ABQG BGP Community Standard @ NMTIE

10/17/2008

ABQG
Connecting Networks Faster

UNM INFORMATION TECHNOLOGY SERVICES

Albuquerque Gigapop

- **Network Aggregation Point in Central New Mexico**
- **Keeping New Mexico network traffic in New Mexico**

Tilte: **ABQG BGP Community standard**

Summary:

ABQG (Albuquerque Gigapop), the New Mexico on-ramp service provider to Internet 2 and National Lambda Rail (NLR), also provides Local peering service for New Mexico local entities. The session will focus on BGP (Border Gateway Protocol) community strings that ABQG uses to serve its own peering partners and how peering partners can take advantage of this attribute for route manipulation.

Co-Presenter:
Ed May - Network Engineer UNM/ABQG **emay@unm.edu**
Paul Chang - Network Engineer UNM/ABQG **pchang@unm.edu**

Session Pre-requisites

- **Interested in peering with ABQG**
- **Somewhat familiar with BGP (Border Gateway Routing Protocol)**

ABQG Connectivity Overview

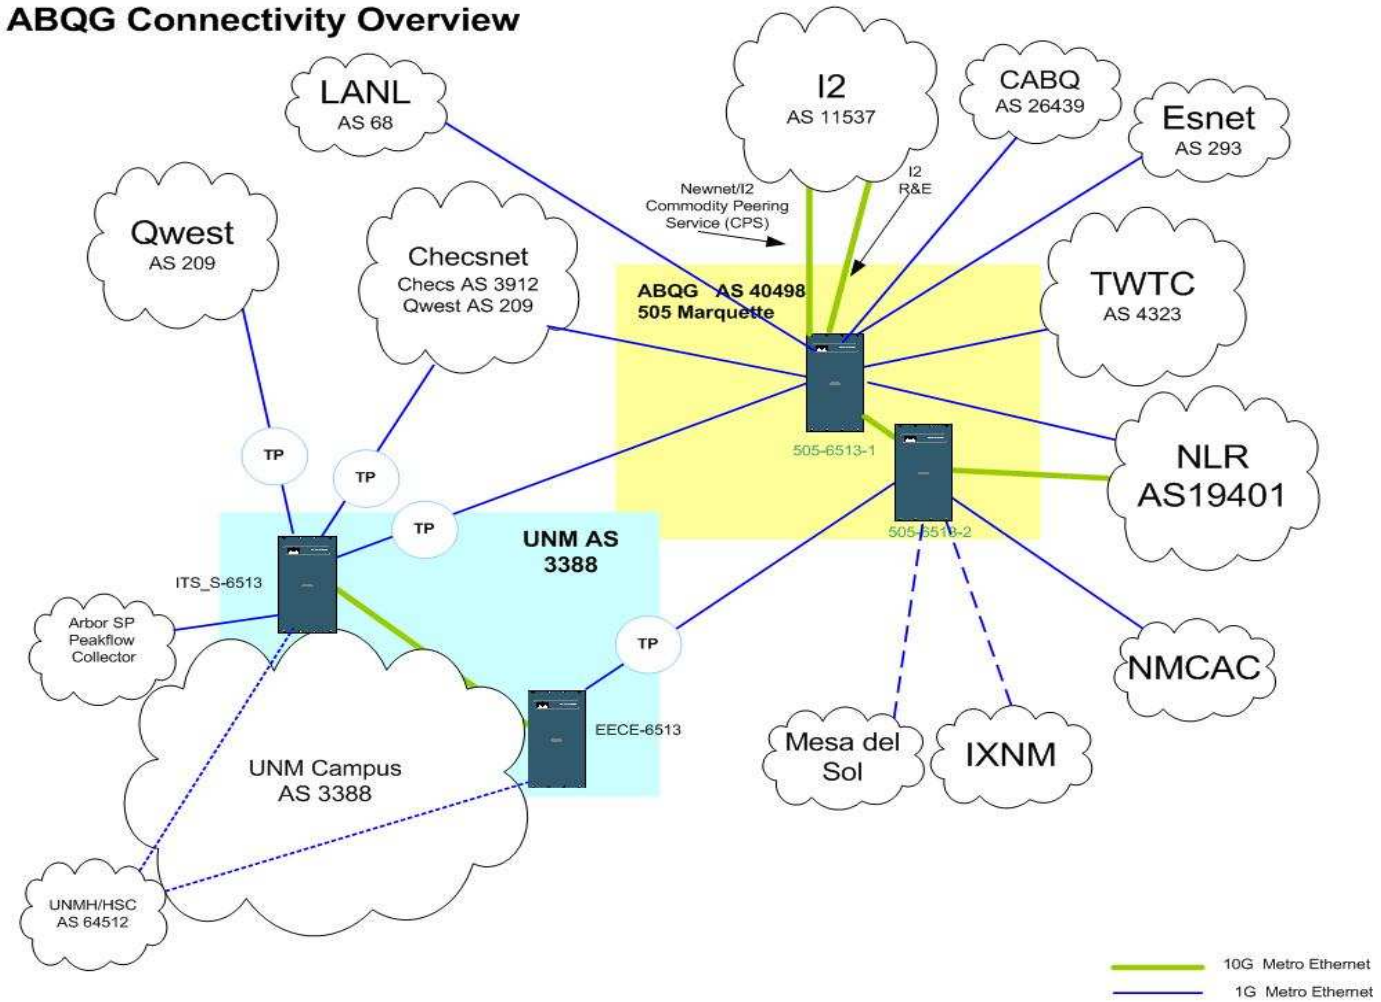

Prefixes in ABQG routing table

BGP peer	Prefix learned
NLR	9668
I2 R&E	11264
I2 CPS	97780
TWTC	1
UNM	16
Checs	120
LANL	5
Esnet	146
CABQG	4
Total	119,004

119,000 +

Too many prefixes !!!!

How to scale this to be manageable?

Typical given requests:

- **Newly peered client A wants to get UNM routes; I2 R&E routes but no LANL Routes ; no I2 CPS routes.**

Using BGP attributes:
AS origin…AS Path,
Next HOP, MED, Local preference….

BGP Community is the Answer!

ABQG BGP Community assignment
ABQG AS: Connector AS ==== 40498:XX

Items	Community Value	Description	Status	Service type
1	40498:4323	TWTC	In production	Commodity Service
2	40498:19401	NLR 10G	In production	NLR Service
3	40498:19402	NLR 1G	In production	NLR Service
4	40498:11537	I2 R&E	In production	I2 Service
5	40498:11538	I2 CPS	In production	I2 Service
6	40498:3388	UNM primary	In production	Peering partner
7	40498:3389	UNM secondary	In production	Peering partner
8	40498:68	LANL	In production	Peering partner
9	40498:3912	Checsnet	In production	Peering partner
10	40498:293	Esnet	In production	Peering partner
11	40498:26439	CABQ	In production	Peering partner
12	40498:17153	NMT	Coming soon	Peering partner
13	40498:64516	NMCAC	In production	NMCAC dedicated

Note

1. By defaults peering partner prefixes are announced to all ABQG peering partners

2. Community string listed above are additive.

How does it work? From ABQG side

- ABQG tags all incoming prefixes learned with ABQG BGP community standard additively.
- ABQG then announce prefix to its clients based on what type of combination services client selects.
- Now the control plant is set up for this particular clients.
- When packet hit ABQG router, data plant will route traffic based on the control plant info

How is it done? From ABQG side

BGP Mantra:

- Incoming policy affects outgoing traffic
- Outgoing policy affects incoming traffic.

- Example of the set up….

How to configure it?

A. Leg work first – some global configuration

```
ip as-path access-list 1 permit ^$
ip as-path access-list 68 permit ^3388_
ip bgp-community new-format
ip community-list 11 permit 40498:4323
ip community-list 12 permit 40498:19401
ip community-list 13 permit 40498:19402
ip community-list 14 permit 40498:11537
ip community-list 15 permit 40498:11538
ip community-list 16 permit 40498:3388
ip community-list 17 permit 40498:3389
ip community-list 18 permit 40498:68
ip community-list 19 permit 40498:3912
ip community-list 21 permit 40498:293
ip community-list 22 permit 40498:26439
ip community-list 23 permit 40498:17153
```

B. Route-map set up secondly so it can later applied to BGP peer:

```
route-map ITSin permit 10
 match as-path 68
 set local-preference 100
 set community 40498:3388 addictive
!
route-map ITSin permit 20
route-map ITSout permit 10
 match as-path 1
!
route-map ITSout permit 20
 match community 11 12 13 14 15 16 17 18 19 21 22 23
```

C. Finally apply to routing policy for this peer:

```
router bgp 40498
address-family ipv4

 neighbor 129.24.198.98 send-community both
 neighbor 129.24.198.98 route-map ITSin in
 neighbor 129.24.198.98 route-map ITSout out
```

D. Check the bgp routes to see whether community is properly applied.

E. Check whether the routes announced to your BGP peer is based on the community string list you applied.

Inspecting prefix sent by AS 668 - DREN

208.77.76.1 - PuTTY

```
  11537 668, (received-only)
    208.77.76.130 from 208.77.76.130 (64.57.28.246)
      Origin IGP, metric 0, localpref 100, valid, external
      Community: 668:100 11537:3000
505-6513-01#show ip bgp 6.4.0.0
BGP routing table entry for 6.4.0.0/16, version 37768977
Paths: (5 available, best #4, table Default-IP-Routing-Table)
  Advertised to update-groups:
     3           6           11          14          16
  19401 668
    129.24.198.110 from 129.24.198.110 (216.24.191.228)
      Origin IGP, metric 1351, localpref 400, valid, external
      Community: 668:100 19401:924 19401:7000 40498:19402
  19401 668, (received-only)
    129.24.198.110 from 129.24.198.110 (216.24.191.228)
      Origin IGP, metric 1351, localpref 100, valid, external
      Community: 668:100 19401:924 19401:7000
  19401 668, (received & used)
    129.24.198.106 (metric 20) from 208.77.76.144 (208.77.76.144)
      Origin IGP, metric 1350, localpref 600, valid, internal
      Community: 40498:19401
  11537 668
    208.77.76.130 from 208.77.76.130 (64.57.28.246)
      Origin IGP, metric 0, localpref 700, valid, external, best
      Community: 668:100 11537:3000 40498:11537
  11537 668, (received-only)
    208.77.76.130 from 208.77.76.130 (64.57.28.246)
      Origin IGP, metric 0, localpref 100, valid, external
      Community: 668:100 11537:3000
505-6513-01#
```

How does it work? From ABQG client side. (Name few possibilities)

- Do nothing
- Filter prefix(es) based on community value
- Filter prefix(es) based on combination of community value and AS origin
- Etc…..

Case Study I

Redundancy backup commodity service breaks between two ABQG partners.

Solution: Use ABQG community string attribute to manipulate inbound route policy to fix it.

(See details…)

What do we achieve here?

1. **Make routing policy task scalable**
2. **Provide easy maintenance and set up for peering arrangement**

What are other alternatives and limitations ?

- **Using MP-BGP VRFs**
- **MPLS-VPN**

ABQG Stats

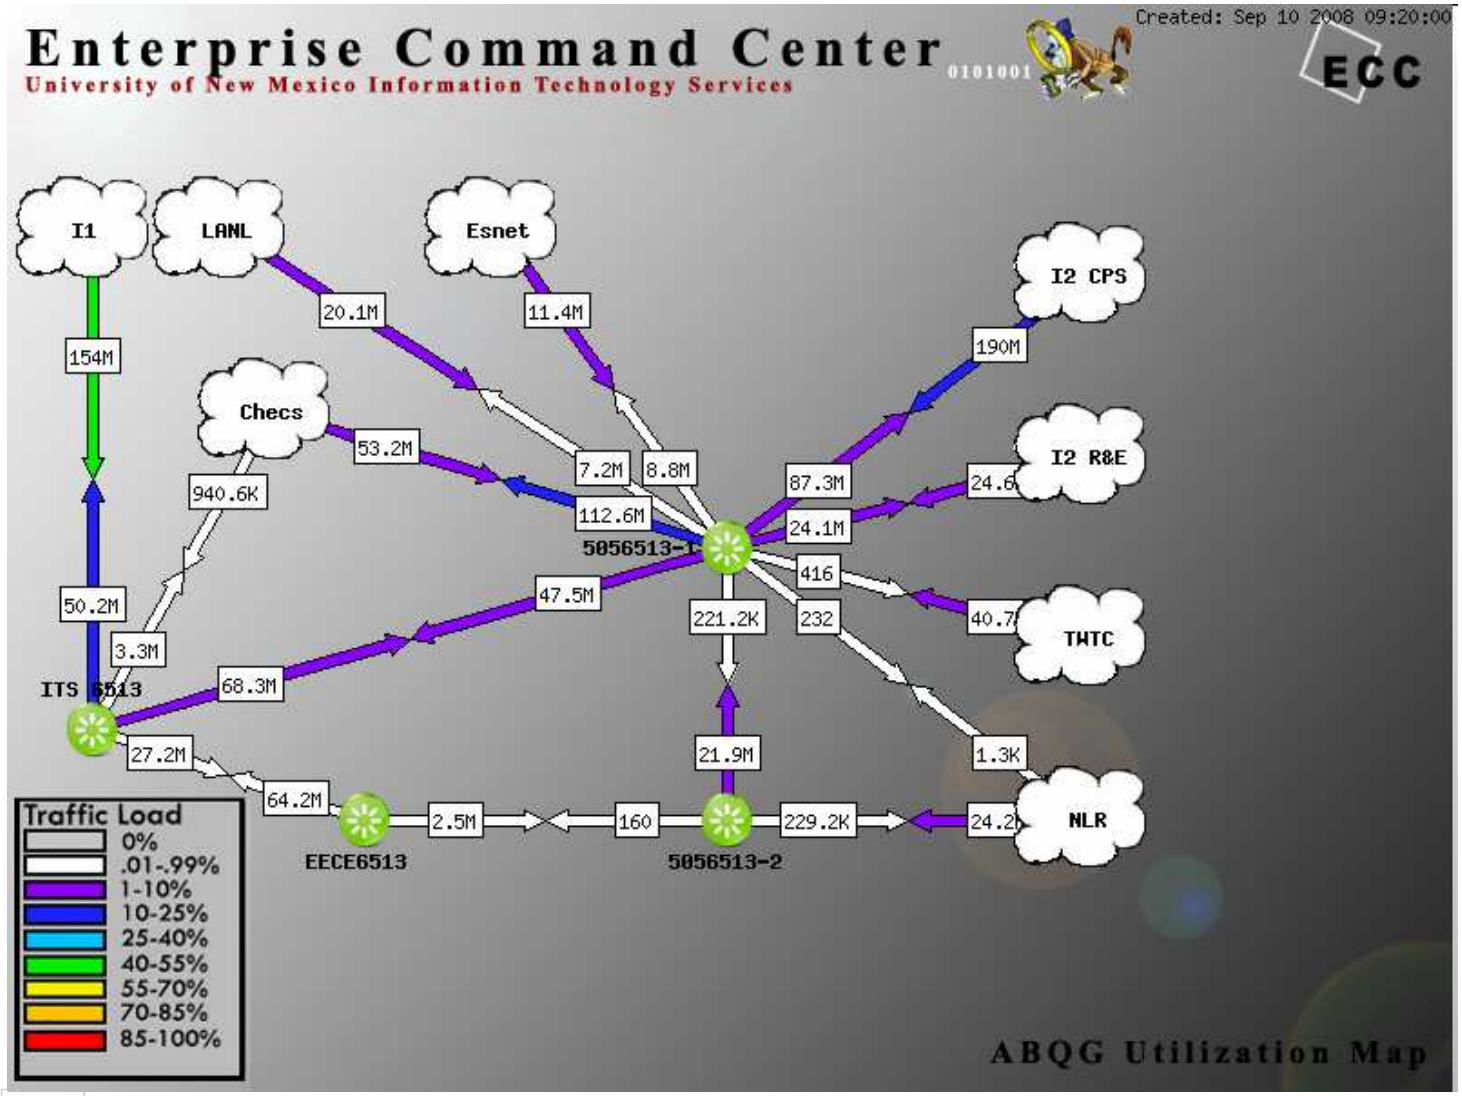
Enterprise Command Center
University of New Mexico Information Technology Services
Created: Sep 10 2008 09:20:00
ECC
I1
LANL
Esnet
I2 CPS
I2 R&E
Checs
TWTC
NLR
ITS 6513
EECE6513
5056513-2
20.1M
11.4M
154M
190M
53.2M
7.2M
8.8M
940.6K
87.3M
24.6
112.6M
24.1M
416
47.5M
50.2M
221.2K
232
40.7
3.3M
68.3M
27.2M
21.9M
1.3K
64.2M
2.5M
160
229.2K
24.2
Traffic Load
0%
.01-.99%
1-10%
10-25%
25-40%
40-55%
55-70%
70-85%
85-100%
ABQG Utilization Map

References:

1. **ABQG web site : http://abqg.unm.edu/**
2. **UNM Enterprise Command Center (ECC) : http://ecc.unm.edu**
3. **ABQG BGP community :**
4. **NLR community : http://noc.nlr.net/nlr/maps_documentation/nlr-packetnet-documentation/nlr-packetnet-bgp-communities.html**
5. **I2 Community : http://noc.net.internet2.edu/i2network/bgp-communities.html**
6. **Verizon Business Community and practice: http://www.cisconet.com/index.php/BGP-Community-Attributes-List/BGP-Community-String-for-Verizon-Business-AS701.html**
7. **BGP best path selection: http://www.ripe.net/projects/ris/docs/bgpcheat.html**
8. **This presentation:**

Q&A