NSI Topology/Pathfinding and Policies

Author: Henrik Thostrup Jensen `<htj at nordu dot net>`

Draft, date: 17/9/2014

Introduction

This document tries to give an overview on policies for transit network and exchange points and how these effect path finding is NSI. The combination of transit networks and exchange points is interesting, as one is often concerned with connectivity through multiple networks, where as exchange point are typically only concernted with local interconnects.

Transit networks

A transit network connects a number of customers together and provide them connectivity with the rest of the internet. This connectivity can either be via a single transit provider, or via a number of peers and a transit providers.

NORDUnet is a transit network for the five nordic countries which are both owners and main customers. It has around 600 peers and uses Telia traffic that is not exchanged via peers. Peers/transit provider cannot exchange traffic via the NORDUnet network. Traffic has to be to or from customers. NORDUnet has a number of other customers that pay for transit to research networks. For some of these customers there are peering in some places and transit in other places (overlapping). Most customers are multi homed, several of them over large geographic distances. Furthermore NORDUnet has shared infrastructure with SUNET and SURFnet (possibly more). Several projects have dedicated links equipment/services that connects or is hosted at NORDUnet (e.g., NDGF, the nordic LHC tier-1). These services have AUPs (acceptable use policy). I am probably missing a few other complications.

The best way to ensure that the policy of the transit network is respected and that resources are used most optimally within the network is to let the NSA of the network handle the path setup itself, by issuing a chain request to the network. Tree requests remove the possibility for the transit networks to setup its own path to or from its customers in the way that is most optimal for the network.

For customer requests, chaining gives the transit network the freedom to optimally route the request to the destination network. For connections coming from peer/transit side, chaining allow the networking to connect to the customer on the link best suited (this case is often a lot simpler).

Hence, for transit networks, chain requests, will typically be the preferred option. Some might only allow chain requests (this can be implemented by not allowing a local request between two demarcation ports). An open issue is how to ensure that a good handover point is chosen. Similarly to BGP, such must be arranged with bilateral agreements, and should possibly be signaled with a mechanism similar to MEDs in BPP.

Exchange Points

An exchange point is a place where multiple network can meet and exchange traffic, typically policy free. Commercial exchanges often solely focus on IP traffic, where exchanges supporting NSI will typically be R&E type, e.g., MAN-LAN, GEANT Open, NORDUnet Open Exchange (NOX).

An exchange point is essentially a switch. As such, it does typically not care about the networks connectivity into the world, and typically doesn't want to know about it. It simply isn't in the scope of its behavior. This means that UPA-only behavior is likely to be the only supported more of operation.

Exchange point will typically only allow networks connecting to it to create connections. Such networks will often be transit networks. This in turn means that end-users will not be authorized to setup connections, and either chain the connection through their transit network, or obtain an authorization token out of band.

As an interconnect on an exchange is a bilateral agreement, the exchange NSA is likely to require an authorization attribute from each side to setup the interconnect. Allowing a single side to setup an interconnect unauthorized, could potentially conflict with already used or reserved resources. It is likely that setting up a cross connect on an open exchange will require some coordination between the two NSAs.