

10 GIGABIT ETHERNET WITH POLICY-BASED QUALITY OF SERVICE

Bandwidth is about to get way out of control, as many IT managers and metro providers know firsthand. Merely pumping up performance won't magically make their problems go away. In fact, bigger pipes can mean bigger headaches—if comprehensive quality of service (QoS) is missing in action. That's why the true test of a switch isn't just how fast it tears through traffic but how smart it is at knowing what to do with the packets it's processing.

Leave it to Extreme Networks® to meld mind and muscle—at an unprecedented 10 Gbps.

Introducing 10 Gigabit Ethernet for Extreme's workhorse BlackDiamond® chassis switch. This latest addition to Extreme's Ethernet arsenal obviously has the bandwidth to blow through any bottleneck. Better yet, it boasts the industry's most advanced bandwidth management, rate shaping and Policy-Based QoS capabilities. The result: Every type of traffic gets exactly the right treatment—without slowing switch performance.

That's not just the best of both worlds, it's an entirely new world for enterprise networkers and metro providers alike. Extreme's intelligent approach to 10 Gigabit Ethernet lets them bump up their bandwidth by an order of magnitude or more, simply by slotting a 10 gigabit module into any BlackDiamond chassis. Actually, make that modules: A fully loaded BlackDiamond 6816™ accepts 16 10 gig blades, each of which has the QoS qualifications to make the most out of every available bit.

AGGREGATION WITHOUT AGGRAVATION

Before digging into the details, let's take on the skeptics in the audience whose response to 10 Gigabit Ethernet may well be: Who needs it?

The Enterprise. Most IT managers who've deployed 1 Gbps Ethernet already know the answer to that one: They do. Start adding up all those 1 gig links and they're probably facing some serious oversubscription problems, if not a full-blown oversubscription crisis. Depending on the size of the enterprise and the complexity of the data center, they may already be struggling with oversubscription ratios of 10:1, 20:1, 30:1, or even 100:1. And given the way the corporate brass has a habit of falling in love with three-letter acronyms that translate into bandwidth-intensive applications like customer relationship management (CRM), network traffic isn't going anywhere but up. Extreme's 10 Gigabit Ethernet solution lets IT pros aggregate links as needed, without breaking a sweat.

Metro Providers. Providers, of course, are expected to haul all this traffic across the metro or wide area—and their customers don't want to hear about fiber exhaustion or timing problems that can be traced back to WAN link aggregation. With Extreme's 10 Gigabit Ethernet solution over dark fiber, they won't have to: Every point of presence (POP) can be a top performer. And incrementally adding throughput as required takes a lot of the agony out of traffic engineering. What's more, Extreme's 10 gigabit modules work with dense wave-division multiplexers (DWDM) and SONET/SDH gear, making end-to-end Ethernet a reality, which is exactly what more and more of your enterprise customers want.

WAKE UP AND SMELL THE QoS

There's another concern that needs to be faced head on: With 10 gigabit of bandwidth to burn, why worry about QoS at all? Actually, the opposite is true. Most high-speed connections are already loaded with a mind-numbing mix of traffic, including delay-sensitive voice-over-IP (VoIP), services, bursty Web data, and legacy applications that were never designed with the enterprise or IP in mind. Extreme's Policy-Based QoS is the only way to guarantee that every packet is processed properly, whether it needs low latency or can't be allowed to exceed a predetermined threshold.

Metro providers can supply the rest of the answer. QoS is the key to putting a premium price on premium services—and making sure that customers get what they pay for. In addition, a growing percentage of businesses are running specialized applications, like voice or video, which are going to wind up dead on arrival without plenty of protection from Extreme's Policy-Based QoS.

The end result is that higher bandwidth links like 10 gig demand sophisticated QoS capabilities because:

- 1) Aggregated traffic is more heterogeneous and when you combine different traffic types without QoS, only the brutes win while the real gems are forced out. UDP-based services and applications force out TCP-based services and applications. TCP-based services and applications that are delay-sensitive lose the most because services and applications that don't work well don't sell and get used.
- 2) The more services and applications that work on the network, the more valuable the network is.

GETTING TO KNOW 10 GIG

So what's the skinny on 10 Gigabit Ethernet? Officially, the spec is the handiwork of the IEEE 802.3ae 10 Gigabit Ethernet Task Force, whose job is to write and ratify the final standard (which is slated to come up for a vote soon). The draft standard defines a full-duplex, fiber-only technology that can be extended to 40 kilometers. It uses the same media access control (MAC) protocol, frame size, and framing format as tried-and-true Ethernet, which means it enables a seamless, tenfold speed increase over its 1 gig predecessor. And thanks to a very smart piece of engineering called the WAN PHY, 10 gig data streams can use legacy SONET links as a physical layer transport (at 9.29 Gbps).

There's something else to know if you are considering this supersonic Ethernet scheme. No vendor's switch is currently capable of delivering a full 10 Gbps—including Extreme's. It's an architectural issue that's being addressed in next-generation hardware. How does Extreme stack up? Its current BlackDiamond tops out at 8 Gbps bidirectional per slot, which is the best anyone in the industry can deliver without resorting to expensive, proprietary kludges that are decidedly nonstandard. But rather than engaging in specsmanship, Extreme has come through with something far more important: the bandwidth management and buffers to smooth potential traffic spikes and keep packet loss at a minimum.

DISSECTING POLICY-BASED QoS

Let's cut to the chase: Policy-Based QoS is how Extreme makes sure that IT managers and metro providers always make the most of their 10 gig pipes.

The place to begin is with the intelligent bandwidth management built into every Extreme switch. Basically, this Policy-Based QoS capability makes it possible to control the interpacket latency of data stream passing through your pipes – whether it's 10/100 Mbps, 1 gigabit or 10 gigabits. That's very good news for delay-sensitive apps and services like VoIP and video, which are notoriously susceptible to these sorts of slowdowns. Voice is actually the more sensitive of the two, since high latency ends up clipping voice packets, rendering conversations unintelligible in the process. That's not to suggest that video gets off much better: Latency and jitter, which is the variation in latency, typically leads to the sort of herky-jerky motions that make streaming video a contradiction in terms.

But latency also can make losers out of other services and applications, like enterprise resource planning (ERP). And it can be a killer for legacy applications like IPX databases, which were designed for LANs and typically time out when they run into enterprise-level latency (when they run at all). Similarly, extending broadcast domains and virtual LANs (VLANs) over long distances can create delay-related difficulties, since some services and applications, like storage, assume they're on a single subnet and expect fairly quick turnaround times. In these instances, latency can lead to false errors and session timeouts or even kill the service or application entirely.

How does Extreme prevent delay from becoming deadly? The answer lies in precisely controlling the bandwidth consumption of data streams, which in turn controls interpacket latency and jitter. How? By providing a way to set and enforce the minimum bandwidth for any packet stream. This ensures that a specific stream will always have access to the wire at a specific rate. And that, in turn, locks down and stabilizes the arrival time between packets. No more garbled conversations, stop-motion video, or session timeouts.

SHAPE UP OR SHIP OUT

But bandwidth management doesn't just control latency. It is, in fact, just the beginning. Extreme also offers extensive traffic-shaping tools that make it possible to set maximum bandwidth consumption for any traffic type, thus preventing bursty data and web sessions from hogging the wire and starving other services or applications. Better yet, packets that exceed the upper limit of a particular stream are buffered and then sent as incoming traffic falls off. This is known as traffic shaping and is far more efficient than simply dumping packets. It also gives much better overall utilization of the link.

Extreme does something else very smart here. It allows bandwidth not being used by one traffic type to be exploited by another, rather than sitting idle because a specific stream isn't gobbling up every last bit allocated to it. Time-division multiplexing (TDM), in contrast, slices up the available bandwidth and reserves it for a particular flow, which is pretty inefficient, even on a 10 Gbps connection.

THE WEIGHTING ROOM

It's important to understand that Extreme's bandwidth management and traffic shaping aren't just another way of implementing weighted-priority schemes (like weighted fair queuing). Simply said, weighted priorities won't cut it on high-speed links.

Basically, weighted priorities are all about ratios; for example, X number of voice packets transmitted for every Y number of data packets. The problem with ratios is that they create a false interdependence between what are essentially unrelated traffic flows. They also raise some troubling questions if traffic flows vary from expected ratios. What happens, for example, if there's a 10:1 ratio between data and voice and fewer than 10 data packets are sent? Does voice get transmitted? How are latency and jitter impacted for voice if the switch maintains ratios between data and voice?

Extreme, in contrast, lets IT professionals decide exactly the treatment every independent data stream needs. It allows unused bandwidth to be shifted where it's required at wire speed. And it makes it possible to manage every queue, configuring minimum bandwidth (for latency control) and maximum bandwidth (for rate shaping).

CLASSIFICATION STATION

The BlackDiamond also offers extraordinary classification for each of the eight queues on every port, so giving the right treatment to the right data stream is never a problem. Classification can be explicit, relying on IP DiffServ code points or IEEE 802.1p tags. Or it can be implicit, triggered by ranges of Layer 4 IP information, IP source or destination address, MAC address or incoming physical port.

Better yet, the BlackDiamond can always apply Layer 3-4 classification, even when it's working as a straight Layer 2 switch. What's more, every queue comes complete with random early detection (RED), so if traffic spikes turn out to be long plateaus, the BlackDiamond will drop packets randomly instead of deleting entire blocks of a specific stream. True, the downstream client will lose a few packets, but all that means is that it will negotiate a smaller TCP window size and slow the session down.

That may not sound like a big deal, but RED is one of the ways that the BlackDiamond overcomes its 8 Gbps per slot architecture in the brave new world of 10 Gbps connectivity.

Let's take it from the top. The BlackDiamond's extensive bandwidth management, coupled with the deep buffers on every queue, let it smooth out transient traffic spikes to minimize packet loss and deliver better overall utilization, while controlling latency and jitter for delay-sensitive services and applications. In the event of sustained loads that will ultimately overflow the buffers, RED kicks in on the affected queues to preserve TCP session integrity while windows are renegotiated. The lesson: Brains will always outperform brute force.

Finally, Extreme offers incredible Policy-Based QoS capabilities that are sure to interest metro providers: A way to remap packets on the fly according to predefined DiffServ code points. When invoked on a BlackDiamond at the edge of a network, this QoS capability (known as "coloring") gives providers a way to automatically enforce premium service-level agreements (SLAs). Basically, coloring can be used to dynamically correct explicit QoS signaling like IP DiffServ and IEEE 802.1p, thereby preventing customers from pumping up priorities and abusing SLAs. That's not just good for business, it makes enforcing end-to-end QoS a no-brainer.

10 GIG FOR ENTERPRISES

Extreme understands that the key to successfully deploying 10 Gigabit Ethernet is using it where and when it's needed, without being forced to buy more capacity or more chassis than is truly needed. That's why Extreme's 10 gig solution is available in a choice of three BlackDiamond chassis, each designed to give IT managers and metro providers an unparalleled degree of redundancy and resiliency. The BlackDiamond 6816, which accommodates up to 16 I/O modules and four management switch-fabric modules (MSMs), is a powerhouse in the network core or heart of the data center. The BlackDiamond 6808 supports up to eight I/O modules and dual MSMs to provide an incredibly intelligent aggregation solution. And the new BlackDiamond 6804, which features up to four I/O modules and two MSMs, delivers significant advantages to the network edge.

Enterprise networkers, for example, can rely on BlackDiamond to extend 10 Gigabit Ethernet across a campus, giving outlying buildings a redundant, high-speed path to required resources. In fact, given Extreme's bandwidth management capabilities, it's possible to establish a so-called serverless buildings that exploit centralized services and applications (and reduce server deployments) while guaranteeing latency.

What's more, those outlying buildings can be a fair distance away, since Extreme offers an optical interface that extends 10 Gigabit Ethernet distances to 10 kilometers. Also, Extreme will soon offer hot-pluggable XENPAK optical inserts, that extend 10 gig distances to 40 kilometers. At that range, IT managers can start reconsidering their real estate options, with an eye to moving data centers or remote sites to less expensive locations. They also can begin giving some serious thought to linking multiple data centers—again, without worrying if latency-sensitive or bandwidth-intensive applications are going to cause problems on their pipes.

Also consider this: Extreme's 10 Gigabit Ethernet capabilities open the way to Ethernet storage area networks (SANs) with native iSCSI interfaces on servers and storage devices. This promises to redefine the conventional wisdom for disaster recovery, remote backups and business continuance. With iSCSI, it's possible to port native SCSI protocols and place them on an IP network. This lets servers with host-based adapters (HBAs) have fast hardware-based iSCSI translators that can connect to any 1 or 10 gig Ethernet/IP network and talk natively to an iSCSI storage device.

In addition, 10 Gigabit Ethernet is well-suited for server farm and supercomputing clusters. Server farms and supercomputing clusters have become immensely popular for enterprises, universities and government research facilities that want to harness the power of a group of computers working as a single, continuously available system. Clustering provides the power needed for the development of solutions that can process complex, data-intensive scientific problems, such as the modeling of automobile crash test simulations.

Of course, it's not all about going the distance. Extreme's take on 10 Gigabit Ethernet is equally at home within the confines of the data center, where it's a natural for switch-to-switch and switch-to-server apps. With a BlackDiamond pumping packets and paying attention to QoS, IT managers can begin building out their 10 Gbps backbones, future-proofing the enterprise for streaming video, CAD/CAM, modeling, imaging, or whatever bandwidth-hungry applications are headed their way.

10 GIG FOR METRO PROVIDERS

What about metro providers who want to put 10 Gigabit Ethernet to work? As noted, the place to start is by deploying BlackDiamonds to any existing POP on the metro ring that might be taking a pounding. This puts an immediate end to the capacity crunch, while giving providers a way to offer premium service-level guarantees and extend the reach of VoIP, real-time video, streaming audio and other revenue-generating services by lighting up dark fiber.

Providers with metro offerings in multiple cities also can count on Extreme for connectivity. In this instance, the BlackDiamond's aforementioned WAN PHY interface will be their ticket to travel. It will work with a long-haul carrier's DWDM equipment or optical regenerators, and enable 10 Gigabit Ethernet to run over SONET/SDH infrastructures at 9.29 Gbps.

THE EXTREME ADVANTAGE

At Extreme Networks, Ethernet Everywhere® isn't merely a motto, it's a blueprint for the future. Extreme's solution for the BlackDiamond scales from 10 Mbps all the way to 10 Gbps, while providing a consistent hardware, software and management architecture across all product platforms. That gives IT managers and metro providers the opportunity to deploy exactly the bandwidth they need, where and when they need it, along with a simple, cost-effective migration strategy that lets them take the next step into tomorrow at their own speed. What's more, thanks to Extreme's industry-leading Policy-Based QoS capabilities, every packet traveling over their networks will always get the precise treatment it requires.

There's something else to consider: Extreme's commitment to Ethernet doesn't stop with its own solutions. Two highly placed members of Extreme's professional staff serve on the IEEE 802.3ae 10 Gigabit Ethernet Task Force, in the roles of vice chairman and secretary. Similarly, another member of Extreme's professional staff is on the board of the 10 Gigabit Ethernet Alliance (10GEA), a nonprofit industry consortium. Extreme is also a member of the XENPAK multisource agreement, which defines fiber optic transceiver modules that conform to the 10 Gigabit Ethernet standard as laid down by the IEEE 802.3ae.

Ethernet Everywhere: It's a way of life at Extreme.



3585 Monroe Street Santa Clara, CA 95051-1450 Phone 408.579.2800 Fax 408.579.3000
Email info@extremenetworks.com Web www.extremenetworks.com

© 2001 Extreme Networks, Inc. All rights reserved. Extreme Networks, BlackDiamond, Summit, Summit7i, ExtremeWare, ServiceWatch, Extreme Ethernet Everywhere, Ethernet Everywhere, Extreme Velocity, Extreme Turbodrives and the color purple are registered trademarks of Extreme Networks, Inc. in certain jurisdictions. Alpine, ExtremeWare Vista, Extreme Standby Router Protocol, ESRP, Summit1i, Summit4, Summit4/FX, Summit5i, Summit24, Summit24e2, Summit24e3, Summit48, Summit48i, SummitLink, SummitGbX, SummitRPS, SummitPx1, PxSilicon, EPICenter, vMAN, the BlackDiamond logo, the Alpine logo and the Extreme Networks logo are trademarks of Extreme Networks, Inc., which may be registered or pending registration in certain jurisdictions. ExtremeWorks, the Extreme Turbodrives logo and the Go Purple-Extreme Solution Partner logo are service marks of Extreme Networks, Inc., which may be registered or pending registration in certain jurisdictions. All other registered trademarks, trademarks and service marks are property of their respective owners. Specifications are subject to change without notice.

L-MO-10GIG-201