Optimizing Carrier Ethernet Latency & Bandwidth Efficiency

Executive Summary

Keeping latency in check is critical for quality of service (QoS) of demanding Ethernet wireless backhaul, wholesale and business services. Given that Ethernet bandwidth is considered relatively 'inexpensive', some providers increase throughput to reduce latency and keep performance within strict SLA tolerances.

However, the relationship between bandwidth and latency is anything but simple in packet-based networks, where increasing throughput can sometimes have no effect on delay, or even increase it under certain conditions.

This paper explores the key sources of delay, and techniques to optimize latency without consuming excess bandwidth. These guidelines allow providers to use their network resources as efficiently as possible, while assuring the QoS of critical, real-time services.

Latency's Role in QoS & Sources of Delay

Low latency is key to delivering reliable, high-performance Ethernet business services and backhaul for 3G and 4G wireless networks. Real-time communications, transactional applications, high-speed roaming, and media streaming are all delay-sensitive. Latency increases of just a few milliseconds can result in dropped calls, garbled voice and unresponsive applications, and can mean significant losses in financial trading.

At times, service providers over-provision bandwidth to keep latency and jitter in check. While increasing bandwidth can sometimes reduce latency, it often has little effect. In packet-based networks the relationship between latency and bandwidth is complex and varied. Consider the four main sources of latency, categorized as:

Sources of Latency

- **Serialization Delay:** time required for a port to transmit a packet, related to frame size and bit-rate;
- **Propagation Delay:** limitations imposed by the laws of physics (speed of light, path length, circuit design);
- **Congestion Delay:** the time a frame idles in the output queue of a network element (NE) while a backlog of packets is being transmitted; can be caused by traffic bursts, larger ingress versus egress bandwidth (e.g. oversubscribed aggregation), or network congestion resulting in paused transmission (flow control).
- **Forwarding Delay:** the time required for the NE to analyze, process and forward a packet in a congestion-free scenario; a function of NE architecture and packet-processing requirements (the number and complexity of operations performed on a packet between receipt & transmission, e.g. switching, rate limiting, shaping, etc).

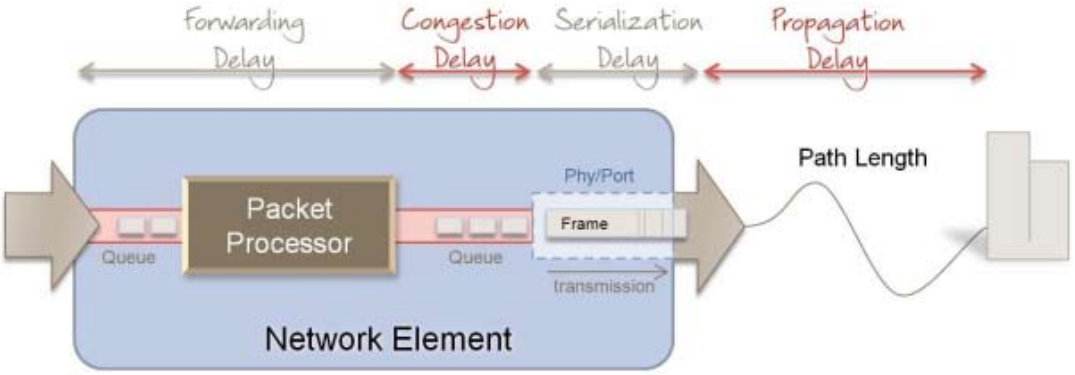

Figure 1: The 4 Components of Delay in Packet-Based Networks

How do NIDs help assure latency and jitter? Can't I just increase bandwidth if there is a problem?

Of these components, serialization delay is the most constant, having only a small influence on end-to-end latency. Propagation delay, typically stable in circuit-switched networks, can be irregular and introduce jitter over routed networks due to path variation; overall, its contribution is usually small, even under heavy utilization.

The more important sources of latency –congestion and forwarding delay –are not entirely independent: as a NE is subject to heavy load (congestion delay), it may need additional queue time to handle and process the increased volume of traffic (forwarding delay). Depending on the NE's design, forwarding delay can be significant when advanced functions such as traffic shaping and multi-flow Ethernet OAM[1] [1]are enabled.

As network congestion can have a large impact on end-to-end latency – impacting both forwarding and pure congestion (queuing-related) delay – reducing traffic bottlenecks is a key part of network management and design. Increasing capacity (available bandwidth) should, at least in theory, help reduce congestion when applied to network "pinch points".

However, increasing throughput does not always lead to the expected decrease in latency, even if congestion is reduced. Results will vary depending on implementation, network architecture, traffic patterns, and a number of other factors explored here.

Bandwidth vs Media Rate

Although Ethernet offers easy access to bandwidth on demand, "dialing up" throughput often has little effect on latency if the link was originally sized correctly. Consider a cell tower connected to an access platform using a 100FX optical link, with a CIR2[2]of 20 Mbps (Figure 2). If tower traffic never exceeds 20 Mbps, increasing the throughput to 30, 50, or even 100 Mbps will have no effect on latency.

Why? Although capacity may have increased, each packet is still bound to the physical link speed negotiated between ports (e.g. 100 Mbps for 100FX media). This means the latency (time of flight) is the same as long as the physical link remains the same.

1 Operations Administration & Maintenance
2 Committed Information Rate, enforced by rate limiting

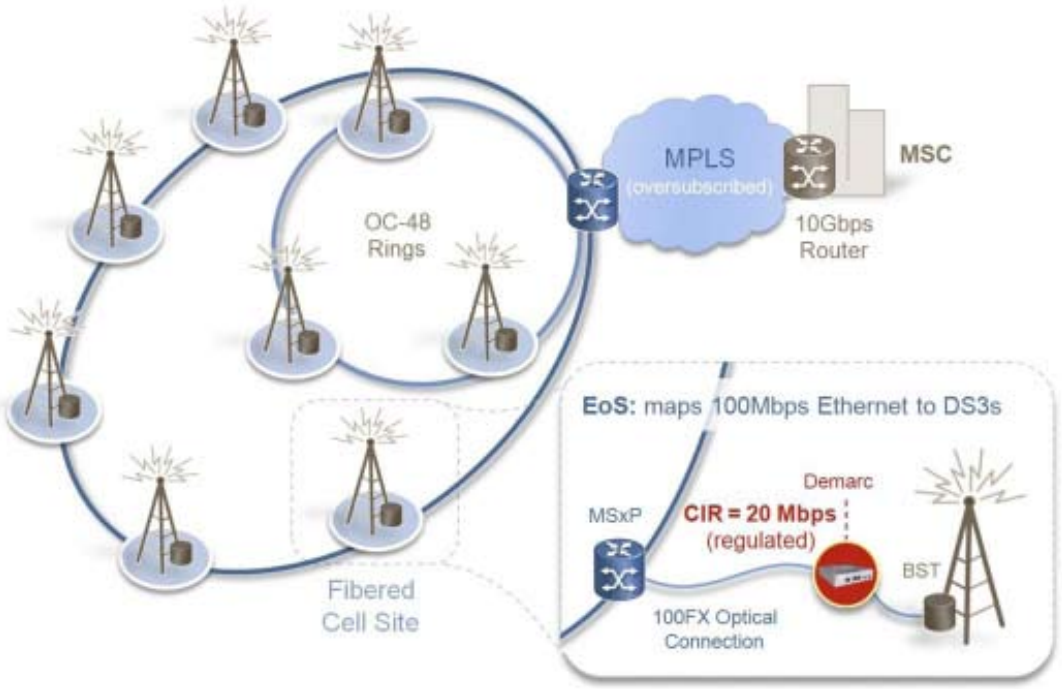

Figure 2: Fibered cell site and common wireless backhaul topology

Dialing-up bandwidth can get more data across a link quicker (a file downloads faster), but each individual packet still travels at the same speed. This is analogous to adding more lanes to a highway without changing the speed limit: unless there is congestion, time to travel from A to B remains the same. Instead of increasing the CIR, the media needs to connect at a faster rate to reduce delay, in this example, changing from the 100FX to a GbE interface would reduce latency ten-fold. In many cases this requires replacing or upgrading network equipment, an expensive and time-consuming option.

Access vs End-to-End Throughput

When congestion is present, increasing bandwidth will reduce latency, but only if throughput is increased end-to-end. For example, if Ethernet-over-SONET / SDH (EoS) is used for transport, increasing the media of the last mile will have little effect on overall latency if packets are still carried over the same TDM "container" (e.g. DS3 / E3). This is analogous to the Ethernet media rate effect: the rate of traffic entering the network has no effect on latency if TDM / core provisioning remains unchanged.

Likewise in all-packet networks (e.g. MPLS core), if the end-to-end network doesn't have sufficient capacity, packets will become congested in the core instead of the access network, simply displacing the problem elsewhere. When this occurs, increasing access bandwidth may result in even longer delays, as already busy core NEs add more traffic to their queues and processing load.

Stability Uncertain

Latency is likely to change over time, even in the best-designed networks. Increased oversubscription, new subscribers, changing traffic patterns and increased service usage can all introduce congestion and delay. Latency can also increase suddenly if a link fails and a protected path is required. In microwave links, inclement weather and radio issues also affect delay. As no service, subscriber base or user behavior is constant, so latency (and jitter) will vary over time.

Flow-Specific Latency

Not all packets are created equal. Critical applications and services typically traverse the network with a higher priority than best-effort traffic, as maintaining low latency and jitter is required for these applications' quality of service (QoS). However, networks do not always respect or maintain a packet's priority.

Some network elements discard or alter the specified class of service (CoS), or re-map Ethernet CoS to the IP equivalent (DSCP), or vice-versa, when bridging layer 2 or 3 network segments. Errors in NE configuration, firmware, and service class mapping rules often result in latency issues, which are commonly introduced when new services or traffic patterns require NEs to be reconfigured.

Optimizing Latency Performance

Making Latency Visible

The causes of latency are complex and non-deterministic. To maintain overall QoS, latency and jitter need to be continuously monitored on a per-service, application, SLA or VLAN-basis. Monitoring latency using coarse port-level or software-based methods (e.g. ping), fail to identify latency issues affecting specific traffic classes, and are unable to isolate whether latency originates at the IP or Ethernet layer.

Likewise, monitoring latency round-trip is insufficient for SLA reporting and troubleshooting delay issues in networks where traffic is often asymmetrical; one-way latency and jitter measurements are required. To provide a complete view of latency performance, measurements should also span the complete service path, end-to-end (the demarcation point and at key nodes in-between). This allows operations staff to isolate latency issues to access, metro, core networks –or perhaps more importantly –the customer's network.

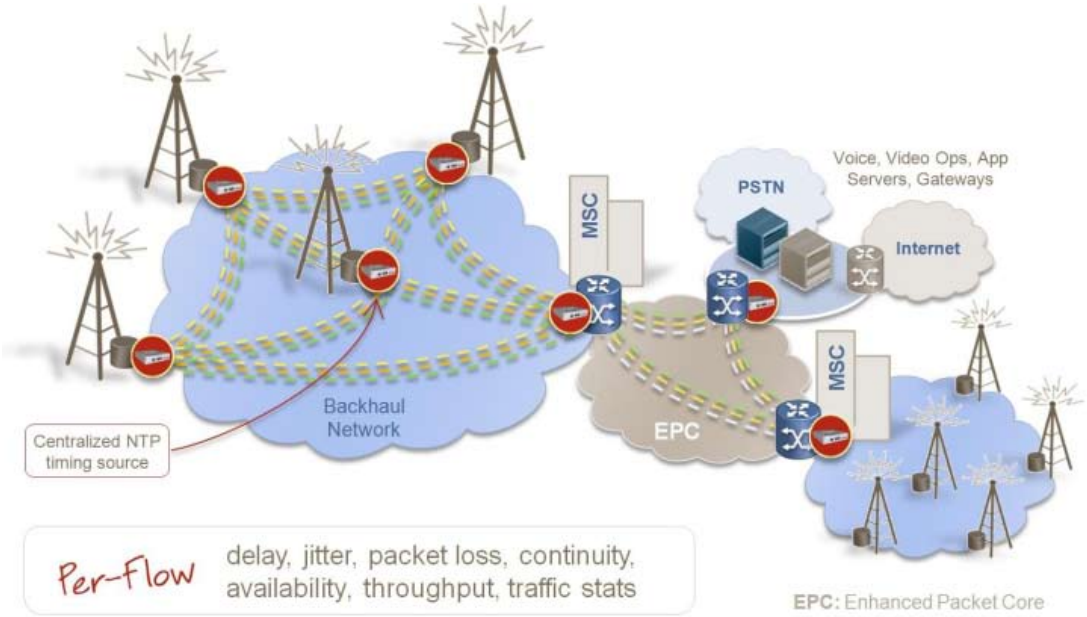

Figure 3: Typical measurements and multi-CoS monitoring paths in a wireless backhaul network

Latency measurements need sufficient precision to detect subtle changes in delay to pro actively identify latency drift that can lead to QoS issues and SLA violations. Precision should be at least one order of magnitude greater than the target threshold to be useful (e.g. sub-millisecond if an SLA specifies a 10ms latency maximum), otherwise measurement error will mask delay issues, create false warnings, or even result in negative latencies in customer SLA reports!

Similarly, measurements need sufficient granularity (frequency) to detect short-term latency problems, which can indicate micro-bursting, ineffective traffic management practices, or mis-configured NEs. The ability to measure each second, for example, provides sufficient information to diagnose intermittent QoS issues.

In summary, to provide complete visibility into latency issues, measurements should:

- Be conducted per-flow, application, service or VLAN, not simply per-port;
- Be conducted one-way, as opposed to round-trip;
- Measure latency end-to-end to accurately account for all sections of the network, as well as to key intermediate nodes for problem isolation;
- Offer sufficient precision to detect gradual shifts in latency, and to provide meaningful alarm and reporting values;
- Provide high granularity if required to troubleshoot transient, delay-related issues.

Validating Congestion

Before trying to correct latency by adding bandwidth, it's important to first determine if and where congestion is occurring in the service path. As provisioning additional bandwidth is time-consuming and 'expensive' from a network perspective, it is best increased where it will most improve performance.

Per-flow packet loss is a good indication of congestion, and can normally be measured along with latency using non-intrusive active testing. As good as packet-loss measurements are at identifying network bottlenecks, they cannot be used to describe the impact congestion is having on customer throughput, or on the latency of a particular application or flow.

In-Service throughput testing[3] is a unique technique that allows providers to determine the capacity available over a path, which can then be compared to a customer's CIR to determine the level of oversubscription present in the network. As in-service throughput testing has no impact on customer traffic, it can be conducted on-demand when issues are identified, or during peak traffic hours to accurately assess a service under worst-case conditions.

Based on the RFC-2544 throughput test, in-service throughput testing outputs latency, jitter, packet loss and throughput as measured one-way over the precise path a customer's traffic traverses. When used with long-term latency monitoring, in-service throughput testing can be a valuable tool in maintaining QoS and optimizing network design.

Optimize at the Edge

We've explored how latency issues can be displaced from one section of a network to another as link capacity and parameters are changed. From a network efficiency perspective, it is better to optimize latency, bandwidth usage and traffic priority as close to the service demarcation point as possible. This ensures the best-possible QoS end-to-end, as intermediate NEs have less traffic to manage. Core and aggregation NEs will have reduced forwarding delay as processing requirements are shifted to the network edge, as well as lower congestion-related delays as less traffic will be queued at all nodes across the network.

3 This technique is exclusive to JDSU Networks' MetroNID™ packet assurance demarcation units, with patents pending on the technology.

Figure 4: In-service testing example; customer traffic is considered part of the test traffic

Effective traffic management is best accomplished using hierarchical QoS (H-QoS), a combination of per-flow rate limiting (bandwidth policy enforcement), filtering and traffic shaping implemented against a service priority hierarchy. To ensure that H-QoS does not impact the latency it aims to improve, it is critical that the NE offers ultra-low pass-through forwarding delay as well as hardware-based shaping performance. As classifying traffic into service flows is required to apply H-QoS, the unit should also have versatile traffic classification and priority marking capabilities. This may involve defining as many as 500 unique flows for Ethernet wholesale hand-offs.

When properly applied, H-QoS can greatly reduce latency by ensuring critical traffic is identified and handled with sufficient priority, reducing packet loss (and re-transmission) through effective shaping methods, and limiting ingress traffic to each service's particular bandwidth profile. A key benefit of H-QoS is that latency for critical applications can often be greatly improved without any increase in access or core bandwidth.

A key benefit of H-QoS is that latency for critical applications can often be greatly improved without any increase in access or core bandwidth.

To summarize the key elements of effective latency and QoS optimization:

- Establish per-service priority and bandwidth as close to the edge as possible;
- Apply H-QoS using a hardware-based device with ultra-low forwarding latency;
- Classify traffic as granular as possible to ensure each application is given its correct position in the service hierarchy.

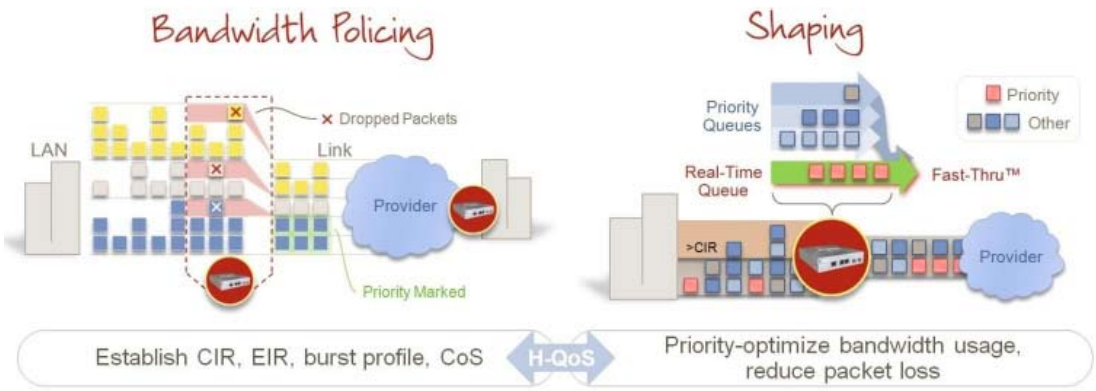

Figure 5: Components of H-QoS to summarize the key elements of effective latency and QoS optimization:

Using NIDs to Improve Latency & QoS

Network Interface Devices (NIDs), provide a cost-effective monitoring and traffic optimization solution. Typically installed at customer premises, cell sites and inter-carrier handoff points, NIDs are uniquely positioned to provide an end-to-end view of the network. When installed at key aggregation and switching locations, integrated test capabilities map out the health of the entire network to deliver per-flow visibility for all key QoS and SLA parameters. If the NIDs also feature embedded H-QoS functionality, a provider not only sees the performance of each service, but can precisely tailor it to meet individual customer and service requirements.

Networks' line of EtherNID™ & MetroNID™ packet assurance demarcation units were designed to address all of these requirements, while also integrating Ethernet OAM and MEF-certified Carrier Ethernet service mapping functionality. Based exclusively on a dedicated silicon, all-hardware design, these units' Fast-Thru™ packet processing engine offers forwarding latency and jitter of 3.3 and 0.1 microseconds, respectively, eliminating the delays associated with most NEs' network processor-based, store-and-forward architectures.

These devices are capable of high-density H-QoS on full line-rate traffic, can monitor up to 100 flows at layer 2 or 3, perform in-service throughput testing, and establish or terminate MEF-certified E-Line and E-LAN services. A unique remote clock synchronization technique provides one-way delay and jitter measurements to 1 microsecond resolution, even over geographically diverse, multi-technology, multi-vendor and multi-operator networks. Tests can be conducted to up to 100 remote sites from each NID, over point-to-point, hub-and-spoke, and full-mesh unicast or multicast services.

A zero-latency traffic shaping feature ensures that no jitter or latency is added to critical traffic, while minimizing packet loss and queuing delays for lower-priority flows. With real-time traffic classification, the units combine rate limiting and shaping into a sophisticated H-QoS function that is currently used in critical, delay-sensitive applications such as financial trading, cellular backhaul, media delivery and large enterprise business services.

Figure 6: JDSU Networks MetroNID™ Feature Set

Service Performance Monitoring

The NetComplete® Performance Management application allows service providers to cost-effectively scale their network and operational resources as service penetration increases by centralizing monitoring and providing automatic analysis of raw performance data. The NetComplete Performance Monitoring OSS efficiently collects both Y.1731 standard statistics and/or proprietary statistics from NID demarcation devices and/or NEs and generates both real-time views and performance/SLA reports with drill down capabilities. Ethernet service providers use these graphical views and reports generated to pro actively monitor network and customer Ethernet quality of service (QoS) to ensure service level agreements (SLAs) are consistently being met. And they use them reactively to help network operations

technicians troubleshoot problems quickly and effectively. NetComplete offers a high degree of automated results interpretation and customization features that allow technicians to rapidly analyze and interpret statistical data, and quickly draw conclusions to support overall service performance and customer satisfaction. These features include automatic association of statistic and SLA profiles to Ethernet circuits, at a glance measurement color coding according to associated SLA profile, 1-way and 2-way result indication and advanced configuration for filtering/grouping of measurements for reporting. SLA reporting sites for each customer, showing only the metrics they need to confirm their services are functioning within spec.

Figure 7: NetComplete SLA with real-time reporting

Carrier-Grade to the Core

Adding an EtherNID™ or MetroNID™ unit to your network is a fail-safe choice. Designed and built from the ground-up as a premium, in-line network element ensures reliability is never impacted, only assured.

Unconditionally certified to MEF 9 & 14 standards these units offer the full functionality required to provision, police & deliver Carrier Ethernet Virtual Circuits (EVCs). The units are also certified to NEBS Level 3, delivering performance & non-interference with other carrier-grade elements. Temperature hardening allows installation in outdoor cabinets and harsh environments.

With no moving parts or fans to fail and consuming only several watts, EtherNID & MetroNID units feature a 52 year Mean Time Between Failures (MTBF) when evaluated by the Telcordia Reliability Prediction Procedure.

A fail over bypass circuit ensures links never go down if a power outage occurs. Ensuring that's a remote possibility, the units feature 3 way redundant power with instantaneous switching between twin 48 VDC and AC-powered feeds.

Link redundancy is provided by 1+1 protected uplinks (optical or copper pairs) with ultra fast-fail over, LAG & LACP support.

Networks' Packet Performance Assurance solutions enable carrier-grade, packet-based wireless backhaul, business services & multi-carrier applications over wireless & wireline networks.

Test & Measurement Regional Sales

NORTH AMERICA	LATIN AMERICA	ASIA PACIFIC	EMEA	WEBSITE: www.jdsu.com
TEL: 1 866 228 3762 FAX: +1 301 353 9216	TEL:+55 11 5503 3800 FAX:+55 11 5505 1598	TEL:+852 2892 0990 FAX:+852 2892 0770	TEL:+49 7121 86 2222 FAX:+49 7121 86 1222	