

January 21, 2008

NTAC Peering Working Group
1/21/08

In attendance: Jeff Bartig, David Farmer, Ron Hutchins, Caren Litvan, Ana Preston, Chris Robb, Linda Roos, Paul Schopis, Joe St Sauver, Bob Stovall, Ryan Vaughn, Steve Wallace, Linda Winkler

Agenda

1. Quarterly technical update
2. International peering

1. Quarterly update

Previously, Matt Davy was the principle contact within the NOC for CPS-related issues and those responsibilities have been transferred to Caren Litvanyi.

Caren talked about the current peers and showed the NOC page that contains the utilization graphs. The NOC continues to work on the statistics to describe the traffic.

There group mentioned previous NTAC netflow data discussions. Steve Wallace will frame a question to the NTAC on sharing netflow data.

Caren encouraged the group to bring ideas regarding entities with which to peer. She will pursue the ideas for peers. It may or may not be possible to peer with the networks that people mention as peers have requirements such as volume requirements, etc.

Caren mentioned that Internet2 has an entry in peeringdb and that is the mechanism for other networks to find Internet2. There is an internal mailing list for peering (peering@internet2.edu). Caren, Steve Wallace, Rob Vietzke read the mail sent to that list.

It was explained that, when connecting to exchange points, we have only one vlan, rather than one for v4 and one for v6. With this single connection and single vlan, there is no clean way to take v6 and connect to VRF. All solutions for enhancing v6 via CPS were expensive, difficult or nearly impossible so it was decided that we would put v6 routes into commercial VRF. It is not necessary to be v4 CPS customer in order to peer with v6. However, all connectors that want to retain v6 transit routes will have to do v6 CPS. Internet2 won't list the connector as a CPS customer on webpages if the connector only buys v6. It was also mentioned that, if a connector has a vlan for CPS v4 service, that vlan can also be used for v6 although it is possible to have a separate vlan for v6.

It was mentioned that there are issues around addressing and that it may be time to consider getting provider independent addresses. Internet2 had proposed using this opportunity to encourage all to obtain provider independent addresses and had approached the NTAC with that plan. The NTAC indicated that they disagreed with the plan. Most want to move to provider independent addresses eventually but not right away.

NTAC has latest plan of record on the v6 migration (also included in these notes---see below). The document currently doesn't contain deadlines. The group suggested that timeframe be proposed to NTAC for a response.

Discussion ensued regarding queuing. Dave Farmer will frame a question to the NTAC regarding queuing.

As a final comment regarding CPS, it was emphasized that, if CPS is successful, Internet2 can react quickly to increase the capacity of the network.

Plan to Migrate Commercial IPv6 Peers from R&E VRF to CPS VRF January 8, 2008

Background

One of the goals of the CPS project is to improve commercial IPv6 connectivity for our R&E connectors and members. In addition to peering with other network providers via IPv4, our presence at these public peering locations has enabled expanded and upgraded IPv6 connectivity. It was our intent to connect the commercial IPv6 peers to the traditional R&E network service (original instance of AS11537, with which all connectors peer), however, due to a limitation in the Juniper VRF implementation, we're recommending a different approach. Rather than have the R&E network peer with commercial IPv6 networks, we're recommending that the CPS peer with them. The CPS network would then contain all unicast commercial traffic, both IPv6 and IPv4. The traditional R&E service would contain non-commercial IPv6 and IPv4, as well as all IP multicast traffic. Connectors may choose to BGP peer with the CPS for either or both IPv6 and IPv4 commercial connectivity.

The net effect of connecting commercial IPv6 peers to the CPS network is that for Internet2 Connectors to realize the improved IPv6 connectivity, they will be required to establish a second BGP session with the Internet2 router. This second session may be used for both CPS as well as the improved IPv6 connectivity, or either independently. All IP multicast connectivity (both v6 and v4) will continue to be provide via the R&E network.

During discussions with the NTAC and the IU NOC, two additional issues were raised:

The desire for continued universal access to IPv6 performance measurement servers and the opportunity to aid in the migration from Abilene-allocated IPv6 prefixes to provider-independent prefixes for all Internet2 Connectors.

I've discussed the IPv6 server access issue with Chris Small and the NOC. We identified three potential technical solutions: connecting each server to a sub-interface that's associated with a blended VRF, connecting each server to both VRFs (either by distinct physical interfaces or using VLAN tags) and running a BGP demon, or connecting each server to both VRFs and use policy routing to cause packets to return over the same interface they entered.

We're also suggesting that we take this opportunity of coordinating a change with all IPv6 connectors to aid their migration from Abilene-allocated IPv6 space to something the connectors can own and control, however this is not being considered a prerequisite to moving the connector to the CPS VRF for commercial IPv6 connectivity. In other words, a connector whose only IPv6 connectivity is through Internet2 (IP network and/or CPS) can use an IPv6 address allocation from the Internet2 block. Connectors that are multiply homed have several choices: 1) use their own address space from ARIN, 2) use a block from their other provider (Internet2 will NOT announce that block via CPS) or 3) use blocks from each provider. Each approach has advantages and disadvantages; the Internet2 NOC can provide advice.

In addition to connectivity to peers of Internet2, the Internet2 network also provides IPv6 transit. Internet2 purchases a small amount of IPv6 transit for this purpose. During the migration of commercial IPv6 routes to the CPS VRF, we'll also be moving the transit offering to CPS VRF. This will allow Internet2 to take advantage of its multiple peering locations to secure more resilient IPv6 transit service.

Desired End-State

We're proposing the following end-state:

- The CPS VRF would exclusively present commercial IPv6 (including IPv6 transit) and IPv4 unicast routes to Internet2 Connectors.
- The R&E VRF would exclusively present non-commercial IPv6 and IPv4 unicast routes as well as all multicast routes to Internet2 Connectors
- Each Internet2 IPv6-enabled connector would have at least an IPv6 peering with the CPS (for purposes of documentation, a IPv6-only CPS connector would not be advertised as a customer of CPS).
- A connector whose only IPv6 connectivity is through Internet2 (IP network and/or CPS) can use an IPv6 address allocation from the Internet2 block. Connectors that are multiply homed have several choices: 1) use their own address space from ARIN, 2) use a block from their other provider (Internet2 will NOT announce that block via CPS) or 3) use blocks from each provider. Each approach has advantages and disadvantages; the Internet2 NOC can provide advice.
- Internet2's IPv6-accessable performance measurement servers would be accessible from both VRFs.

Tasks

1. Chris Small - determine method for providing universal connectivity to IPv6 measurement servers. Suggest that this task proceed independently.
2. Caren Litvanyi enable IPv6 routing on the CPS VRF.
3. NOC (in coordination with Caren and Steve) communicate with each IPv6 connector the needs to a) establish an IPv6 peering session to the CPS, and b) if needed, to migrate to their own v6 prefix
4. Caren Litvanyi, after all #3 completed, start the migration of commercial IPv6 peers to CPS.