Cisco IOS: All the Small Things

PRESENTED BY:

JASON BOMAR, CCIE #9316

LAnet

Loopbacks

Use them. Even if you don't know why, add them, then figure out why you added them, and use them. A few things that they are good for:

- Sourcing Logs
- Sourcing TACACS
- Sourcing SNMP
- Sourcing Gateway registration (UC)

See a trend? Having a well defined management address to source traffic from adds security and stability. Use them!

Loopbacks

Interface loopback0

ip address 10.1.1.1 255.255.255.255

Router eigrp 1

eigrp router-id 10.1.1.1

network 10.1.1.1 0.0.0.0

ip tacacs source-interface loopback 0

snmp-server source-interface traps loopback0

logging source-interface loopback0

NTP

NTP is critical for so many reason:

- Assists in troubleshooting
- Assists in forensics
- Correlation of events
- Etc.

Pop Quiz: what interface will you be sourcing NTP from?

NTP can be authenticated if you so desire.

NTP

ntp server 1.1.1.1 source loopback0 prefer

ntp server 2.2.2.2 source loopback0

Show ntp status

Show ntp associations <detail>

SNMP

If you have a network device and a management system, chances are you have SNMP configured. Some thoughts and recommendations:

- Do you NEED SNMP RW? If not do NOT enable it.
- Consider SNMPv3, it is more secure.
- Any SNMP access MUST be secured with an ACL.

Pop Quiz: What interface will we source SNMP from?

SNMP

```
snmp-server community <something tough> RO RO-SNMP
snmp-server community <something tough> RW RW-SNMP
snmp-server trap-source Loopback0
snmp-server location LA Office
snmp-server contact LeBron James
snmp-server enable traps snmp authentication

ip access-list standard RO-SNMP
 permit ip host 10.2.2.2
 permit ip host 10.2.2.3

ip access-list standard RW-SNMP
 permit ip host 10.2.2.2
```

LAnet

Logging

Logging is incredibly critical to your environment – first you need to figure out how many logs to send, then where to send them, and lastly you MUST review them.

- Logs can eat bandwidth, be aware of this.
- You should use a centralized logging service
- Determine your policy for how long to retain logs

Pop Quiz: What interface will we source logs from?

Logging

logging 1.1.1.1

logging buffered warnings

logging source-interface loopback0

logging trap notifications

logging on

DNS

To be clear I am not talking about running DNS on your routers and switches, I'm talking about adding them to DNS, and giving them the ability to resolve DNS.

Pop Quiz: What Interface gets added as the IP for the resolution when adding your devices to DNS?

So why do this? Let's pretend you having a routing loop – when you do a traceroute – whether from the router or a PC – wouldn't it be nice to know which routers/cities it is hitting/bouncing between? Lets also pretend you have 30 routers, not 3 =). Bottom line, names are easier to attach significance to than numbers, so do so please!

DNS

To enable the router to do lookups is easy:

ip domain-name la-networks.com

ip domain-lookup

ip name-server 10.6.6.6

ip name-server 6.6.6.10

SSH

Enable it.

```
ip domain-name yourdomain.com

crypto key generate rsa modulus 1024

ip ssh time-out 120
ip ssh authentication-retries 3

line vty 0 4
 transport input ssh
```

Access Control Lists

They are used for a lot of things, not just security. QoS uses ACL's, Route-map's use them, PBR uses them, etc. So it is important to think about the best way to deploy ACL's and set a standard ahead of time.

My recommendation is to use NAMED ACL's only, and to remark them in a reasonable manner. What is a reasonable manner? Well that depends on your environment, but I would say a good guideline is to remark for each major portion of an ACL rule set.

Access Control Lists

Sample Standard ACL

```
ip access-list standard route_filter
  remark *** Permit only 1918 addresses ***
  permit 10.0.0.0 0.255.255.255
  permit 172.16.0.0 0.15.255.255
  permit 192.168.0.0 0.0.255.255
```

Sample Extended ACL

```
ip access-list extended QoS_Gold
  remark *** Permit SSH ***
  permit tcp any any eq 22
  remark *** Permit Citrix ***
  permit tcp 10.0.0.0 0.255.255.255 any eq 1494
  permit udp 10.0.0.0 0.255.255.255 any eq 1604
  remark *** Permit to/from ERP VIP ***
  permit tcp any host 10.100.100.1 eq http
  permit tcp host 10.100.100.1 any eq http
```

LAnet

Banners

Banners can be applied to a number of processes: login, MoTD, exec, etc – I suggest that you check the docs on the why's and wherefores:

http://www.cisco.com/en/US/docs/ios/fundamentals/configuration/guide/cf_connections_ps6441_TSD_Products_Configuration_Guide_Chapter.html#wp1001226

Regardless of which way you go (I often use exec) you SHOULD use a banner which warns people away if they don't belong there.

Banners

Banner exec ^

WARNING - PRIVATE ELECTRONIC DEVICE - ACCESS PROHIBITED

This device is a private network device. Access to this device is

not authorized. Any attempt for unauthorized access will be logged

and appropriate legal action will be taken.

^

Services

no service tcp-small-servers
no service udp-small-servers
no ip finger
no service dhcp
no service pad
no ip http server
no ip http secure-server
no service config

service password encryption
service tcp-keepalives-in
service tcp-keepalives-out
service timestamps log datetime msec show-timezone
service timestamps debug datetime msec show-timezone

LAnet

Passwords and Secrets

Never use ‘passwords’ for securing your router, only use secrets.

Secret strength should be 8-9, which is the equivalent of SHA-256 and SCRYPT. Please do NOT use 4 ever, and if we see it, we should recommend changing it. Even secret type 5 is better than 4 as their was a flaw in the algorithm that allows it to be cracked quite quickly.

Unicast RPF

This is a security feature that should be used with caution (all of these features should be used with caution). It ensure that when a packet is received on an interface, it has a valid ***source*** address before forwarding the packet on. There are two modes (three really – but lets focus on just two) that affect this behavior:

- Strict Mode: When a packet arrives on an interface, the router checks the routing table, and ensures that the received interface would be the forwarding address to reach the source. If not, it drops the packet.
- Loose Mode: When a packet arrives on an interface, the router checks that the source network is in the routing table.

TCP Synwait

Ever try to SSH to 10.1.1.1 and accidentally type “SSH 10.1.1.11” where that is not a host running SSH? Know that really annoyingly long wait period (90 seconds) where you are waiting for the router to give up waiting for a SYN-ACK? Well you can tweak that down to anywhere between 5 seconds and 3 minutes (who wants to tweak it UP?!?!?!). It is a simple global command:

ip tcp synwait-time 6 <tweak tcp time to wait for a syn-ack to 6 seconds>

LAnet

Names, Descriptions, Remarks – oh my!

I am a big fan of using names, remarks, descriptors of any sort whenever possible. Need a static route? Give it a name! Want to enable an interface? Give it a description! Enabling an ACL? Remark it! Creating a VLAN? Name it!

- ACL's
- Route-maps
- Critical interfaces

Graceful Restart

Graceful Restart is the ability to route ***through*** an outage/interruption in service. Why would this happen? You typically want to do this when you have multiple routing engines, and one fails, and the other is taking over – think dual supervisors in a chassis.

- Graceful Restart Capable – This means that the device in question actively supports Graceful Restart.
- Graceful Restart Aware – This means that the device (usually upstream/downstream) is aware of how Graceful Restart works and will assist to the best of its ability.
- Supports OSPF, IS-IS, EIGRP, LDP and BGP

During a RP failover, the router will NOT learn new routes, but will keep forwarding along the pre-failure paths that existed until the new adjacencies are formed.

BFD

BFD is a very useful protocol for detecting failures in a L2 network. As opposed to Graceful Restart, BFD routes ***around*** a failure rather than through one. It is very handy because:

- The protocol is very simple to configure, and easily achieves sub-second failure.
- It works with the routing protocol (EIGRP, OSPF, BGP, IS-IS and HSRP), but does not rely on the protocol itself.
- It informs the protocol of the loss of connectivity and the routing process can them immediately expire all timers and fail to a redundant path.
- You must check with your ISP to see if they support this feature.

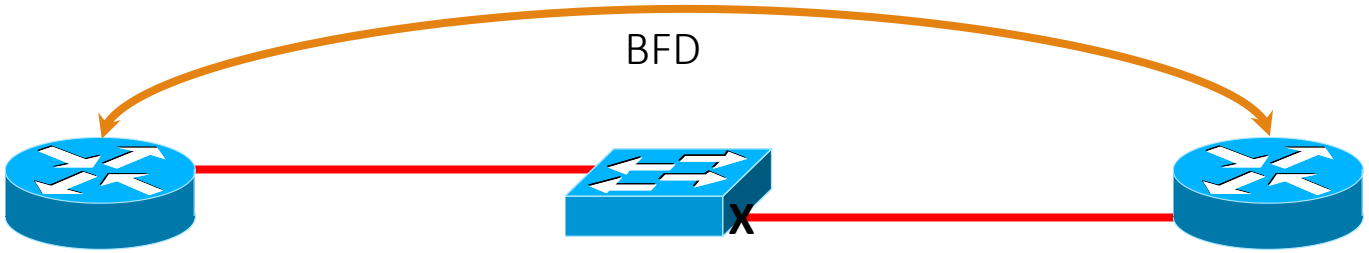

ASA – TCP Ping

Ever want to see if not just a host is up, but that it is listening on a service? You can use TCP ping on the ASA to assist with this:

Ping tcp mywebserver.com 80

This will send Syn packets to see if it gets a syn-ack back on port 80. Very handy when it comes to troubleshooting a FW and it can be used in conjunction with ‘capture’ command.

Gimme some links!!

Graceful Restart
BFD
Unicast RPF

LAnet

Thank You!