BGP and peering at ABQG

June 4, 2020

Ed May UNM DNG/ABQG

Overview

- A few basics of Routing
- Review of BGP in Routing
- Example BGP configuration
- WRN and peering for New Mexico via ABQG
- Some basic BGP commands
- Basic troubleshooting of BGP and link state for routing
- Description of Peering connections, and other future connections

Basics of Routing

- Routing is Destination-Based
  - When you attempt to reach another node/host/website on the internet, you do not enter your own node name, or address. You enter the one you are trying to reach. This is typically accomplished by unicast routing.
- Unicast routing
  - Delivers a message from one node to another single specified node, based upon a DNS (Domain Name Service) or IP (Internet Protocol) address specified
- A couple references for Routing
  - http://en.wikipedia.org/wiki/Routing
  - http://www.cisco.com/en/US/tech/tk365/technologies_white_paper09186a008014f8a9.shtml

Basics of Routing

- Don’t forget, that Routing is Destination-Based,
  - This means that a communication between two network-attached hosts occurs in TWO directions which can be mutually exclusive based upon routing policies.
  - Internet packets destined for a network outside of your own can take your default path out, and the response packet can return by another path based upon the distant network routing policy.
  - We will discuss this some more in BGP configuration
- Routing Policies
  - You base your routing decisions upon information you have about your networks and where your packets need to go, and as much as you can on how they will return(return control is limited).
  - Other Internet-connected organizations will do the same from their end.

Routing with BGP(Border Gateway Protocol)

- What is BGP?
- What are AS numbers?
- How do we use BGP?

What is BGP?

- BGP
  - An acronym for Border Gateway Protocol.
  - Is the accepted standard for Internet Routing between separate networks.
  - BGP bases these networks upon AS (Autonomous System) numbers
- BGP is a Routing protocol, not a Routed protocol.
  - A routing protocol is a protocol that specifies how routers communicate with each other to disseminate information that allows them to select routes between any two nodes (routers) on a network or connected networks.
  - A routed protocol is a protocol used with data embedded for host to host communication.
- BGP is a path vector routing protocol
  - It keeps track of the path a route is available based upon AS (Autonomous Systems) that the route is advertised through.
  - http://en.wikipedia.org/wiki/Path_vector_protocol

What are AS numbers?

- AS – Autonomous System
  - Autonomous System is a collection of connected IP routing prefixes under the control of one or more network operators that presents a common clearly defined routing policy to the Internet.
  - from RFC 1930
  - http://en.wikipedia.org/wiki/Autonomous_system_(Internet)

How do we use BGP?

- We use BGP to handle routing decisions between networks based upon policies that we set. These decisions are made both inbound and outbound on a network interface.
- BGP is a path vector protocol
  - It keeps track of the path a route is available based upon AS (Autonomous Systems) that the route is advertised through.
  - http://en.wikipedia.org/wiki/Path_vector_protocol
- For inbound updates, the order of preference is:
  - route-map
  - filter-list
  - prefix-list, distribute-list
- For outbound updates the order of preference is:
  - prefix-list, distribute-list
  - filter-list
  - route-map

How do we use BGP?

- We use BGP commands to accept certain Autonomous system paths and deny others. This allows us to receive only the routes that we want to see, while denying others that might be received.
- We use BGP commands to set preferences based upon AS paths that are announced to us, which allows us to prioritize one path over another.
- We use BGP to express the prefixes that we are responsible for to other networks that we have connections with.

BGP Configuration

- The basic commands needed to establish a BGP session
  - Router BGP <AS Number> establishes BGP running on the router
  - Neighbor statements to establish BGP connection peers
  - Network statements to establish which ip networks The AS is responsible for
  - BGP commands that make interoperability happen
  - BGP command to establish which networks will be advertised and received
  - Basic security configuration for BGP

BGP Configuration

```
Router BGP
local-as 40498
auto-shutdown-new-neighbors
neighbor 55.14.121.49 remote-as 36455
neighbor 55.14.121.49 description vl55-ABQG_Carbo-9aug11-ELM
neighbor 55.14.121.49 ebgp-multihop  2
neighbor 55.14.121.49 remove-private-as
neighbor 55.14.121.49 password 2 $IzgyIlNLdG4ra25DPWRX
neighbor 55.14.121.49 soft-reconfiguration  inbound
neighbor 5005:1900:2100::c55 remote-as 36455
neighbor 5005:1900:2100::c55 description vl55-ABQG_Carbo_v6-9aug11-ELM
neighbor 5005:1900:2100::c55 shutdown
neighbor 5005:1900:2100::c55 password 2 $IzgyIlNLdG4ra25DPWRR
neighbor 5005:1900:2100::c55 soft-reconfiguration inbound
```

BGP Configuration

```
address-family ipv4 unicast
network 29.29.29.29/8
redistribute connected
redistribute static
neighbor 55.14.121.49 route-map in Carbo_in
neighbor 55.14.121.49 route-map out Carbo_out
exit-address-family

address-family ipv6 unicast
network 2020:48e0::/32
redistribute static
neighbor 5005:1900:2100::c55 activate
neighbor 5005:1900:2100::c55 route-map in From-Carbo-v6
neighbor 5005:1900:2100::c55 route-map out To-Carbo-v6
exit-address-family
```

Configuration used by BGP

```
route-map Carbo_in permit 10
match ip address prefix-list Default
route-map Carbo_in deny 20
match ip address prefix-list Bogon-List
route-map Carbo_in permit 30
match ip address prefix-list Carbo-Accept-List
set local-preference 120


route-map Carbo_out deny 10
match ip address prefix-list Default
route-map Carbo_out deny 20
match ip address prefix-list Bogon-List
route-map Carbo_out permit 30
match ip address prefix-list OurPrefixList
```

<<NOTE: with Route-maps, many routers have an implicit DENY ALL at the end>>

Basic Configuration...

- We still have Prefix lists
- We still have Communities
- We still have access lists
- We still have more configuration options and different operating systems

WRN and peering for New Mexico via ABQG

- Internet2 in Chicago
- CENIC in California
  - Digital California
  - CAL-REN
- Level3(bought by Centurylink)
- Centurylink
- Google Global Cache
- Netflix

WRN Map

- Northern path passes through FRGP in Colorado and PNWGP in Washington
- Southern path passes through El Paso

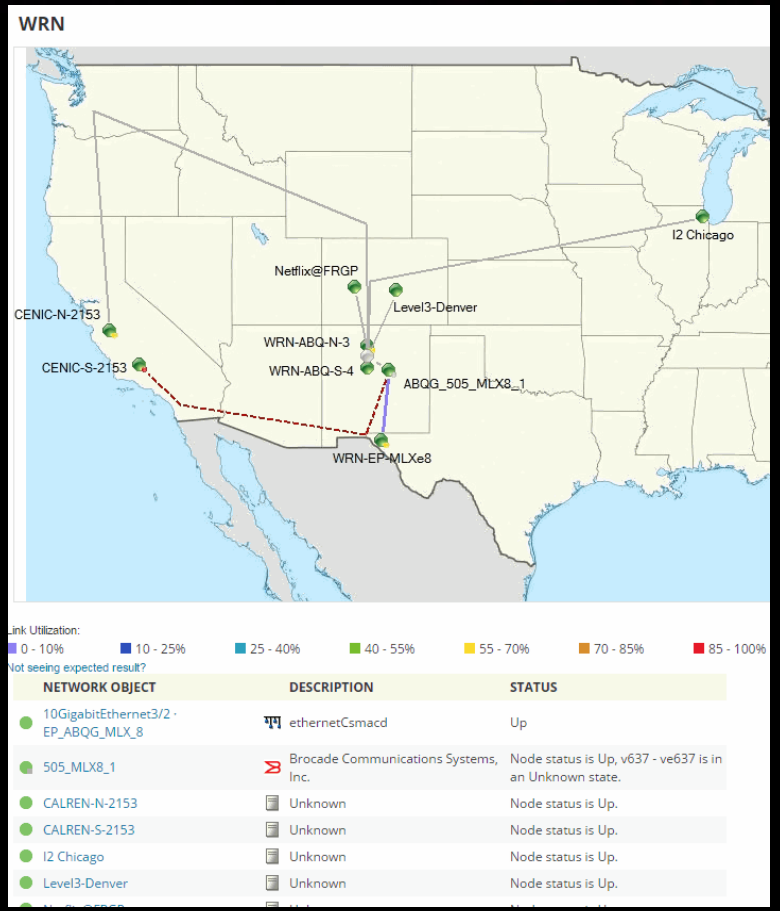

BGP Commands

- Show ip bgp summary – gives a summary of bgp sessions, how long the connection has been established, and the number of prefixes that have been received from that session.

```
RouterName# show ip bgp summary
BGP router identifier 55.14.121.50, local AS number 40498
BGP table version is 215716, main routing table version 215716
9983 network entries using 1168011 bytes of memory
9983 path entries using 519116 bytes of memory
1726/1723 BGP path/bestpath attribute entries using 241640 bytes of memory
1413 BGP AS-PATH entries using 41844 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 1970611 total bytes of memory
Dampening enabled. 24 history paths, 0 dampened paths
BGP activity 40868/30884 prefixes, 67415/57431 paths, scan interval 60 secs

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
55.14.121.49    4 36455   65592   29032   215716    0    0 2w6d        9958
RouterName#
```

BGP Commands

- Show ip bgp – requests a list of all routes that are in the BGP routing table for this router, includes routes from all peers, if you have more than one.

```
RouterName#show ip bgp
BGP table version is 215720,  local router ID is 55.14.121.50
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 6.1.0.0/16       55.14.121.49                         0 40498  19401  668 i
*> 6.2.0.0/22       55.14.121.49                         0 40498  19401  668 i
*> 6.3.0.0/18       55.14.121.49                         0 40498  19401  668 i
*> 6.4.0.0/16       55.14.121.49                         0 40498  19401  668 i
*> 6.5.0.0/19       55.14.121.49                         0 40498  19401  668 i
*> 6.6.0.0/16       55.14.121.49                         0 40498  19401  668 i
*> 6.7.0.0/16       55.14.121.49                         0 40498  19401  668 i
*> 6.8.0.0/20       55.14.121.49                         0 40498  19401  668 i
*> 6.9.0.0/20       55.14.121.49                         0 40498  19401  668 i
*> 6.10.0.0/15      55.14.121.49                         0 40498  19401  668 i
*> 6.14.0.0/15      55.14.121.49                         0 40498  19401  668 I
... continues for full list of Networks in the table...
```

BGP Commands

- Show run | begin bgp – The show run command will show the entire configuration for a Cisco router page by page. This can get quite lengthy, and adding **| begin bgp** causes the showing of the configuration to start at the first line with bgp in it, and shows the configuration page by page from there.

```
RouterName#show run | begin bgp
Router BGP
local-as 40498
auto-shutdown-new-neighbors
neighbor 55.14.121.49 remote-as 36455
neighbor 55.14.121.49 description vl55-ABQG_Carbo-9aug11-ELM
neighbor 55.14.121.49 ebgp-multihop 2
neighbor 55.14.121.49 remove-private-as
neighbor 55.14.121.49 password 2 $IzgyIlNLdG4ra25DPWRX
neighbor 55.14.121.49 soft-reconfiguration inbound


…Configuration continues by page…
```

BGP Commands

- Show ip bgp neighbor – gives a detailed summary of BGP information for a bgp neighbor that you name

```
RouterName#show ip bgp neighbor 55.14.121.49
BGP neighbor is 55.14.121.49, remote AS 36455, external link
 Description: vl55-ABQG_Carbo-9aug11-ELM
  BGP version 4, remote router ID 55.14.121.49
  BGP state = Established, up for 2w6d
  Last read 00:00:55, last write 00:00:42, hold time is 180, keepalive interval
is 60 seconds
  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Address family IPv4 Unicast: advertised and received
    Address family IPv4 Multicast: advertised and received
  Message statistics:
    InQ depth is 0
    OutQ depth is 0

                 Sent     Rcvd
    Opens:          5        5
    Notifications:      0        0
    Updates:            6    36359
    Keepalives:     29033    29245
    Route Refresh:      0        0
    Total:          29044    65609
…Continues on for the rest of the information for this BGP neighbor…
```

BGP basic troubleshooting

- Ping – locally and through Internet Service providers
- Traceroute – locally and through Internet Service providers
- BGP commands – locally, and through Internet Service providers
  - Ping, traceroute and route servers around the Internet can be accessed from many sites.
  - http://www.traceroute.org is a good resource for finding a few to select for testing.
  - https://www.internet2.edu/products-services/performance-analytics/performance-tools/ performance tools available for testing
  - I like https://sourceforge.net/projects/winmtr/ WinMTR
- Perfsonar.unm.edu

BGP basic troubleshooting

- Local Ping to verify link state

RouterName#ping 55.14.121.49

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 55.14.121.49, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

- Using an remote server; a ping from the remote server back to our router verifying connectivity to the remote network

……..>ping 55.14.121.50

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 55.14.121.50, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 33/33/34 ms

<<Note the difference in round-trip time for the ping. The test device is multiple hops away, thus more time required for the response>>

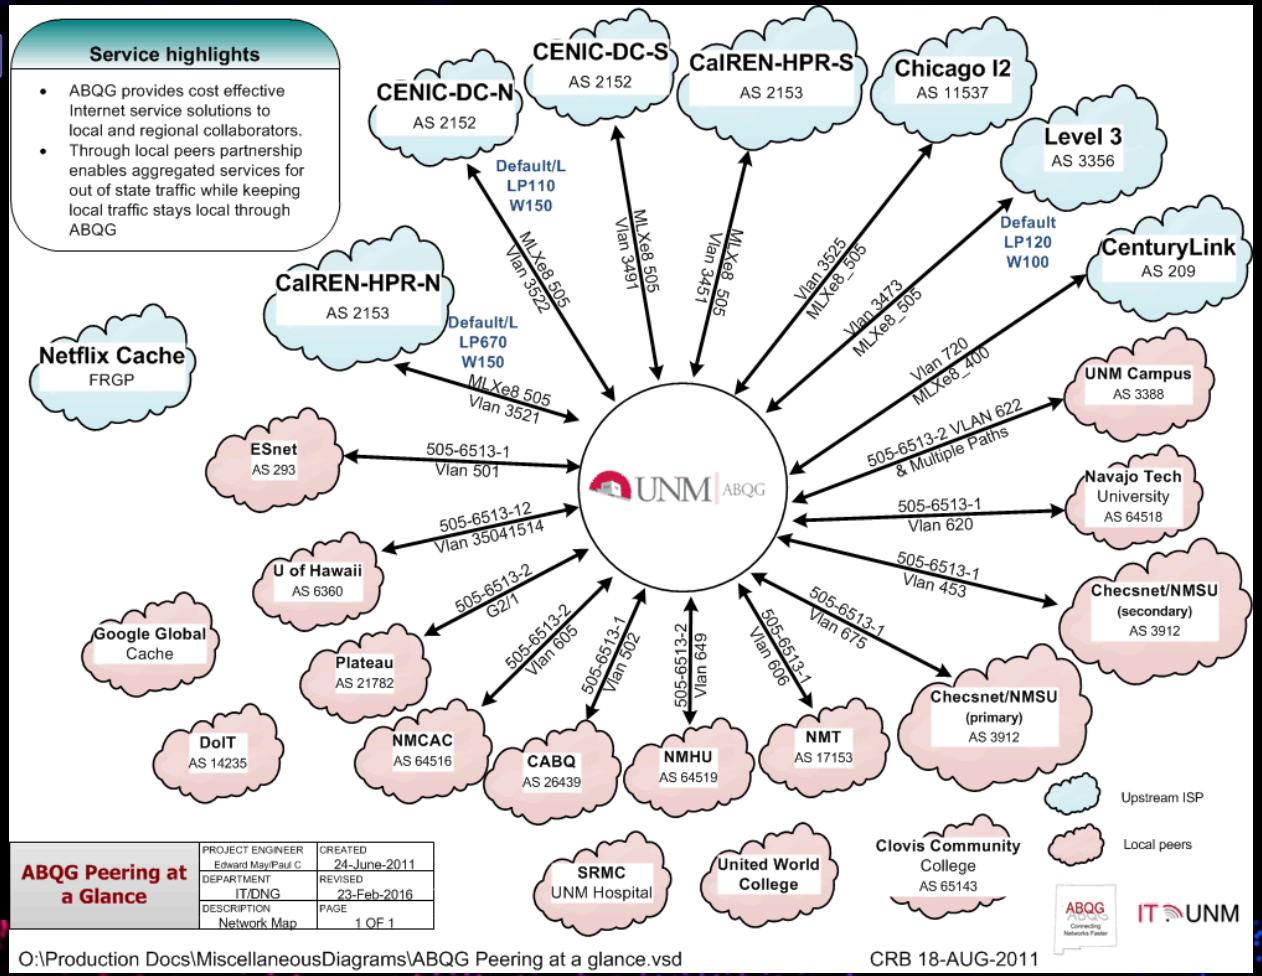
Service highlights
ABQG provides cost effective Internet service solutions to local and regional collaborators.
Through local peers partnership enables aggregated services for out of state traffic while keeping local traffic stays local through ABQG
CENIC-DC-N
AS 2152
CENIC-DC-S
AS 2152
CalREN-HPR-S
AS 2153
Chicago I2
AS 11537
Level 3
AS 3356
CenturyLink
AS 209
CalREN-HPR-N
AS 2153
Netflix Cache
FRGP
UNM Campus
AS 3388
Navajo Tech University
AS 64518
Checsnet/NMSU (secondary)
AS 3912
Checsnet/NMSU (primary)
AS 3912
ESnet
AS 293
U of Hawaii
AS 6360
Google Global Cache
Plateau
AS 21782
DoIT
AS 14235
NMCAC
AS 64516
CABQ
AS 26439
NMHU
AS 64519
NMT
AS 17153
SRMC
UNM Hospital
United World College
Clovis Community College
AS 65143
UNM ABQG
Default/L LP110 W150
MLXe8 505 Vlan 3522
MLXe8 505 Vlan 3491
MLXe8 505 Vlan 3451
Vlan 3525 MLXe8_505
Vlan 3473 MLXe8_505
Default LP120 W100
Default/L LP670 W150
MLXe8 505 Vlan 3521
Vlan 720 MLXe8_400
505-6513-2 VLAN 622 & Multiple Paths
505-6513-1 Vlan 501
505-6513-1 Vlan 620
505-6513-12 Vlan 35041514
505-6513-1 Vlan 453
505-6513-2 G2/1
505-6513-1 Vlan 675
505-6513-2 Vlan 605
505-6513-1 Vlan 502
505-6513-2 Vlan 649
505-6513-1 Vlan 606
Upstream ISP
Local peers
ABQG Peering at a Glance
PROJECT ENGINEER Edward May/Paul C
CREATED 24-June-2011
DEPARTMENT IT/DNG
REVISED 23-Feb-2016
DESCRIPTION Network Map
PAGE 1 OF 1
O:\Production Docs\MiscellaneousDiagrams\ABQG Peering at a glance.vsd
CRB 18-AUG-2011

ABQG Recent Connections and Possible Future

- Windstream built a lateral fiber into our suite for NNMCC Complete
- MRGPTC E-Rate network is connected via Zayo fiber, multiple 10g connections
- CNM peering connected via fiber at the APS datacenter
- Comcast in H5 Datacenter – peering to NMSU and NMHU
- Hurricane Electric is offering Internet service out of the H5 Datacenter
- Sacred Winds is building a suite a couple rooms up from UNM, no current plans for a cross-connect
- Navajo Tech is planning to peer with FRGP in Denver directly for Research Grants
- WRN is upgrading to Juniper routers
- ABQG will have dedicated 100g on WRN to Denver and California as part of the WRN upgrades

Thank You

Questions? Please e-mail to backbone-l@unm.edu