WHITE PAPER

Multi-Level Traffic Management

Abstract

Allot Communications' NetEnforcer offers an intelligent, comprehensive, policy-based approach to traffic prioritizing, load-balancing, cache redirection and accounting. The NetEnforcer controls the most critical areas of the network using an architecture that allows users to optimize their traffic flows. This approach provides granular per-flow bandwidth management and can shape both incoming and outgoing traffic flows. Through its advanced traffic control algorithms, the NetEnforcer assures guaranteed traffic flows, minimal packet loss or retransmissions, optimizes WAN throughput and assures that the network is always used to its maximum potential.

Class-based Queuing Approaches

Many router products today use various queuing approaches such as weighted-fair queuing or class-based queuing. These queuing algorithms provide fairness between different classes or priorities of traffic.

However, connection flows that are of the same priority class have no consistent fairness policies. If a connection comes in with a given priority or guaranteed bandwidth, it will be put on a certain queue. As traffic on the router begins to queue up and more connections arrive with that priority class, the new connection will always go to the back of the queue and must wait until all previously queued packets are sent out over the line before it can begin sending its packets. The end result is inconsistent and unpredictable delivery of traffic.

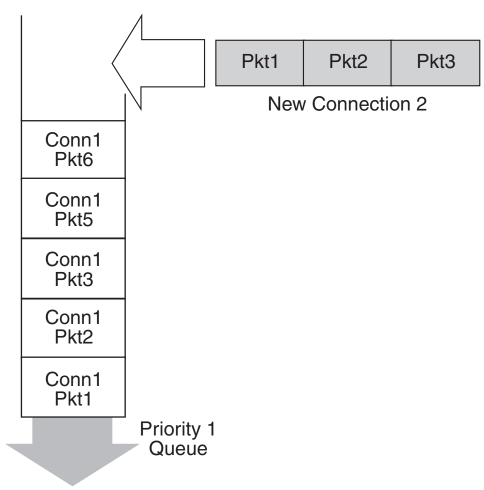

Figure 1. With the traditional queuing approaches, there is no "fairness" between connections of the same class or priority. A new connection will simply go to the back of the current queue, resulting in inconsistent or unpredictable results.

The Allot PFQ Approach

The NetEnforcer uses a unique approach to queuing called **Per-Flow Queuing (PFQ)**. With PFQ, each new connection flow gets its own queue. This new queue can then be treated equally with other flows having its same priority policy class. New connections will wait a fair, proportionate amount of time relative to all other connections. The end result is that connections in a given class policy will get predictable results on the network.

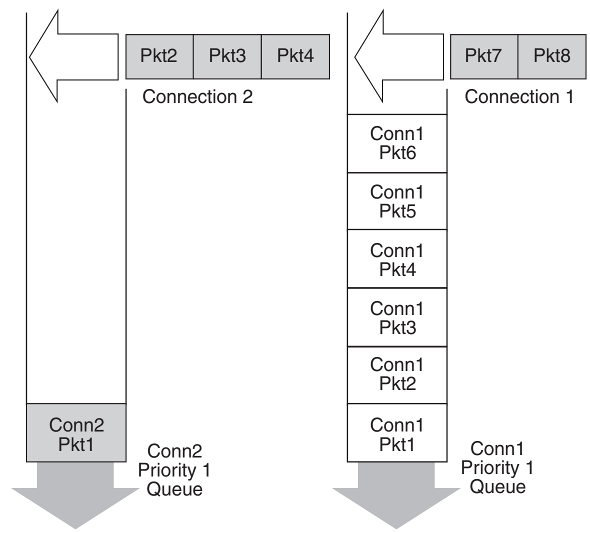

Figure 2. With the Allot NetEnforcer and PFQ, each connection flow receives it's own per connection flow queue. This means that the QoS algorithm can be tailored for each specific flow of traffic.

Advanced Queuing Approach

The NetEnforcer will perform a superset of traditional weighted-fair and class-based queuing algorithms to control the flow of connections between different priority classes. Based on user-defined policies, traffic will be put in various queues and sent to the network based on bandwidth requirements. Using advanced control algorithms, both prioritized and guaranteed traffic can be fairly allotted without starving lower priority connections. The Allot queuing mechanism is built using the hierarchical approach, thus subclasses of traffic could be independently prioritized within the class, while traffic classes maintain priority among them.

Within this class-based mechanism the per-flow queue is maintained where each connection or communication flow is allocated with it's own queue. These queues are combined with bandwidth guarantee and limitation mechanisms that provide preferred treatment to flows with traffic

Allot Communications

guarantee policies. These bandwidth guarantee policies are implemented using combination of leaky bucket and token bucket techniques. These techniques provide precise bandwidth allocation control as well smoothing and delay guarantees.

Allot products use these mechanisms to support DiffServ definition, not only managing the local bandwidth, but marking and remarking IP datagrams according to their conformance to the traffic policy profiles. This allows the maintaining of end-to-end Service Level Agreements in a DiffServ enabled environment.

Traditional Queuing Approaches

The advantage of traditional queuing approaches to quality-of-service is very precise control of traffic when traveling from a high speed LAN to a slower WAN. Traffic is classified according to policy and then assigned to various priority queues. Although this technique is very accurate for outgoing traffic, it is not effective for control of inbound WAN traffic. Queuing is a localized solution that only takes into account the state of the network at the point of the prioritizing system. If traffic suddenly becomes congested, packets will begin to accumulate in the various queues. Very quickly, packets will be dropped and end-to-end retransmissions will occur. This inevitably causes more traffic on the congested network and results in an uneven end-to-end flow of traffic on the network. In summary, bursty data traffic, caused by delays throughout the network, are not easily dealt with using local queues. In addition, this queuing system focuses primarily on outbound traffic. Since the queuing system only has a local view of traffic, traffic flowing from a slow-speed connection to a high-speed connection cannot be accurately controlled because, by the time the traffic reaches the bandwidth management device, there is nothing to queue.

TCP Rate Control Approaches

The traditional TCP rate control performs arbitrary window size tampering, which cause the underutilization and waste of the your most valuable resource - the WAN link. This is because:

- It doesn't take into account the congestion on the link (and this is exactly when you need bandwidth management) and even worsen the situation by multiplication of ACK packets
- It cause slow recovery for connections that was previously slowed down.

For the traffic that comes from the high-speed link to the slow-speed link - the window size tampering is not the right method to control the rate.

Unfortunately, this gives little real-time control of a connection. If there is too much traffic sent, congesting the network, it will cause excess retransmissions resulting in more traffic on the already overloaded network. If the system prevents too much traffic from getting through, the network will be underutilized. If there is a sudden burst of data on several connections, it will take several seconds for individual TCP connections to recognize this and for the throttling mechanism, inherent in the protocol, to slow things down. The end result is often unnecessary delays and oscillations in controlling traffic between clients and servers. In addition, protocols, such as UDP, which are not connection oriented and have no end-to-end flow control mechanism, cannot be controlled.

Controlling Traffic Flow – The Allot Multi-Level Approach

The NetEnforcer performs traffic shaping using an intelligent PFQ approach to controlling traffic flow. This flexible system allows both short-range and long-range control of both inbound and outbound traffic. Each connection throughput is metered and the exact amount of bandwidth is allocated to it. The end result is that the WAN connection is maximized to its full potential.

Both inbound and outbound sides of the flow are continuously estimated in real-time, and bandwidth is allocated according to actual flow throughput and policy applied to the flow. The mechanism can smooth particular flows, as well as allow some level of burstiness. Delay boundaries can be assigned to the flow allowing best performance for delay-sensitive applications such as audio and real-time.

Achieving Maximum Performance

The NetEnforcer's multilevel control is able to achieve wire-speed control through an advanced packet analysis architecture. The NetEnforcer is divided into two logical layers: the Data Knowledge Model (DKM) layer and the Fast Wired Switch (FWS) Layer.

Data Knowledge Module Layer

The Data Knowledge Module (DKM) Layer handles the advanced analysis required to determine the proper class of service to assign to the connection. It takes into consideration the state of the entire system and will interface to various higher-level analysis functions in order to determine the correct policy and class-of-service to assign to specific connections and traffic flows.

Fast Wired Switch Layer

The Fast Wired Switch Layer (FWS) presents wire speed IP routing and bridging software. The FWS will communicate with the application analysis layer to build a high-speed table that defines the physical connection initiated by the DKM Layer. The DKM Layer will instruct the FWS to pass up specific traffic that contains information necessary to make the application layer decisions. This includes connection setup information as well as specific data that require parsing as part of specific rules. Once a traffic flow is established, it will be handled exclusively by the FWS layer, which is optimized to transparently pass traffic, at wire speeds, between the NetEnforcer physical interfaces.

Summary

The key to the Allot architecture is maximizing the use of the WAN and delivering the most predictable service to users. Allot's PFQ technology uniquely assures fairness both between priority classes as well as for multiple connections within a single class. Because of its PFQ approach to bandwidth control, the highest granularity and precision of bandwidth control is achieved. Combining this with a two-tier internal architecture, the NetEnforcer can assure full wire-speed traffic flow. The end result is a predictable, policy-based user environment that offers the most optimal system for controlling WAN traffic.

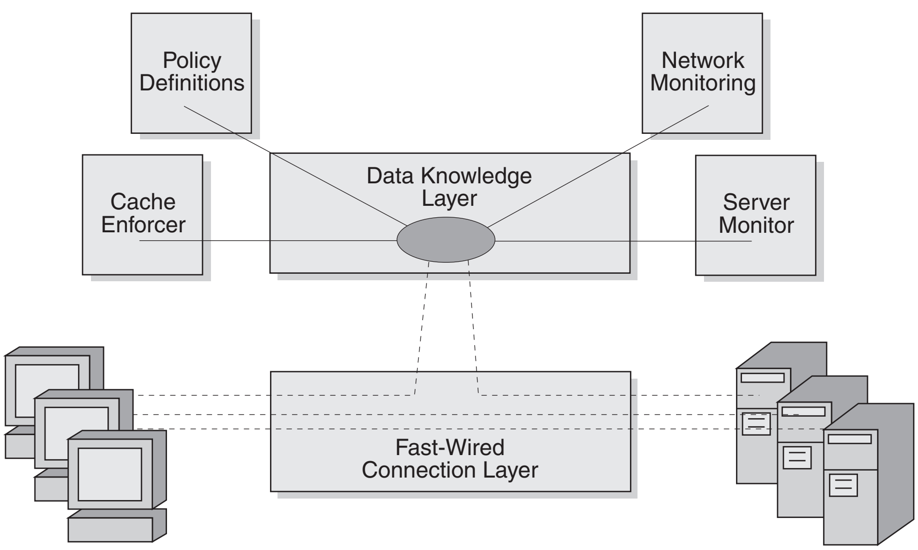

Figure 3. The NetEnforcer achieves wire-speed performance by presenting a two-tier approach to processing traffic. For the initial connection setup and advanced packet analysis requirement, the Data Knowledge Module will interface to various functions to determine resources needed for the connection. For most data flow, the traffic will travel through the fast-wire connection layer that will pass the traffic at wire-speed.

WHITE PAPER

US Offices
433 Airport Blvd., #303
Burlingame, CA 94010
Tel (650)-401-2244
Fax (650)-401-2277

Web www.allot.com

International Offices
5 Hanagar Street
Industrial Zone
Hod-Hasharon, 45800, Israel
Tel 972-(0)9-744-3676
Fax 972-(0)9-744-3626

Email info@allot.com

Europe
World Trade Center
1300, Route Des Cretes
BP 255
Sophia Antipolis Cedex
France 06905
Tel 33-(0)4-92-38-80-27
Fax 33-(0)4-92-38-80-33

Japan
Kowa Shinjuku Building 7F
2-3-12 Shinjuku
Shinjuku-ku
Tokyo 160-0022
Tel 81-(0)3-3355-0450
Fax 81-(0)3-3355-2445

Allot Communications

P/N D004001 REV 1A 05/00