Connection-oriented Ethernet

Connection-oriented Ethernet refers to the transformation of Ethernet, a connectionless communication system by design, into a connection-oriented system. The aim of connection-oriented Ethernet is to create a networking technology that combines the flexibility and cost-efficiency of Ethernet with the reliability of connection-oriented protocols. Connection-oriented Ethernet is used in commercial carrier grade networks.

Traditional carrier networks deliver services at very high availability. Packet-switched networks are different, as they offer services based on statistical multiplexing. Moreover, packet transport equipment, which makes up the machinery of data networking, leaves most of the carrier-grade qualities such as quality of service, routing, provisioning, and security, to be realized by packet processing. Addressing these needs in a cost-efficient way is a challenge for packet-based technologies.

The IP-MPLS approach aims at providing guaranteed services over the Internet Protocol using a multitude of networking protocols to create, maintain and handle packet data streams. While this approach solves the problem, it inevitably also creates a great deal of complexity.

This has resulted in the emergence of connection-oriented Ethernet which includes a variety of methodologies to utilize Ethernet for the same functionalities otherwise based on extensive IP protocols. The challenge of carrier Ethernet is to add carrier-grade functionality to Ethernet equipment without losing the cost-effectiveness and simplicity that makes it attractive in the first place. To meet this challenge, common connection-oriented Ethernet solutions have chosen to rid themselves of the complex parts of packet transport to achieve stability and control. Key connection-oriented Ethernet technologies used to achieve this include mainly IEEE 802.1ah, Provider Backbone Transport and MPLS-TP, and formerly T-MPLS.

Contents

PBT and PBB

Provider Backbone Transport (PBT) is connection-oriented switch operation scheme and network management architecture. PBT was invented by British Telecom (BT) and developed by Nortel (now Avaya). It defines methods to emulate connection-oriented networks by providing "nailed-up" trunks through a packet-switched network. Key data-plane differences from PBB include the static configuration of forwarding tables within Ethernet switches, dropping of multicast packets and the prevention of "flooding" of frames to unknown destination addresses. Configuration is performed by a centralized

management server like in SDH networks, though in the future a control plane may be added. PBT has been presented to IEEE802 and a new project has been approved to standardize it under the name of Provider Backbone Bridge Traffic Engineering (PBB-TE) (IEEE 802.1Qay), a modification to PBB.

Provider Backbone Bridges (PBB) is an Ethernet data-plane technology invented in 2004 by Nortel Networks (now Avaya). It is sometimes known as MAC-in-MAC because it involves encapsulating an Ethernet datagram inside another one with new source and destination addresses (termed B-SA and B-DA). IEEE802 is standardizing the technology as (IEEE 802.1ah), currently under development. PBB is the original data-plane chosen by British Telecom for their new PBT-based Ethernet transport.

PBB can support point-to-point, point-to-multipoint and multipoint-to-multipoint networks. PBT focuses on point-to-point connectivity, and may be capable of extension to point-to-multipoint, a key technology for advanced data applications such as IPTV. PBT avoids trying to address multipoint-to-multipoint networking, as in the opinion of some of its supporters guaranteed levels of service in multipoint-to-multipoint networks are impossible.

Additionally Ethernet is being reinforced with operations, administration, and maintenance (OAM) capabilities through the work of various standard bodies (IEEE 802.1ag, ITU-T Y.1731 and G.8021, IEEE 802.3ah).

PBT/PBB equipment leverages economies of scale inherent in Ethernet, promising about 30%–40% cheaper solutions compared to T-MPLS equipment with identical features and capabilities,[1] making PBT a better overall return on investment.[2]

T-MPLS (Transport MPLS) / MPLS-TP (MPLS Transport Profile)

T-MPLS, as its name implies, is a derivative of MPLS that renounces all MPLS signaling features and, like PBT, uses a centralized control-plane to perform routing and traffic engineering. T-MPLS is currently being standardized only at ITU-T and enjoys strong vendor support but little carrier support.

As a native MPLS derivative, T-MPLS can be easily implemented over existing MPLS routers. However, T-MPLS has been stripped of the characteristics which originally made it attractive to carriers—control-plane automation, signaling, and QoS—and therefore has yet to prove its benefits for the transport network. T-MPLS OAM, defined in ITU Y.1711, is different from MPLS OAM and lacks powerful management tools that carriers typically expect. T MPLS was abandoned by the ITU-T in favor of MPLS-TP in December 2008.[3]

MPLS-TP or MPLS Transport Profile is a profile of MPLS developed in cooperation between ITU-T and IETF since 2008 as a connection-oriented packet-switched (CO-PS) extension. Based on the same architectural principles of layered networking that are used in longstanding transport network technologies like SDH, SONET, and OTN, MPLS-TP provides a reliable packet-based L2 technology that is comparable to circuit-based transport networking, and thus aligned with current organizational processes and large-scale work procedures similar to other packet transport technologies.

Achieving the promise of carrier-grade Ethernet

Services in the data network are typically classified into 2 major categories: Committed Information Rate (CIR) and Excess Information Rate (EIR). A CIR service guarantees its user a fixed amount of bandwidth, whereas an EIR service offers best-effort only transport. Both types of services share a single capacity-constrained infrastructure. Both are further defined by additional parameters.

A carrier's return on investment is directly related to its ability to transport more service instances over a fixed capacity-constrained infrastructure, keeping Quality of Service high. It is further associated with its ability to offer a broad range of added-value services, such as IPTV, Voice, and VPN, whose requirements can widely vary and pose technical difficulties when sharing the same infrastructure.

With the above in mind, the carrier's objective is to offer a maximum amount of best-effort EIR services over its network while reliably serving its committed CIR services. To achieve this PBB/PBT and T-MPLS approaches largely under-provision network resources, in order to avoid a situation where a burst in best-effort traffic would jeopardize the ability to serve committed traffic, leading to costly penalties. An additional issue with best-effort access on data networks is fair allocation among clients. With PBB/PBT and T-MPLS, the amount of bandwidth available to a particular client greatly depends on the client's location and the prevailing traffic conditions. This limits the value customers attach to EIR services and undercuts carriers' opportunities to offer differentiated access to its excess capacity.

Performing traffic engineering in real-time is thus key to next-generation Ethernet transport. Additional qualities are required to make Ethernet a carrier-grade technology:

- Rich OAM: capacity optimization, path calculation and configuration, iterative optimization.
- Path protection: sub-50ms failover.
- Support for next-generation services: efficient provisioning of point-to-multipoint and multipoint-to-multipoint services.

Multi-vendor support, the ability to support a variety of Ethernet switches in the core, is a desirable attribute as it allows carriers to use inexpensive switches to build their metro transport network. Vendors such as Tejas Networks, Ethos Networks, and Nortel offer solutions which meet the above requirements, yet preserve Ethernet's simplicity and flexibility.

See also

- Ethernet in the first mile
- Metro Ethernet

References

1. The T-MPLS vs. PBT debate (http://www.networkworld.com/supp/2007/ndc4/061807-tmpls-vs-pbt.html)
2. Traffic engineering for Ethernet: PBT vs. T-MPLS (http://www.lightwaveonline.com/featured-articles/traffic-engineering-for-ethernet-pbt-vs-t-mpls-53438082.html)
3. - ITU-T newslog, 10 March 2009 (http://www.itu.int/ITU-T/newslog/Clarifications+On+MPLS+For+Transport+Network.aspx)

External links

- Nortel's PBB/PBT page (https://web.archive.org/web/20070129051009/http://www2.nortel.com/go/solution_content.jsp?segId=0&catId=0&parId=0&prod_id=55120)
- Tejas Networks' page (https://web.archive.org/web/20080314053108/http://www.tejasnetworks.com/technology/tech_carrier-ether.shtml)
- IEEE 802.1Qay (http://www.ieee802.org/1/pages/802.1ay.html)

This page was last edited on 12 February 2018, at 16:51 (UTC).