report network issues (e.g. trouble ticket system interface instructions), etc. Initial versions of this information SHOULD be provided in advance of the meeting.
- The network provider MUST NOT view the IETF network as an experimental facility at the risk of impacting the IETF attendee experience. (Do not experiment with his/her favorite pet technology.)

- The network provider SHOULD have attended at least one prior IETF to observe the IETF network deployment and operation.
- The network provider SHOULD supply the IETF network design to an IETF technical review team for comments.

10. Acknowledgements

These requirements are born out of the pain and experience of past hosts and volunteers. Contributors of particular note are (in no particular order):

- Jim Martin
- Karen O'Donoghue
- Chirs Elliott
- Joel Jaeggli
- Lucy Lynch
- Bill "wej' Jensen
- Chris Liljenstoipe
- Bill Fenner
- Hans Kuhn

References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.