Transit Load Sharing

Overview

Definitions

Load Sharing: Load sharing provides Transit Service Providers (SPs), with two or more transit services, the ability to manipulate how we will prefer routes on each connection and will deliver traffic on each connection.

When load sharing, you will have multiple transit connections that provide access to the Internet for its routes. By nature, services that possess this load sharing characteristic are not redundant, unless separate unused resources have been purchased and allocated to address the requirement. Load sharing is achieved through route manipulation using BGP attributes.

We will provide pre-defined, proprietary, BGP community tags that can be used by you to change the normal preference given to your routes by our routing architecture. These tags can be used by you to achieve load sharing or to provide redundancy across two or more of PSINet Transit services.

Additionally, you may be able to use this functionality to coordinate traffic flows across our Transit service(s) and the transit service(s) you purchase from another provider.

BGP Community: A BGP attribute that, when utilized, contains a list of 32-bit community values used to identify a route belonging to a category of routes, all of which are treated the same with respect to a routing policy.

Internal Backup: Internal Backup describes the situation where your routes are to be preferred only if there is not another customer route to the same destination available within our network. Typically, this situation exists when load sharing is executed among our Transit services only.

External Backup: External Backup describes the situation where your routes are to be preferred only if there is not another route to the same destination available on the Internet. Typically, this situation exists when load sharing is executed among our Transit service(s) and those of another provider(s).

PSINet Transit Super Core (TSC): Our router that aggregates traffic to and from our Transit customers.

Requirements

Service Requirements

This Load Sharing service is provided and supported with the Transit service. You will normally have at least two Transit services with us. Your topology can be subcategorized into the following:

- Different Transit POPs to two or more routers at our POP. See Figure 2.2.1 below
- Same Transit POP to one router at our POP. See Figure 2.2.2 below.

Transit Load Sharing (cont'd)

Transit Customer Provisioning Requirements

Physical Circuit: Multiple physical circuits (Transit Interconnect Facilities) can be provisioned, either to different Transit POPs, or to the same Transit POP.

- **Different Transit POPs**. If you decide to use multiple circuits in different geographic areas, you should provide different Transit Interconnect Facilities to our Transit POPs in each of the different geographic areas. See Figure 2.2.1 below

Different Transit POPs

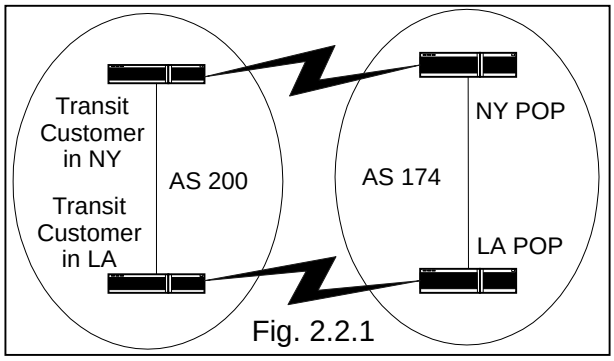

Fig. 2.2.1

- **Same Transit POP**. If you choose to connect multiple circuits to the same Transit POP, there must be multiple different physical circuits to the Transit POP. See Figure 2.2.2 below

Same Transit POP

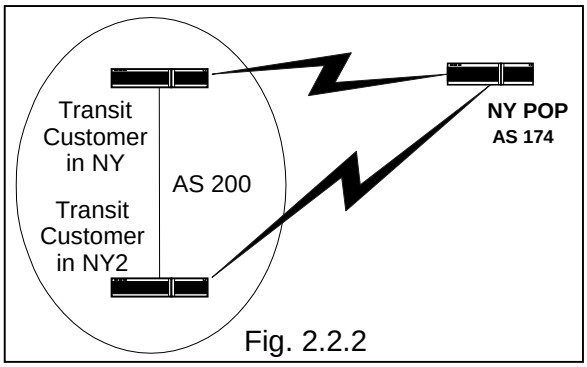

Fig. 2.2.2

Transit Load Sharing (cont'd)

Customer Requirements – Layer 3

To implement Load Sharing on two or more Transit services, you must:

- Have two or more, separate, Transit service(s)/Interconnect Facilities.
- Implement your Transit service(s) using the BGP4 routing protocol. Load sharing of Transit services under static routing is not supported.
- Choose one of the following load sharing methods for all of your Transit services:
  1. AS-based (AS-Path Padding) Load Sharing Method
  2. Our BGP Community-based Load Sharing Method

Transit Customer Load Sharing Implementation

AS-Path Padding Method

BGP's AS-Path attribute provides a list of the AS numbers traversed by a route to reach a destination.

You may elect to manipulate route preferences across our Transit service links via BGP and AS-Path Padding to implement load sharing.

Sample Customer Configuration:

```
Router bgp AS#
Neighbor PSI.ADD.TSC.ROU remote-as 174
Neighbor PSI.ADD.TSC.ROU ebgp-multihop 2
Neighbor PSI.ADD.TSC.ROU updat-source loopback0
Neighbor PSI.ADD.TSC.ROU route-map setpath out

Ip as-path access-list 10 permit ^$

Route-map setpath permit 10
Match as-path 10
Set as-path prepend AS# AS# AS#    (will prepend AS# 3 times)
```

BGP Community Method

With the use of our BGP communities, you can implement load sharing on your Transit Services with BGP by controlling its routing policy configuration. We will get the BGP community values from your routing announcements and apply the appropriate weight (Juniper routers use Preference) values to these routes.

Community Attributes

Community Tag	Description
No community set	PSINet prefers this route over others. **Default** Transit Customer Route.
174:26000	Treat this route equal to a PSINet peer route.
174:27000	Internal Backup (2 or more connections with PSINet, this is the backup community tag).
174:27999	Do not announce this route to any peers, internal backup.
174:29000	External Backup (2 or more connections, one with PSINet and the other(s) with other providers. This is the backup community tag for PSINet).
174:29999	Do not announce this route to any peers.

Transit Load Sharing (cont'd)

Route Selection will be based on tagged community routes. These tags will identify route preference based on our number plan for these communities.

Sample Configuration for BGP Communities

```
router bgp AS#
neighbor PSI.ADD.TSC.ROU remote-as 174
neighbor PSI.ADD.TSC.ROU ebgp-multihop 2
neighbor PSI.ADD.TSC.ROU update-source loopback0
neighbor PSI.ADD.TSC.ROU send-community
neighbor PSI.ADD.TSC.ROU route-map cust-tsc out
!
! do not send community
! Identifies customer route as Normal TRANSIT
route-map cust-tsc permit 10
match ip address 1
!
! Identifies customer route as Internal Backup
route-map cust-tsc permit 20
match ip address 2
set community 174:27000
!
!
! Identifies customer route as External Backup
route-map cust-tsc permit 30
match ip address 3
set community 174:29000

access list 1 permit Net.work.num.ber 255.255.255.0
access-list 2 permit Net.work.num.ber 255.255.255.0
access-list 3 permit Net.work.num.ber 255.255.255.0
```

Transit Load Sharing Examples

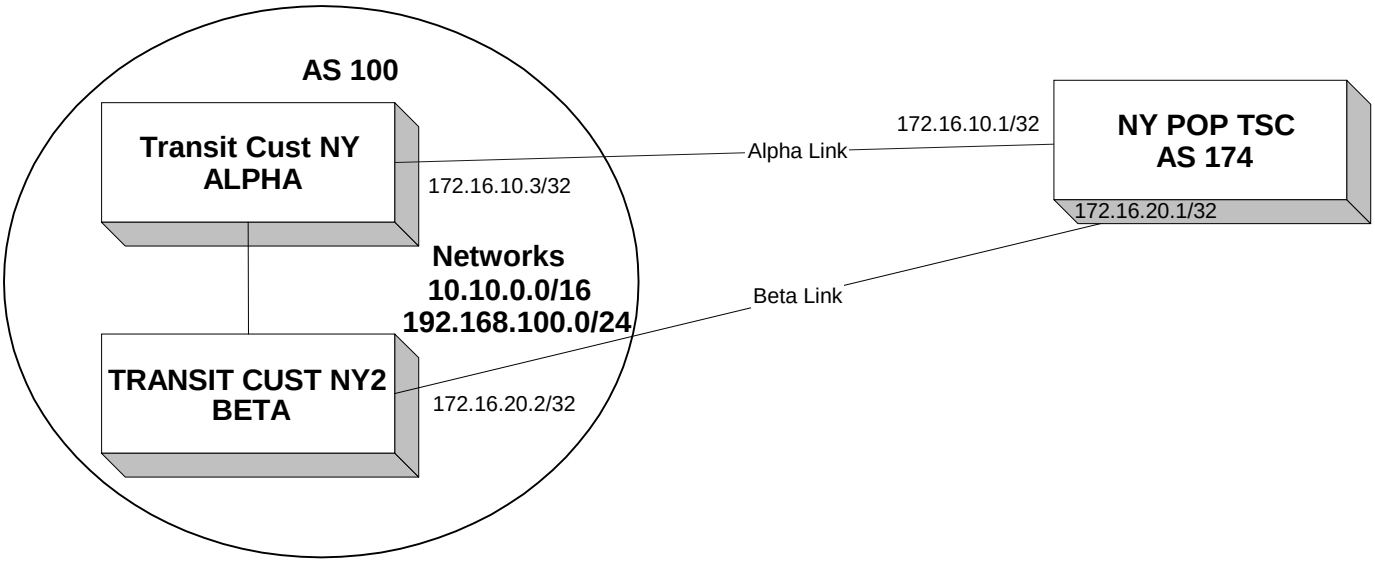

Transit Load Sharing (cont'd)

Example: Load Sharing via AS-Path Padding

Transit Customer has Transit Services (Alpha & Beta), implementing Load Sharing for the situation depicted above:

Sample Configuration for your NY router (ALPHA):

```
router bgp 100
neighbor 172.16.10.1 remote-as 174
neighbor 172.16.10.1 ebgp-multihop 2
neighbor 172.16.10.1 update-source loopback0
neighbor 172.16.10.1 route-map setpath out

ip as-path access-list 10 permit ^$

access-list 4 permit 10.10.0.0 0.0.255.255
access-list 5 permit 192.168.100.0 0.255.255.255

route-map setpath permit 10
match as-path 10
match ip address 4
set as-path prepend 100 100

route-map setpath permit 20
match as-path 10
match ip address 5
```

Sample Configuration for your NY2 Router (BETA):

```
Router bgp 100
Neighbor 172.16.10.1 remote-as 174
Neighbor 172.16.10.1 ebgp-multihop 2
Neighbor 172.16.10.1 update-source loopback0
Neighbor 172.16.10.1 route-map setpath out

Ip as-path access-list 10 permit ^$

Access-list 4 permit 10.10.0.0 0.0.255.255
Access-list 5 permit 192.168.100.0 0.255.255.255

Route-map setpath permit 10
Match as-path 10
Match ip address 5
Set as-path prepend 100 100

Route-map setpath permit 20
Match as-path 10
Match ip address 4
```

NOTE: From the diagram above, using the AS-path attributes, you will prefer to announce 10.10.0.0/16 to us via the Beta link and Alpha will serve as the backup for the 10.10.0.0/16 route to us via the Alpha link. Alpha will prefer to announce 192.168.100.0/24 via Alpha link to us and Beta will serve as the backup for the 192.168.100.0/24 route to us. Load sharing will be accomplished by you and your manipulation of your routing via the multiple links.

Transit Load Sharing (cont'd)

Example: Load Sharing via Our BGP Communities

Using the same situation diagrammed above, i.e., illustrating the same example, but utilizing the BGP Community Attribute to implement Load Sharing:

Sample Configuration for your NY router (ALPHA):

```
Router bgp 100
Neighbor 172.16.10.1 remote-as 174
Neighbor 172.16.10.1 ebgp-multihop 2
Neighbor 172.16.10.1 update-source loopback0
Neighbor 172.16.10.1 send-community
Neighbor 172.16.10.1 route-map cust-tsc out

Access-list 4 permit 10.10.0.0 0.0.255.255
Access-list 5 permit 192.168.100.0 0.255.255.255

Route-map cust-tsc permit 10
Match ip address 4
set community 174:27000

Route-map cust-tsc permit 20
Match ip address 5
```

Sample Configuration for your NY2 Router (BETA):

```
Router bgp 100
Neighbor 172.16.10.1 remote-as 174
Neighbor 172.16.10.1 ebgp-multihop 2
Neighbor 172.16.10.1 updat-source loopback0
Neighbor 172.16.10.1 send-community
Neighbor 172.16.10.1 route-map cust-tsc out

Access-list 4 permit 10.10.0.0 0.0.255.255
Access-list 5 permit 192.168.100.0 0.255.255.255

Route-map cust-tsc permit 10
Match ip address 5
set community 174:27000

route-map cust-tsc permit 20
match ip address 4
```

NOTE: With this configuration and the use of our community attributes, your Alpha router will announce 192.168.100.0/24 to us with community 174:27000 and 10.10.0.0/24 with no community. The Beta router will announce the route for 192.168.100.0/24 to us, as preferred over the Alpha link, and for the 10.10.0.0/16 it will announce community 174:27000.