



W H I T E P A P E R

BANDWIDTH MANAGEMENT — Scalable, Granular Resource Control for Web Data Centers

■ INTRODUCTION

■ WEB OS BANDWIDTH MANAGEMENT IN DETAIL

■ WHY BANDWIDTH MANAGEMENT IS BETTER

■ BANDWIDTH MANAGEMENT APPLICATIONS IN E-BUSINESS

■ CONCLUSION

Alteon WebSystems, Inc.

50 Great Oaks Boulevard
San Jose, California 95119
408-360-5500
408-360-5501 fax

<http://www.alteon.com>

94004.25/02-01

INTRODUCTION

With more and more users, applications and organizations sharing resources in a Web hosting center, the ability to monitor and control bandwidth utilization by different traffic categories has become increasingly important. Web hosting operators need it to offer their customers Service Level Agreements (SLAs). Web site operators need it to ensure the best performance for their money or to avoid paying over-subscription penalties.

While this requirement is not new, many Web data centers are still operating under a “best effort” model, because the tools to monitor and control traffic at the volume and granularity levels required were simply not available.

Bandwidth management or Quality of Service (QoS) solutions on routers and LAN switches do not offer the flexibility and precision necessary to manage the variety and combinations of virtual services housed in today’s Web data centers. Dedicated bandwidth management appliances, on the other hand, lack performance and scalability to keep pace with growing traffic in today’s data centers.

Hence, Alteon has introduced high-performance bandwidth management service to give Web data center operators fine-grained control over bandwidth utilization. Integrated with server load balancing, traffic redirection, access control and content processing services on a scalable Web switching platform, Alteon’s Web OS bandwidth management is completely compatible with existing QoS mechanisms on routers and LAN switches. For service providers and enterprises with installed QoS networks, Alteon’s switch-based bandwidth management extends QoS to the “last mile” – to servers, applications and content – for end-to-end control without application changes.


Even for organizations that can afford to “throw bandwidth at the problem,” Web OS bandwidth management offers benefits in security monitoring and control.

This paper describes how Web OS bandwidth management works and the applications it enables.

WEB OS BANDWIDTH MANAGEMENT IN DETAIL

What it is

Bandwidth management simulates multiple “virtual pipes” within one or more physical interfaces (ports). Depending on the model, Alteon Web switches can support either 256 or 1,024 virtual pipes, or traffic classes. Each traffic class defined by a broad range of policies including physical port, VLAN, source or destination IP address, TCP or UDP port number, URL, HTTP cookie and other access control filters. Complex classifications such



as a source-destination flow, a QoS class, an HTTP cookie within a virtual service, etc., can be defined via the rich set of Web OS traffic filters.

Administrators can also define the order of precedence when handling packets that fit the criteria for multiple virtual pipes; by default, the switches select according to virtual service (or IP), then filter, then VLAN, and finally, physical port.

For each virtual pipe, the administrator applies a bandwidth policy which must include three data rates: a Committed Information Rate (CIR), a Soft Limit, and a Hard Limit.

- The CIR is the data rate that the switch will guarantee the virtual pipe; for a given sample period, the data rate will never be lower than this amount.
- The Web switch constantly works to move traffic flow rates to the soft limit. It dynamically adjusts the frequency with which it forwards traffic from the virtual pipe according to the volume of traffic and the rate it is trying to simulate.
- Hard limit defines the “cut off” rate over which traffic will be discarded. The hard limit can always be set to the maximum rate of the physical port(s) if traffic discard is to be avoided.

Availability of the soft limit, in addition to the minimum and maximum rates, allows hosters, service providers, and IT administrators to regulate each traffic flow to a desirable target rate, while leaving the systems a degree of flexibility necessary to respond to temporary bursts and congestions.

How it works

For each virtual pipe, the switch simulates a physical link by pacing packet transmission at a simulated rate. Initially, the simulated rate is set to the soft limit. When a packet is transmitted, the switch determines the difference between the time that it would have taken to transmit the packet at the simulated rate and the actual time it took to transmit the packet at the physical port speed. It then waits for that length of time before sending another packet from the same virtual pipe. When a virtual pipe spans multiple physical ports, the switch automatically calculates the correct governing rates at each port.

Packet buffers at each egress port are divided into a queue per virtual pipe. Different queues can have different amounts of buffering (queue depth) allocated, defined by the administrator as part of the bandwidth policies. For example, applications that are not loss-sensitive such as voice should be in shallow queues, while loss-sensitive applications such as FTP should go in deep queues.

If a queue is approaching its queue depth over time, the switch will increase its simulated rate, up to the hard limit, to allow for temporary bursting. When the hard limit is reached, new packets will not be queued until the simulated rate drops below. In other words, if an application is sending data more quickly, the data rate will increase up to the maximum rate. As the queue depth decreases, the simulated rate is ratcheted down accordingly, although it will never drop below the CIR.

Transmission for each queue is regulated individually, until global congestion occurs. When the global queue depth reaches a congestion threshold, the simulated rates are reduced on each queue according to its bandwidth policy. This scheduling scheme effectively allows each virtual pipe to borrow bandwidth up to its hard limit, if bandwidth is available.

Monitoring and managing traffic

Statistics tracked by bandwidth management allows administrators to identify when and how much bandwidth each virtual pipe has borrowed for billing and reporting use. Critical information such as queue depth, bytes sent and bytes dropped per virtual pipe is sampled every second. When a virtual pipe spans multiple physical ports, the switch automatically surmises the information from each port. Statistics can be polled via standard SNMP interfaces, or they can be automatically download to a pre-configured email address for history and accounting records.

Enabling end-to-end QoS

While the load balancing client requests to server farms, the same Web switch is ideally positioned to meter outbound server traffic and control its utilization of upstream resources such as the WAN router links. An optional feature in bandwidth management, IP Type Of Service (ToS) tagging, further allows administrators to mark the packets transmitted below and above the soft limit with different QoS settings. This enables upstream, QoS-capable, switches and routers to prioritize packet delivery according to the target bandwidth rate set for each traffic class, without incurring the cost of classifying and tagging traffic themselves.

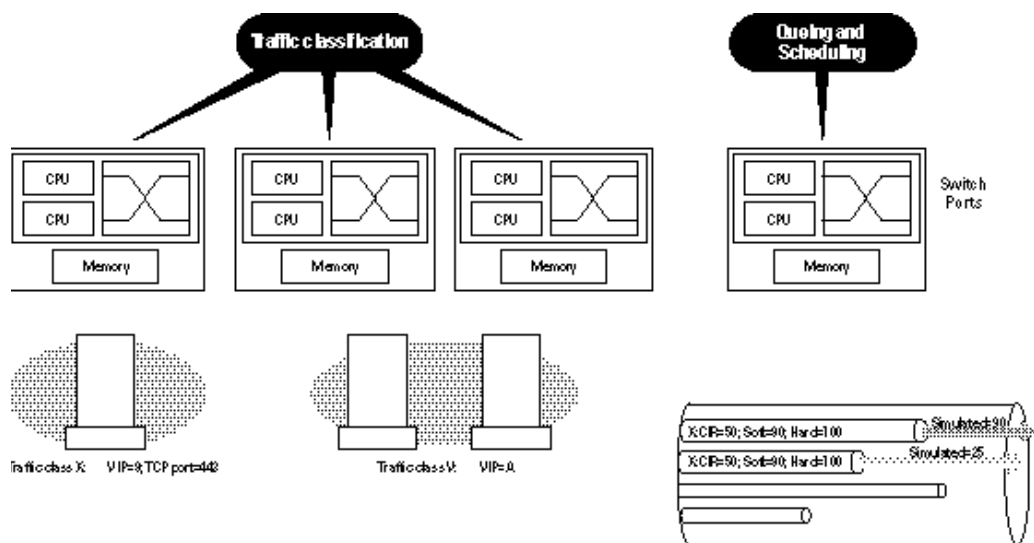
WHY BANDWIDTH MANAGEMENT IS BETTER

Traditional QoS models in packet networks provide only coarse bandwidth control. For example, the Committed Access Rate model that is available on some router platforms divides traffic into “in-spec” and “out-of-spec” packets per interface. Traffic that exceeds the committed data rate is outside the specified commitment and forwarded at a lower priority or discarded. The routers do nothing to regulate traffic flow or correlate utilization for a single traffic class across multiple physical interfaces.

QoS is also complex. With bandwidth management, service providers and application hosters are speaking a simple, common, well-understood language-link rates, in simple 64Kbps increments. Rather than applying traffic shaping algorithms like Random Early Discard or Weighted Fair Queueing, which are all imprecise, there is no ambiguity or confusion with Alteon’s class-based bandwidth management solution.

Optimized by Alteon’s distributed processing architecture

The unprecedented flexibility, granularity and performance with which bandwidth management performs traffic classification and bandwidth enforcement is attributable to Alteon’s distributed processing architecture. With bandwidth management, traffic classification and queuing are managed by the WebIC-embedded network-processors at each ingress port, while transmission scheduling is performed by the network-processors at the egress port(s). The parallelism reduces the processing time for complex classification and scheduling operations, minimizing delays.



Traditional packet switches with centralized processor architectures can offer gigabit-speed QoS by integrating queuing mechanisms in hardware. But unlike ATM switches, the complexity of managing different scheduling rates for variable size packets on queues that may span multiple ports simply makes implementing a large number of queues cost prohibitive. Implementing this in software in a centralized processor, on the other hand, results in poor performance.

The same distributed processor model that enables Alteon's industry-leading Virtual Matrix Architecture (VMA,) is required to deliver the flexibility, scalability and performance demanded bandwidth management in Web data centers.

PRACTICAL APPLICATIONS OF BANDWIDTH MANAGEMENT IN E-BUSINESS

QoS technologies hold real promise for managing traffic growth, but until now complexity and compatibility issues have slowed their broad adoption. With the simplicity and power of bandwidth management, many critical e-business problems can now be solved.

Protecting sites from burst penalties. Traditionally, a company bought capacity from a service provider and had the right to saturate the link for which they paid. In today's densely populated Web, smaller sites don't need – and cannot afford – an entire dedicated link. It is also more difficult to provision dedicated physical circuits for a large number of collocated customers. Consequently, service providers offer a connection whose physical capacity exceeds the contracted average link rate for the connection. If a subscriber exceeds the average link rate, it causes over-subscription on an upstream router and potentially impacts other users. Service providers mitigate this by charging hefty premiums for traffic above the agreed-upon rate in order to encourage subscribers to buy the maximum bandwidth they need. But in today's information economy, traffic fluctuations can be huge. bandwidth management lets site managers control their utilization and scale capacity on their terms, rather than paying unacceptable premiums for heavy load.

Enabling service providers to aggregate profitably. Today's collocated environments rely on virtualization – the ability to simulate multiple logical systems with a single real system. In a hosting environment, many sites share a single physical WAN link; in an access ISP's point-of-presence, different types of remote access servers share a single connection to the Internet. Providers depend on "law of averages" in provisioning shared resources to maximize profits while maintaining customer satisfaction. But without careful regulation of individual traffic streams, bandwidth is not shared equitably. A Web site experiencing flash-crowd can overwhelm others Web sites, and fast DSL modem traffic can impact slower dial-up connections. Bandwidth management not only allows providers to enforce fairness, they can offer flexible, usage-based services to increase revenue.

Increasing Security. Denial of Service attacks can come from a variety of sources, under a variety of guises. Sophisticated attacks are difficult to prevent as attack traffic appears as legitimate traffic. A way to thwart denial of service attacks is to set rate limits at traffic ingress points for "suspect" traffic, such as TCP SYN packets, PING packets and gateway broadcasts packets. Universities can limit outgoing attack traffic without stripping users' freedom by imposing a rate limit on all packets destined for unknown application ports. Bandwidth management provides a monitor mode whereby statistics on each traffic class are collected without actual bandwidth enforcement. This allows administrators to observe the typical amount of bandwidth consumed by each type of suspect traffic under normal conditions, then set bandwidth limits that can prevent and alert them of potential attacks.

CONCLUSIONS

Alteon's bandwidth management solution solves real problems for e-business innovators in a simple, powerful manner. It leverages Alteon's distributed processing model to offer fine-grained classification and dynamic link simulation at unprecedented data rates. And it underscores the future-proofed nature of Alteon's products, needing no hardware changes.