TOTALTELECOM
daily news for global communications

Overcoming the stumbling blocks to NNI agreements

By Shankar Narayanaswamy, Reliance Globalcom
27 February 2008

Carriers can benefit from creating new types of network-to-network interface agreements, despite lack of standards.

As the global Ethernet services market grows, carriers and service providers are expanding their effective footprints by building network-to-network interfaces (NNIs). A well-designed NNI is critical because it enables easy integration with other carrier Ethernet networks, thus driving seamless communication across borders. For example, a U.S.-based carrier and a European carrier can leverage each others' networks to serve their respective customers better and more cost-effectively through an NNI than through more traditional means. Similarly, any Ethernet carrier can partner with selected carriers to build effective NNIs, establishing a global footprint with minimum capital and operational cost.

The stumbling blocks

The actual implementation of an NNI faces several stumbling blocks due to the lack of standards specifying how to interconnect two carrier Ethernet networks, although the Metro Ethernet Forum (MEF) is working on this. Without interconnection standards, each carrier NNI is unique and must be jointly negotiated and designed by both sides. This poses several problems.

The first problem is operational scalability. A carrier with multiple unique NNIs to partner networks must train its operations personnel on all of the unique designs. There is no standard troubleshooting procedure. The delineation of responsibilities between the respective carriers' operations teams is unique to each NNI design and can potentially impede the ability to resolve customers' issues.

The second problem is mismatched features and requirements. For example, one carrier may support five different classes of service (CoSes) while another may support only three. Unless the carriers have the same number of CoSes, each with identical parameters, it is very difficult to translate services between carriers at the NNI in a way that meets the end customer's service level agreement (SLA). The MEF is developing a standard CoS implementation agreement template for use between carriers, but carriers need solutions today.

The third problem is mismatched inter-carrier logical and physical interfaces. A strand of fibre or copper can transport either native Ethernet or TDM services encapsulating Ethernet. In addition, some carriers are comfortable with IEEE 802.1Q- based Q-in-Q encapsulation, whereas others do not support this technology. Unless both carriers agree on the physical media and the protocols to be exchanged, it is not possible for them to interoperate.

The fourth problem is service quality expectation. Although this is not an interface issue, it is relevant to effectively servicing end customers' needs. For example, carriers experience various levels of oversubscription, which affects congestion and packet loss. Carriers may support different maximum transmission unit (MTU) sizes, which may cause Ethernet frames to be dropped if they exceed this parameter. Some carriers may prefer geographically redundant handoffs, whereas others think a single handoff is just fine.

The light at the end of the tunnel

Despite the many problems caused by lack of interconnection standards, there is a small universe of NNI designs which can cover 90% of the interconnection requirements. Reducing the NNI design universe to a small number simplifies operational scalability in the same way as a standard. The problems of mismatched features and carrier interfaces are addressed by the specific NNI designs used in any implementation. From a technical perspective, these designs are distinguished by VLAN encapsulation, aggregation, and redundancy.

There are some common threads in how carriers prefer to design their NNIs. For example, copper Gigabit Ethernet handoffs are by far the most common physical interface. Some carriers prefer fibre Gigabit Ethernet handoffs, which is a minor variation of the design. This reduces the physical interface issue to a simple parameter in a design template.

Another common thread is that very few carriers offer more than four classes of service. Furthermore, carriers typically specify CoS behaviour based on strict priority only. This is another de-facto standard which simplifies NNI template designs by allowing very simple mapping of the CoS bits between carriers.

One significant difference in NNI designs revolves around the way carriers carry traffic and the logical interfaces they expose to NNIs: whether they can encapsulate traffic by tag stacking or tag swapping. Some carriers are able to transport Ethernet frames transparently by stacking their own tag on top of their partners' tags, whereas others must swap tags due to MTU size or equipment limitations. We therefore need a major variation of the NNI design template for each of these two cases.

Another significant difference in design templates is aggregation. Some carriers are able to aggregate multiple customers' traffic onto a single Gigabit Ethernet or 10-Gigabit Ethernet handoff, whereas other carriers' equipment limitations require separate physical handoffs for each end customer. Customized NNI design templates are required for each case.

The third and final significant difference in designs is redundancy. Some carriers are able to provide redundant handoffs in multiple locations, whereas others provide either single-location redundancy or no redundancy at all. By designing NNI templates to account for each of these three variations, it is possible to cover all these cases.

Templates speed NNI turn-up

It is by now apparent that there are only three major architectural parameters distinguishing NNI design templates: encapsulation, aggregation and redundancy. This lends itself to a small universe of NNI designs that can be tweaked for each application. Testing each variant is relatively simple, and the technical risk is low since most of the functionality is already tested in the base template.

The one stumbling block which is not entirely solved via NNI design is service quality expectation. This must be negotiated by the carriers. Standardization can help by defining a set of network performance metrics which carriers publish with the NNI guidelines. With all NNI design variants accommodated within a limited set of templates, carriers no longer have to waste time negotiating unique NNIs for every relationship with another Ethernet carrier. Instead, they can quickly choose the template the works for a particular relationship, thus speeding NNI turn-up and enabling a flexible, scalable service.

Shankar Narayanaswamy is vice president of network architecture at Reliance Globalcom.