

Ethernet as Carrier Transport Infrastructure

David Allan and Nigel Bragg, Nortel

Alan McGuire and Andy Reid, BT

ABSTRACT

New innovations in the Ethernet space promise to significantly enhance both the scalability and capability of Ethernet as a networking technology. This article outlines how the convergence of hierarchy, OAM functionality, and enhanced forwarding capability combine to permit Ethernet to assume a much larger role in carrier networks with substantial economic and operational benefits.

INTRODUCTION

Ethernet-based technology has become ubiquitous in both the enterprise and home broadband arenas. The combination of simplicity and rigorous specification has permitted a degree of integration and commoditization that other networking technologies have been unable to achieve.

On the other hand, the current service-provider infrastructure is based on a legacy circuit-based infrastructure, and private line services are the basis for the majority of frame relay, ATM, and IP services and interconnect. This has placed service providers in a difficult position, as they face both the costs of supporting multiple technologies and a service arbitrage situation — they sell the same service on multiple technology platforms.

Ethernet is the technology of choice in the customer domain and is therefore a desirable choice in the service-provider domain to eliminate potential interworking problems and leverage the customer-driven investment. However, every technology transformation in the service-provider space is time-consuming and also represents a major commitment; consequently, comprehensive functionality is required as a prerequisite to mass deployment. From a carrier's perspective, Ethernet still has deficiencies with respect to OAM, reliability, traffic management, and scalability.

It turns out that many of the fundamental issues with Ethernet are well understood, and are currently being addressed with the same rigor and drive for simplicity that has been the hallmark of Ethernet to date. This article delves into some of the challenges faced, and how existing Ethernet behaviors can be combined with

standards in progress to provide a comprehensive network infrastructure that will address the carrier's concerns.

After a summary of the challenges to Ethernet, the remainder of this article is structured as follows:

- The role of hierarchy
- New Ethernet forwarding modes
- Robustness and OAM
- Engineering and resiliency
- Security

We end the article by concluding that Ethernet is capable of taking on a transport networking role.

CHALLENGES TO ETHERNET

Ethernet has faced a number of scaling challenges due to its nature. The combination of a flat address plan in the form of MAC addresses, the use of broadcast and multicast auto-discovery mechanisms, and limited ability to leverage the breadth of connectivity in the network (driven by the need for a loop-free topology) has placed limits on the upper bounds of sizing an Ethernet network. Previous attempts to address this have been via mechanisms to partition or segment the network into multiple "loop-free" domains. This was initially achieved via the use of Virtual LANs or VLANs, a technique primarily of utility to the enterprise. Extensions to VLAN tagging specified under the provider bridging (IEEE 802.1ad [1]) effort went further in allowing providers to impose their own partitioning without disturbing customer partitioning via stacking of the VLAN tags, where the inner tag is referred to as the customer VLAN (C-VLAN) and the outer tag as the service tag or (S-VLAN). However, this was done via stacking identifiers designed primarily for single Enterprise-scale requirements. In a carrier context this VLAN tag stacking involves overloading a transport separator function with a customer-identification function. The service tag identified both a community of interest and a specific topology within the provider network. The net consequence is the instantiation of service specific state in the Ethernet core, as each individual group of interest (identified by an 802.1ad S-VID) is both mapped to a spanning tree instance, and identifies an individual service instance.

Ultimately the applicability of 802.1ad is limited, as this does not address the fundamental scaling issues associated with broadcast domains and flat addressing, and offers no real solutions for true virtualization of services over a common Ethernet based infrastructure. This is because VLANs only partition the network into the equivalent of closed user groups (CUGs), in essence adding state to the core with a CUG identifier designed for the Enterprise and not for carriers (only 4094 VLAN IDs). This has scaling limitations when contrasted with a hierarchical or client/server layer relationship, where the client and server layer networks are functionally decoupled. Nor does simple partitioning provide any real control over how a carrier's facilities are loaded; the use of multiple spanning trees only increases the degree of indeterminism in network operations.

To date the only real solution for carriers has been to severely limit the scale and scope of applicability of Ethernet networks in a service context, and/or to use it as server-layer interconnect for other (higher) layer networks, such as IP.

HIERARCHY AS THE FUNDAMENTAL BUILDING BLOCK

With the advent of work on 802.1ah [2] provider backbone bridges (also known as MAC-in-MAC), Ethernet gains the tools to permit true hierarchical scaling, virtualization, and full isolation of provider infrastructure from customer broadcast domains. This is a significant step in making Ethernet suitable for carriers. The use of hierarchy is a well-understood mechanism for achieving scalability and security.

Hierarchy has particular utility when the customer base consists of a large number of relatively small communities of interest (the primary situation that will be faced by carriers) that can be overlaid upon a common transport network. It reduces the amount of provisioning and forwarding state in the network core and correspondingly reduces the load and ongoing cost of performing service assurance and fault management.

Figure 1 illustrates the evolution of Ethernet hierarchy as specified by the IEEE. Whereas the original 802.1q (now 802.1Q [3]) and subsequent 802.1ad simply partitioned the forwarding plane of a common single-layer network, 802.1ah implies complete recursion such that customer Ethernet subnetworks are completely encapsulated and isolated from the provider Ethernet network. The server layer, unless modified, will still inherit limitations associated with autolearning and the associated requirement for a loop-free topology (subsequently referred to as bridging). The remainder of this article will focus on addressing these limitations.

NEW FORWARDING MODES

Hierarchical isolation of customers from carrier operations permits the carrier to use different forwarding modes and gain determinism and predictability in operations via the ability to

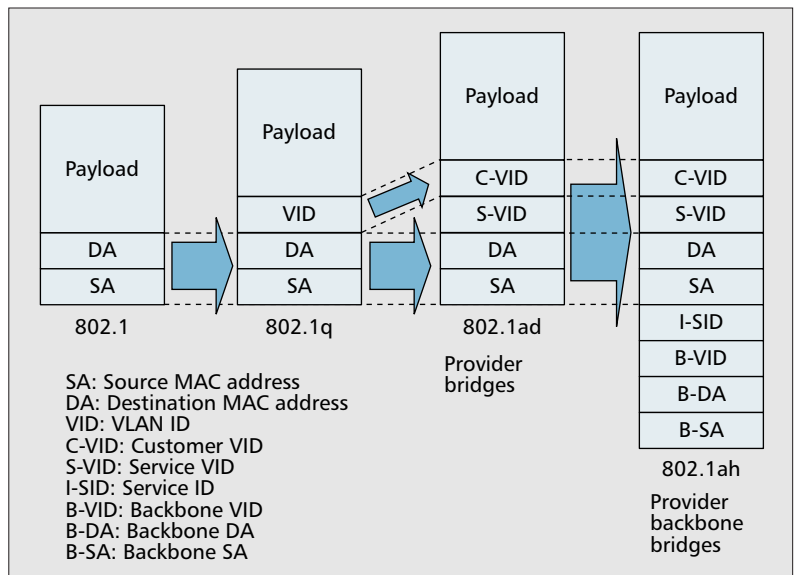
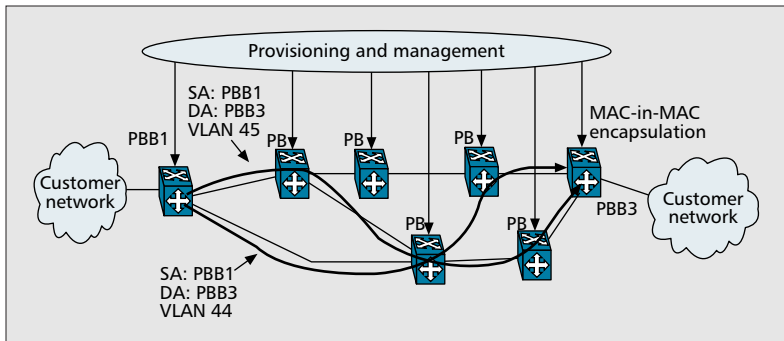


Figure 1. Evolution of Ethernet hierarchy.

engineer the network. This is a critical requirement for carriers so that they can exercise deterministic traffic/fault/performance management of customer services. Whereas Ethernet networks were previously limited to trivial physically loop-free topologies or limitations imposed by spanning trees, alternative provisioning and control-plane options now allow the Ethernet infrastructure to be engineered. As the carrier is in complete control of their own MAC and VLAN space, and is not obliged to share it with their customers, latent abilities hidden within existing Ethernet hardware can be exploited to significant advantage.

One useful carrier requirement not addressed to date is the ability to scaleably and simply set up and pin connection-oriented packet trunks across an Ethernet network. Configuration of VLAN tags as connection-forwarding entries has been proposed, but the use of global tags imposes a network limitation of 4094 connections. Per port translation of VLAN tags (as was most recently incorporated into 802.1ad) would alter the scaling attributes to 4094 VLAN IDs per platform, but for P2P connections this is still a very limited quantity and inferior to first-generation ATM switches. As VLANs ultimately are designed to implement multicast spanning trees, it makes sense to explore alternative mechanisms to scale the connection space instead of consuming a scarce resource not designed for the purpose. Fortunately, there is an alternative...

Independent VLAN learning (IVL) Ethernet switches forward on the basis of a full 60 bit lookup of both the VLAN tag and the destination MAC address in each packet. That this results in connectionless behavior is a consequence of the MAC learning mechanism, and other more deterministic forwarding modes are possible with the same hardware simply by turning some bridging functions off using controls (the majority of which already exist in standards). Bridging populates the forwarding tables in Ethernet switches via well known MAC learning



■ **Figure 2.** Example of configuring an Ethernet switched path.

procedures and flooding, but switches also explicitly allow for control plane or management configuration of the forwarding tables as well.

It is important and useful to preserve the Ethernet paradigm of globally unique MAC addresses, but it is not necessary to maintain global uniqueness for the entire VLAN ID (VID) range. The fundamental unicast forwarding behavior of IVL switches does not require VID to be globally significant; the requirement for global VID comes only from the need to constrain broadcast behavior. Therefore we can allocate a range of VID (say “n” VID) as only locally significant to a given MAC addresses. It is only a smaller step to then consider a VID in that range as an individual instance identifier for one of a maximum of “n” connections terminating at the given MAC address. The combination of VID and MAC is globally unique, with the MAC address identifying the logical administrative owner of the specific VID value.

We can then use a separate control or management plane to populate the switch forwarding tables for the designated VID range [4] being an exemplar of such a control plane. When we configure a VID/MAC-tuple into a contiguous sequence of IVL capable switches, this will create a unidirectional connection. It is again a small leap to combine mirror image connection pairs to create bi-directional connectivity. At this point we have the ability to define Ethernet switched/engineered paths side-by-side with traditional 802.1D [5] bridged behavior.

Figure 2 provides an example of configuring two distinctly routed paths using this technique. In Fig. 1, PBB1 and PBB3 correspond to the network edge and are 802.1ah compliant devices that offer customer facing ports. PBB1 and PBB3 adapt customer traffic onto configured Ethernet switched paths. In this example, B-VIDs 44 and 45 fall into the range set aside for configured behavior.

Two paths have been configured from PBB1 to PBB3, so that the MAC address of PBB3 will be the MAC component of the forwarding table entry for the paths. For the first, a path has been computed, B-VID 44 has been assigned, and the forwarding tables have been configured in the intervening PB switches mapping B-VID = 44/MAC = PBB3 to the appropriate egress ports of each device to define a contiguous path. For the second path, the same process resulted in a path configured in the switches using B-VID = 45/MAC = PBB3. Via a similar process, sym-

metrical return paths from PBB3 to PBB1 would also be configured.¹

In the example illustrated here, the paths deliberately cross to show that it is the combination of both B-VID and MAC that provide the unique forwarding entry. It is the concatenation of the two that determines the forwarding path. Collisions in either space such as B-VID 44 or 45 used in conjunction with another MAC address or as in the example above where paths 44/PBB3 and 45/PBB3 cross are still uniquely resolved to a single egress port.

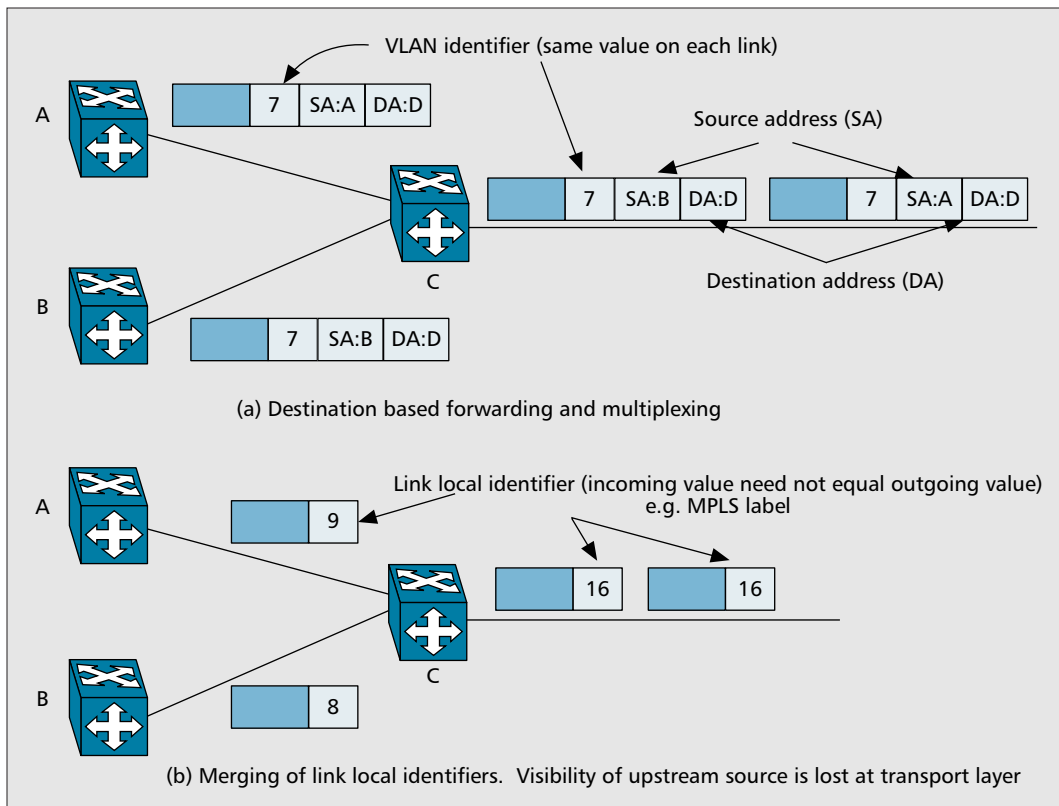
This technique provides the interesting combination of both a destination-based forwarding paradigm and multiplexing. When multiple sources are configured to send to the same VLAN/Destination MAC-tuple, knowledge of source is preserved in the form of the source MAC address and this is present for both bearer and data plane OAM PDUs. This differs from other “lossy” MP2P approaches such as merging, as illustrated in Fig. 3.

Preservation of knowledge of the source of a packet has significant implications for fault and performance management. When considering the quality of information available on the data plane, an Ethernet Switched Path is a multiplex of P2P connections that scales O(N) at transit switches while preserving the useful OAM properties of a full mesh at the terminating device. A path-terminating interface maintains fault and performance management state in proportion to the number of connected peers, and is able to disambiguate the source and path of both bearer traffic and OAM PDUs on the basis of the embedded source MAC address, and VLAN identifier. This is a distinct improvement on the “merge” example outlined above whereby additional information is added to the OAM PDUs in order to identify the source, and actual bearer traffic cannot be associated with a specific source device. What this means is that true traffic counts and availability state can be maintained on a per path pairwise basis and real-time correlation of performance counts is possible.

Replacing traditional flooding and learning with configuration requires a few extra steps to make the technique fully robust:

- Discontinuities in the forwarding-table configuration in the path of a connection will normally result in packets being flooded as “unknown.” This is potentially catastrophic in a meshed topology, so for this reason unknown flooding must be disabled for the designated VID range. This is similarly disabled for broadcast/multicast traffic.
- MAC learning is not required, and may interfere with management/control population of the forwarding tables. For this reason MAC learning is disabled for the delegated VID range.
- This approach does not need a loop-free topology for the delegated VID range, so spanning tree is disabled for the delegated VID range.
- The application of connection admission control (CAC) to connection traffic requires a higher 802.1p priority than STP best effort and in the engineering the network, a reserve must be set aside for STP traffic.

¹ Asymmetrical routing of paths is not precluded and in many cases can be considered a scalability adjunct, as there are fewer constraints to consider when routing a given path. However, many applications benefit from symmetrical routing and, in the case of an Ethernet client layer, are largely dependent on fate sharing of both directions of connectivity for robustness. Asymmetrical failures in Ethernet are catastrophic for spanning tree algorithms to the point where much care has gone into Ethernet standards to ensure that failures are bidirectional. Configured paths do not use spanning tree, but 802.1ah client layers frequently will.



Delegating a small portion of the VID range to configured rather than learned behavior has a trivial effect on the scaling properties of VLAN partitioned bridging while simultaneously permitting a large number of p2p or mp2p Ethernet switched paths to be supported.

Figure 3. Source visibility for a destination-based forwarding and multiplexing solution vs. merging.

This is discussed further below in the section titled Engineering and Resiliency.

- Data plane OAM provides both fault and performance management. As the actual packet transfer function across an Ethernet switch is unmodified (either learned or configured), the OAM currently defined can also be reused without modification of the protocols, just slightly altered semantics. In particular, Y.17ethoam includes unicast PDUs for fault management, which will directly exercise and share fate with a configured path identically to how they would exercise a learned one. What is exercised is the forwarding based on the table contents (independent of the method of table population).

Delegating a small portion of the VID range to configured rather than learned behavior has a trivial effect on the scaling properties of VLAN partitioned bridging while simultaneously permitting a large number of p2p or mp2p Ethernet switched paths to be supported. Setting aside as few as 16 VID values permits 16 uniquely routed mp2p multiplexed paths to the full set of destination PBB devices in the network, while leaving 4078 VID values for bridged behavior. This offers a theoretical network maximum of some 2^{52} ESPs fully meshing (with resiliency) 2^{48} devices. The theoretical data-plane address space significantly exceeds the capacity of existing implementations; however, when used in this mode, existing implementations exceed the scaling requirements of deployed networks.

In contrast to the configuration with flooding and learning, devices are only configured with

knowledge of peers on a “need to know” basis. Most metro architectures are a hybrid of “low-touch” p2p private line and p2p backhaul to either hub switching nodes or “high-touch” layer 3 service-edge devices. The implications of such architectures are that an edge device typically only requires data-plane state in direct proportion to the number of customer facing ports (frequently under 100) and not in linear proportion to the number of peer devices (which in a national scale network can reach 30,000 or more devices). This has implications not only with respect to scaling limitations using current technology but also with respect to content addressable memory requirements, powering requirements, device footprint, and ability to be deployed in outside plant.

ROBUSTNESS AND OAM

Ethernet offers an inherently robust data plane when contrasted with other packet technologies. Use of link local-path identifiers such as MPLS labels, ATM VCI/VPIs, and so forth introduces a “level of indirection” into data-plane forwarding. The current fascination with label swapping that has persisted for several generations of WAN technology actually has a detrimental effect on overall network reliability.

Successful forwarding of a given packet is dependent on a “chain of correctness” of link local relative identifiers across the network. Not only do link local identifiers need to be administered and configured, but the intermediate switch transfer functions need to tie link local identifiers together. This usually takes the form

The configured mode of operation only uses a subset of bridge behavior such that only a subset of the currently defined OAM procedures will be sufficient to instrument configured paths. All that is required is an end system design for implementation subtleties with respect to the number of OAM flows both originated and terminated.

of both a lookup, which identifies an egress port, and a new link local value with which to overwrite the existing value in the packet. It is possible for the consequences of misconfiguration of the transfer function or software problems to be concealed, for example, when a corrupted identifier value collides with one already in use at the current or a nested layer.

This layer of indirection between the path and the chain of identifiers that constitute path forwarding has modes of failure that require explicit monitoring. This has added complexity to OAM tools to perform basic verification of network operation. This was not addressed in ATM and has dominated MPLS OAM tool development (e.g., LSP-PING [6], Y.1711 [7], Y.1713 [8]).

Ethernet uses globally unique interface identifiers in the form of MAC addresses. This means that for connectionless best-effort operation, the true test that “a packet dropped anywhere in the network will find its way home” holds true; the network is thus self-healing (although with degraded performance, as the path is likely less optimal). For configured forwarding, leakage outside the configured path due to a fault or misconfiguration self-identifies as a problem immediately, thus permitting the network to respond. These are desirable attributes, especially so when they come more or less for free.

Significant progress has been made in the field of data-plane OAM with efforts in both the ITU-T (Y.17ethoam [9]) and the IEEE (802.1ag [10]). There is strong agreement as to the functionality required, and a comprehensive suite of tools is already emerging. This is a side benefit of the degree of rigor associated with the specification of Ethernet in terms of frame formats and transfer functions. When the data-plane forwarding and relay behavior is clearly understood and well specified, procedures for verification become a similarly well-understood problem. The end result will be an increasing degree of front-line reliability and much lower mean time to detect and repair when failures happen. It will also increasingly decouple the need for Ethernet networks to depend on transport OAM (e.g., SONET/SDH) or client OAM (e.g., ICMP Ping [11]), neither of which are truly native to the maintenance entity of real interest. In the longer term it will permit additional resiliency options to be added and for comprehensive instrumentation of network performance.

The addition of connection management to Ethernet via reuse of existing data-plane transfer functions means that currently specified OAM protocols will require no modification in order to be applied. Using configuration of existing forwarding addresses fundamental requirements for data-plane OAM and permits reuse of the standards in progress:

- Fault management OAM PDUs (such as continuity check, loopback, link trace, etc.) must utilize and exercise the same forwarding components in intermediate switches as the bearer path. For unicast forwarding based on known or configured VLAN/MAC, the transfer function is common regardless of how the forwarding table was populated.

- For real-time correlation of counts, performance management OAM PDUs must utilize both the forwarding components and specific queuing discipline at intermediate switches so that packets “in flight” can be correctly accounted for. Again, for unicast forwarding based on known or configured VLAN/MAC, the transfer function is common regardless of how the forwarding table was populated.

The configured mode of operation only uses a subset of bridge behavior such that only a subset of the currently defined OAM procedures will be sufficient to instrument configured paths. All that is required is an end system design for implementation subtleties with respect to the number of OAM flows both originated and terminated.

ENGINEERING AND RESILIENCY

Configured Ethernet switched paths have complete route freedom. The ability to define multiple paths to any end system combined with route freedom has a number of implications.

There may be numerous viable paths between any two points in the network, and spanning tree and bridging/autolearning are only able to use one of them. The configuration allows more than one bearer path between any two points to be defined, and criteria beyond simply the shortest path to be used in selecting the routing of any individual path. Multiple metric techniques and graph-splitting algorithms exist to permit optimal sets of paths with specific e2e attributes and without common points of failure to be computed and instantiated. What this means is that:

- Paths can be engineered.
- Capacity mismatches between physical network build and offered load (a frequent occurrence) can be compensated for.
- First line resiliency in the form of protection switching can be delegated to the data plane. Further emerging standards in G.ethps [12] provide for dataplane synchronization of Ethernet protection switching.

When considering a network that utilizes both learning/bridging and configured behavior, the traffic associated with each forwarding behavior needs to be differentiated such that best effort traffic does not degrade ‘engineered’ traffic, especially as the traffic matrix may dynamically change in response to failures and/or maintenance activities. This could be achieved via specifically instantiating “connection” state in transit switches, or can be achieved via intelligent use of class-based queuing and edge policing of class markings.

Ethernet implementations already have class of service capability via the use of 802.1p marking. This is extended in 802.1ad to include discard eligibility that align behaviors more with carrier sensibilities.

This is similar to current-layer 3 QoS mechanisms which are designed around a LAN-style environment where a small amount of QoS-sensitive traffic needs priority over a large bulk of other traffic at a few bottlenecks through an environment of abundant spare

capacity. However, it is desirable to support a much more richly connected network at much higher allocated fills. The result is that packet marking and class-based queuing are in themselves insufficient to offer any acceptable QoS except across a relatively narrow range of network load.

When class-based queuing is combined with Ethernet switched paths, CAC, and admission policing so that there is correspondence between offered load and committed resources, the ability to offer real and measurable guarantees is created. In simplest terms, CAC'd and engineered traffic is given a higher priority class than bridged traffic. In particular, CAC'd connections can be configured to ensure zero congestive packet loss, a key enabler for Ethernet as a convergence vehicle.

SECURITY

Hierarchy has a number of security advantages as well. Isolation of untrusted traffic via encapsulation and containment means that significantly less complexity is required at the network edge to police untrusted behavior, as by definition behavior limitations are rigorously enforced by the connectivity model. It can also be said that autodiscovery (one of Ethernet's strong suits) and security are polar opposites, yet from the point of view of operational complexity, much of Ethernet's autodiscovery capability is still desirable. Hierarchy allows trusted domains to be established that employ autodiscovery; for example, it allows for automated establishment of management connectivity with NEs. The transport layer interconnecting the network's edge devices can operate in a plug and play mode, as the network edge is secured physically. The network edge encapsulates and constrains traffic received from outside the secure boundary.

The addition of connection-oriented engineered trunks further isolates a significant portion of the network from the malicious or the incompetent.

THE TRANSPORT METAPHOR

The ability to configure diversely routed connections in Ethernet combined with OAM provides Ethernet with a strong analogy to the operation of existing ATM and SONET/SDH networks. The server layer in the MAC-in-MAC hierarchy can now offer similar services to these technologies to its own Ethernet client with the advantage that the two layers can be provided by a single platform. This convergence should reduce both operational costs and capital costs. The simplest of these services is Ethernet Private Line, in which the server offers a transparent fixed bandwidth service to the client. With the introduction of engineered connections, the Ethernet server can now offer a transport service with a service-level agreement by which connectivity, latency, loss, and jitter can be specifically monitored and reported. This service can also be offered to clients other than Ethernet (a result of the decoupling between layers that exists in a client/server relationship).

To distinguish this form of Ethernet from varieties that require bridging functions to be enabled, we refer to it as Provider Backbone Transport. Furthermore, the packet-grooming capabilities of Ethernet allow oversubscription in the core, and class-based queuing options exist to more effectively utilize bandwidth. Provider backbone transport may also be used in combination with connectionless Ethernet client layers to provide more advanced services such as hub and spoke applications that use MAC learning or multipoint-to-multipoint services that operate in an overlay that uses Ethernet connections to provide topology.

When operating such a network, it is possible to decouple the details of "service take" from the capacity planning exercise; the old "assign and design" paradigm of network planning can be used in conjunction with the much less operationally intensive "observe and fix" model. This permits carriers to leverage existing skills sets in their operational personnel as they transform their networks from circuit to packet while streamlining the overall service turn up process. It also allows the service suite and service-assurance process to be direct extensions of existing operations.

Operations can be simplified via the transport paradigm, as strict hierarchy (inheritance between client and server layers) combined with well-specified data-plane OAM allows alarm management/suppression and fault correlation to be pushed down into the network elements. As a result, root-cause analysis and alarm suppression become an inherent part of data-plane operations, rather than a bolt on afterthought. Further, hierarchy typically means that connectivity to support many service instances is almost fully congruent with a given PBT trunk. This reduces the need for OAM "chattiness" at the service layer. The availability and performance attributes of the trunk are directly inherited by the services.

CONCLUSION

Traditionally Ethernet has been weak in a number of areas, which has limited the utility of Ethernet to carriers and limited the application of Ethernet technology as infrastructure. These have been deficiencies in the areas of scaling, virtualization, security, OAM, and determinism.

The current round of Ethernet standardization largely addresses this list of outstanding issues, and the addition of the ability to engineer packet trunks across an Ethernet network provides a comprehensive infrastructure solution. The fundamental innovations are the addition of hierarchy, and the consequent ability to use, instrument, and engineer connections in a simple and efficient manner.

So the end result is that, for carriers, Ethernet is ready to substantially expand the range of applications it can address — not just as a link technology, not just as a customer presentation, no longer needing OAM assistance from its neighboring layers — but as a full fledged and relatively complete and "fit for purpose" transport layer in its own right.

Hierarchy typically means that connectivity to support many service instances is almost fully congruent with a given PBT trunk. This reduces the need for OAM "chattiness" at the service layer. The availability and performance attributes of the trunk are directly inherited by the services.

For carriers, Ethernet is ready to substantially expand the range of applications it can address — not just as a link technology, not just as a customer presentation, no longer needing OAM assistance from its neighboring layers — but as a full fledged and relatively complete and “fit for purpose” transport layer in its own right.

ACKNOWLEDGMENTS

The authors would like to thank Neil Jefferies, Neil Harrison, and Brian McIntosh of BT, and Simon Parry and Robert Friskney of Nortel.

REFERENCES

- [1] IEEE 802.1ad, “IEEE Draft Standard for Local and Metropolitan Area Networks, Virtual Bridged Local Area Networks, Amendment 4: Provider Bridges.”
- [2] IEEE 802.1ah, “IEEE Draft Standard for Local and Metropolitan Area Networks, Virtual Bridged Local Area Networks, Amendment 6: Provider Backbone Bridges.”
- [3] IEEE Std. 802.1Q, “Local and Metropolitan Area Networks, Virtual Bridged Local Area Networks,” 2003.
- [4] IETF Internet draft, “GMPLS Control of Ethernet I/VL Switches,” Oct. 2005 (work in progress), available at draft-fedyk-gmpls-ethernet-ivl-00.txt
- [5] IEEE Std. 802.1D, “Local and Metropolitan Area Networks, Media Access Control (MAC) Bridges,” 2004.
- [6] IETF Internet draft, “Detecting MPLS Data Plane Failures,” May 2005 (work in progress), available at draft-ietf-mpls-lsp-ping-09.txt
- [7] ITU-T Rec. Y.1711(2004), “Operation and Maintenance Mechanism for MPLS Networks.”
- [8] ITU-T Rec. Y.1713(2004), “Misbranching Detection for MPLS Networks.”
- [9] ITU-T Draft Rec. Y.17ethoam (Y.1731), “OAM Functions and Mechanisms for Ethernet Based Networks.”
- [10] IEEE Draft Std. 802.1ag, “Local and Metropolitan Area Networks, Virtual Bridged Local Area Networks, Amendment 5: Connectivity Fault Management.”
- [11] J. Postel, “Internet Control Message Protocol,” IETF RFC 792, Sept. 1981.
- [12] ITU-T Draft Rec. G.ethps (G.8031).

BIOGRAPHIES

DAVID ALLAN (dallan@nortel.com) is a graduate of Carleton University, Ottawa, Canada (B.Eng. 1978). He has been active in data telecommunications standards for the past 10 years, including WG chair roles in the DSL Forum and IETF. He has been active for more than 25 years as an architect, design engineer, and developer of real-time systems in diverse areas of technology ranging from process

control and avionics to financial transaction processing. His current role at Nortel is focused on MPLS and Ethernet standards.

NIGEL BRAGG holds an M.A. degree from Trinity College, Cambridge University, and an M.Sc. degree from Southampton University. He has spent 20 years in the telecommunications industry, initially in a contract product development environment, and for the last 10 years with Nortel Networks. He has contributed to a wide variety of projects during that time, ranging from design leadership on a high-performance voice switching system, to many aspects of data switching and routing, to automated power optimization in optical line systems. Most recently, he has been focusing on the requirements for and realization of multiservice packet transport for the carrier environment.

ALAN MCGUIRE [M] graduated from the University of St. Andrews, Scotland, in 1987 with a First Class Honours degree in physics and an M.Sc. degree in biomedical physics from the University of Aberdeen, Scotland, in 1988. Since joining BT he has been involved in network architecture, statistical networking, optical networks, SDH, network management, and control plane technologies. He currently manages a team responsible for current and next-generation Ethernet systems. He has represented BT in a number of standards bodies and has been editor of a number of ITU-T Recommendations, and was a prime contributor in the development of the ASON architecture. He was a Guest Editor for *Journal of Selected Areas in Communications*’ Special Issue on Protocols and Architecture for Next-Generation Optical WDM Networks. He is a Chartered Physicist and a member of the IEE and the Institute of Physics.

ANDY REID is currently the chief network services architect in BT’s Group CTO Office and is responsible for defining the way network services are supported on BT’s 21st Century Network. Prior to this, he worked in various roles in technical, product portfolio, and regulatory strategy. His professional interests include telecommunications technology, traffic management and QoS, functional and enterprise architecture, service pricing, microeconomics applied telecommunications, and regulatory/competition law economics. In the past he has made major contributions to the seminal standards development of SONET/SDH and has co-authored two books on the subject.