Bandwidth Management for Peer-to-Peer Applications

With the increasing proliferation of broadband, more and more users are using Peer-to-Peer (P2P) protocols to share very large files, including software, multi-media files, and applications. This trend has exponentially increased traffic flows across a very wide area network.

If you are coping with excessive bandwidth consumption due to P2P traffic such as BitTorrent, eMule, numerous Gnutella clients, DirectConnect, Kazaa, etc., conventional rate shaping techniques such as limiting bandwidth by TCP port number may not do the trick. A more powerful technique based on application signature identification via packet inspection may be needed.

Traditional rate shaping techniques may not be sufficient to control new breeds of applications. For example, BitTorrent is a protocol that is typically used by simple desktops to transfer user files via broadband connections. However, using BitTorrent to transfer high volumes of data puts huge pressures on the broadband operators' network. Unfortunately, prohibiting BitTorrent traffic has become routine for some broadband operators and is now a key area of contention between users and broadband operators.

This paper describes how you can use F5 BIG-IP® iRules™ and the Rate Shaping feature of BIG-IP® Local Traffic Manager (LTM) to identify different types of traffic for individualized control that can return double-digit capacity without spending a dime on additional bandwidth. Through the combination of iRules and Rate Shaping, you can:

- Ensure that critical applications are not impacted by non-priority traffic.
- Deliver optimal application performance by allocating more bandwidth for higher priority applications
- Eliminate special purpose Rate Shaping products for simplified, centralized traffic management capabilities
- Provide flexible bandwidth limits, bandwidth borrowing, and traffic queuing
- Control rate classes based on any traffic variable
- Enable application bandwidth to be shared across similar priority applications for better resource sharing
- Ensure that specific types of application traffic stay within authorized boundaries

Re-gaining Network Traffic Control

Rather than using a one-size-fits-all approach to controlling network traffic, network managers need a more application-oriented way to transmit and distribute network data. In the case of BitTorrent traffic, F5 suggests:

Step 1 – Identifying BitTorrent traffic via packet inspection

Step 2 – Implementing a rule to isolate BitTorrent traffic

Step 3 – Assigning a rate shaping policy that only applies to BitTorrent traffic

With BIG-IP iRules and the Rate Shaping feature in the BIG-IP Local Traffic Management system, you can control the bandwidth usage of any type of traffic. Figure 1 shows how Rate Shaping can control the bandwidth usage of just BitTorrent traffic.

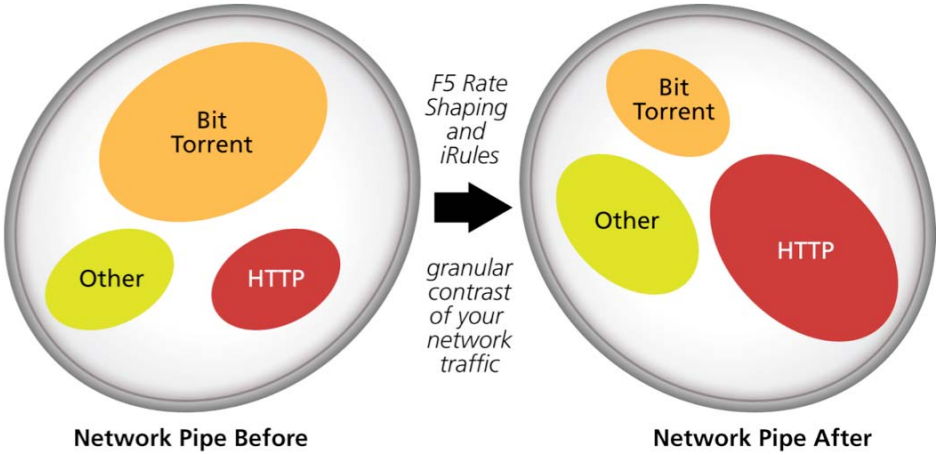

Figure 1: Controlling BitTorrent Traffic

The following sections describe each step of the process, provide a sample iRule to identify the BitTorrent application signature, and describe your options for controlling virtually any type of traffic.

Detection of BitTorrent Traffic

According to a recently published paper by AT&T Labs[1], inspection of the data packets that are transmitting between clients is a good way to detect BitTorrent traffic. The communication between BitTorrent clients starts with a handshake followed by a never-ending stream of length-prefixed messages. The header of the BitTorrent handshake message uses the following format:

<a character (1 byte)><a string (19 byte)>

The first byte is a fixed character with value "19," and the string value is "BitTorrent protocol." Based on this common header, you can use the following signatures to identify BitTorrent traffic:

- The first byte in the TCP payload is the character 19 (0x13)
- The next 19 bytes match the string "BitTorrent protocol"

Using BIG-IP iRules to Detect BitTorrent Traffic

BIG-IP iRules is a powerful yet simple tool you can use to identify and isolate the application traffic you want to direct, filter, or persist on. BIG-IP iRules gives you the ability to customize application switching based on business needs, optimizing the handling of traffic – where and when to send it for the fastest response based on application type, category, and priority.

The following example uses an iRule to intercept traffic and pinpoint when a TCP connection has initiated BitTorrent communication and manage only that traffic without affecting any other type of traffic.

[1] Accurate, Scalable In-Network Identification of P2P Traffic Using Application Signatures – AT&T Labs – Research

```
when CLIENT_ACCEPTED {
    TCP::collect 0 0   // start data collection after client TCP handshake connectic
is
                       // completed
 }
 when CLIENT_DATA {
    append payload [TCP::payload] // assign the collected contents in “payload”
    if {[string length $payload] < 6} { // pass directly if collected contents is le
                                        // than 6
          TCP::release
          return   //end Rules operation, and not carry out subsequent statements
      }
     TCP::release //release the collected contents and go along subsequent work
    binary scan $payload cc5 bt_size bt_protocol //analyze packet content obtained.
    if {($bt_protocol == "66 105 116 84 111") && ($bt_size == 19)} {
    log "Torrent traffic from [IP::remote_addr]" // add Log if it needs to record IP
    Rate Class p2p_bt //if pattern matches, put it in Rate Class ‘p2p_bt’ for
processing
    }
}
```

Once a TCP client is accepted, BIG-IP inspects the first packet’s payload of a TCP connection and looks for a match with the BitTorrent protocol signature. Using the BIG-IP Rate Shaping feature, you can assign a *Rate Class* that corresponds to the policy you define to control traffic with the BitTorrent protocol signature. In this example, if the TCP payload is a BitTorrent payload type, it is assigned to the Rate Class “p2p_bt.”

You can also target BitTorrent traffic for special processing, isolating it from all other traffic on the network including routing all BitTorrent traffic through a separate WAN link, limiting the amount of bandwidth devoted to BitTorrent traffic, or any combination of bandwidth control techniques described in this paper.

Once the connection is built, you can designate all the subsequent packets in the same client session as “p2p_bt,” using the BIG-IP s*ession persistence* feature. BIG-IP minimizes the degradation of switching efficiencies due to packet inspection because it doesn’t need to process every packet of a session beyond the first few bytes of the first payload packet.

Bandwidth Control

Some of the key bandwidth control functions used to manage user traffic by broadband operators include:

- Bandwidth limit of Peer-to-Peer (P2P) application to an individual user (IP)
- Bandwidth limit of P2P application to a group of selected users
- Bandwidth limit of specific application (BitTorrent, WWW, FTP, etc.) to an individual user (IP)
- Bandwidth limit specific applications to a selected user group
- Bandwidth limit the exit traffic according to application types

With the Rate Shaping feature, BIG-IP gives you the ability to:

- Limit bandwidth
- Control bandwidth bursting
- Limit bandwidth by direction

BIG-IP controls bandwidth per Rate Class, so you can control the traffic in a single Rate Class type to obey any and all rate shaping rules independent of the rules you specify for any other Rate Class. By combining bandwidth control functionality with an iRule that identifies and isolates specific types of traffic, you can control traffic in the following ways:

- Base throughput rate
- Absolute limit on the rate at which traffic is allowed to flow when bursting or borrowing
- Maximum number of bytes that traffic is allowed to burst beyond the base rate, before needing to borrow bandwidth
- Direction of traffic (any, client, server) to which the Rate Class is applied
- Rate class from which this class can borrow bandwidth
- Method that the Rate Class uses to queue and dequeue traffic

You can also define policies in each Rate Class for traffic flowing through any single or group of virtual servers and/or pools.

The following example shows the interface and properties for a basic rate class.

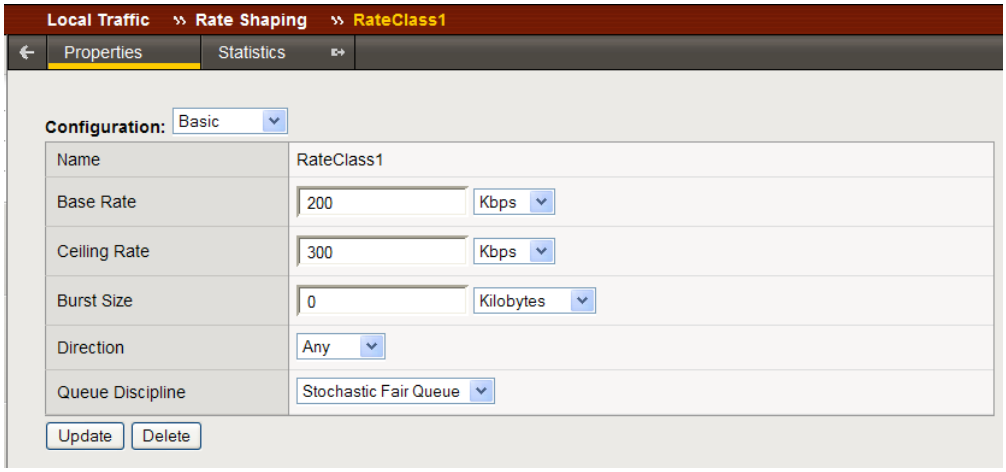

Limit Excessive, Non-critical Traffic

The typical service provider environment that facilitates P2P conversations includes a complex network of connections. These connections start from one end-user to an access network through the backbone of the core network to another access end user network, and then to the destination end users at a distant location.

A key junction point in these P2P connections, like those used with BitTorrent, is the junction point between the Metropolitan Area Network (MAN) router and the router connecting to the service provider's network backbone. From a traffic management perspective, these junction points are high-impact traffic management locations; thousands of users transverse this junction to access the service provider backbone to complete P2P file transfers. F5 traffic control at this junction point enables the network operator to manage thousands of users from one network device that is physically located at this junction point.

The customer referenced in this paper used an in-line deployment of BIG-IP LTM as a bridge between the MAN router and backbone router to manage BitTorrent traffic.

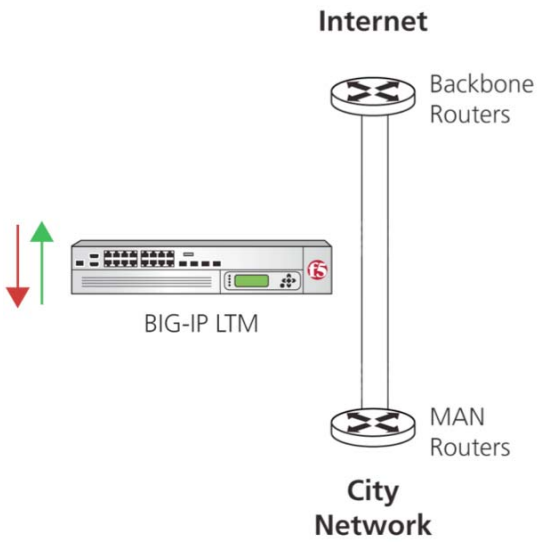

Figure 2: Network Topology Structure

In this configuration, BIG-IP LTM bridges Giga fiber interfaces. The direct uplink port of the original MAN switch is connected through BIG-IP LTM to the backbone router, whereas the BIG-IP LTM switch acts as a bridge. A direct line connects the two routers, which are configured as a low-priority backup.

In the case of broadband operators, the ability to limit excessive non-critical traffic from gaining access to the trans-city backbone dramatically improves nationwide traffic efficiency. For a specific application with difficult–to-identify characteristics, like BitTorrent, the ability to prioritize and limit the bandwidth consumption across thousands of users from a single location is extremely valuable. And since the location of these junction points is typically in central office type facilities, using BIG-IP LTM to define traffic policies is conveniently done from a single graphical user interface.

Measuring Performance Improvements

Baselining

Implementing traffic policies starts with measuring and documenting the baseline performance of your “untuned” network. You can use any monitoring solution (MRTG, Cacti, Cricket) to baseline the performance of your network.

Bandwidth Consumption Baseline

Prior to creating policies to manage specific types of traffic, measure the traffic load passing through the BIG-IP switch against the baseline performance. You can configure a simple monitoring solution to draw curve diagrams of input/output traffic change through the system. Figure 3 shows an example of a data traffic diagram collected by a typical traffic performance monitor system.

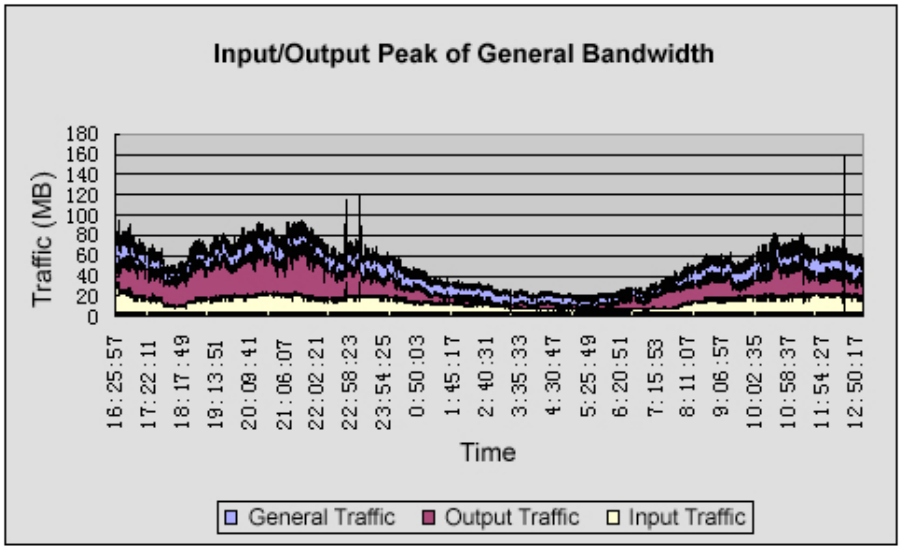

Figure 3: Input/Output Peak of General Bandwidth

Note that the traffic through this network junction is high from 8:00 to 24:00, and during this time period, traffic exceeds 60MB/s quite often.

Bandwidth Consumption of Specific Applications

Baselining traffic across the network also involves the graphing of traffic throughput by application. Monitoring software can analyze typical Internet applications to determine what type of applications merit their own Rate Class. Figure 4 shows the traffic consumption of FTP and WWW traffic that you can use to determine if these types of applications are using a disproportionate amount of bandwidth.

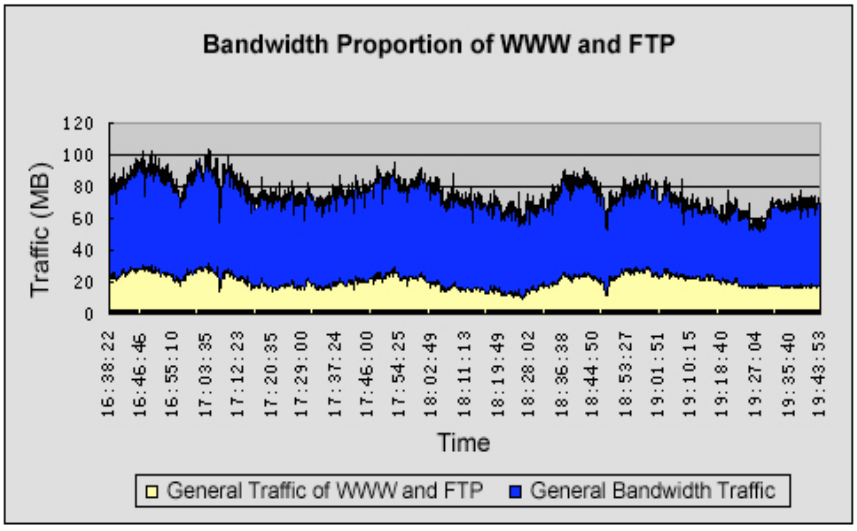

Figure 4: Bandwidth Proportion of WWW and FTP

In many situations, you may want to prioritize some types of applications as higher or lower priority. For instance, you may want to prioritize WWW and FTP traffic as lower priority traffic. However, look at the volume of WWW and FTP traffic before applying this kind of traffic policy to understand the impact of changing the traffic priority of these types of applications.

Bandwidth Control Implementation

After baselining, you can use a combination of F5 BIG-IP iRules and the Rate Shaping feature of BIG-IP LTM to improve application performance. The following sections describe policies you can create to limit bandwidth for:

- P2P traffic
- WWW applications
- Multiple types of applications

Bandwidth Limiting P2P Traffic

A common traffic management rule is limiting bandwidth for specific applications that are lower priority or consume excessive bandwidth when not controlled. In the case of service providers, bandwidth limiting P2P applications in one or more network segments and/or users is an effective way to manage this type of traffic.

With F5, you can bandwidth limit only P2P traffic and select the users to which this traffic policy applies using:

- BIG-IP iRules to analyze P2P traffic in certain IP network segments and assign certain users running P2P applications to a unique Rate Class
- BIG-IP Rate Shaping to define a policy that limits the P2P application bandwidth of the IP network segment for that Rate Class

What if you want to restrict P2P traffic to a limited amount of bandwidth only during peak use? Once configured with a Rate Class, you can watch the traffic change on the user's monitoring system, as shown in Figure 5.

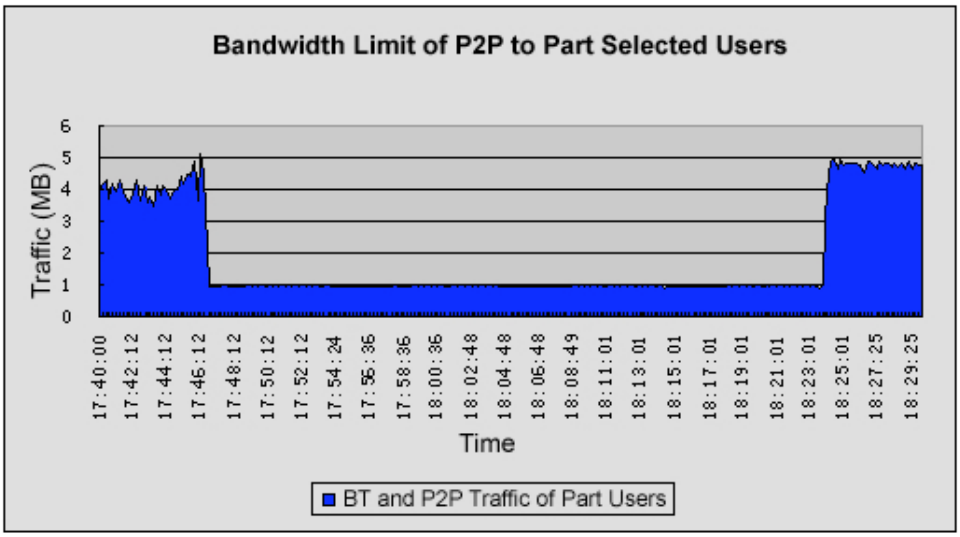

Figure 5: Bandwidth Limit of P2P to Part Selected Users

In Figure 5, users' P2P download speeds were limited to less than 1K/s per user from just after 17:46:12 to just past 18:23:01. During the time period, total download traffic for all users was limited to 1M /s for the network segment under control.

Bandwidth Limit of WWW Applications

With BIG-IP LTM, you can bandwidth limit only WWW applications. To analyze the effect of this kind of traffic policy, do the following:

- Write an iRule to identify WWW applications and assign HTTP traffic to a separate Rate Class.
- Using BIG-IP Rate Shaping, create a Rate Class to limit user application bandwidth for HTTP traffic.
- Watch the traffic change on a network monitoring system.

In Figure 6, the user's HTTP traffic remains within a pre-defined range between 19:22 and 19:46.

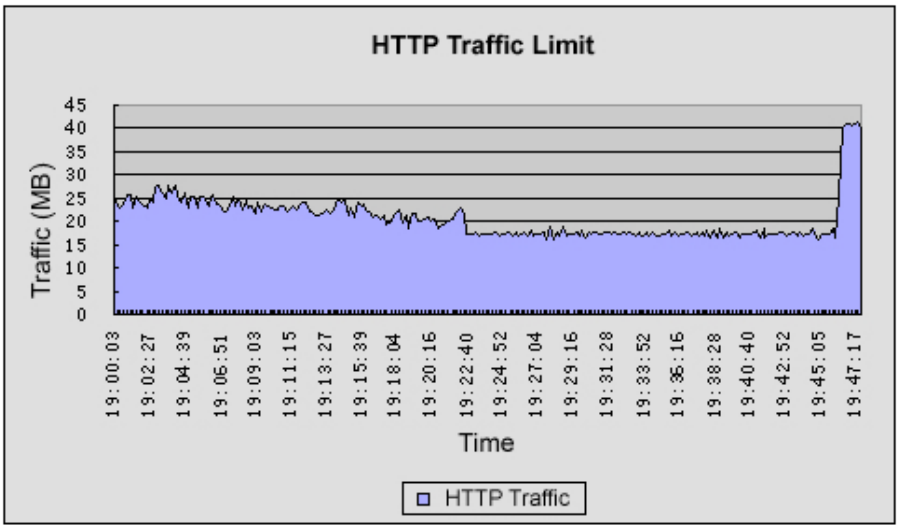

Figure 6: HTTP Traffic Limit

When HTTP is limited, pages open more slowly and HTTP application performance decreases. Once the bandwidth limit is lifted, users' HTTP traffic rises to around 40MB/s when the limit is canceled at around 19:46. HTTP traffic was limited to 18M/s while the traffic policy is in force to prevent degrading users' non-HTTP applications during the test.

Bandwidth Limiting Multiple Applications

With BIG-IP iRules and BIG-IP Rate Shaping, you can customize traffic throughput for different types of applications. To do this, write an iRule to identify each type of application, create a Rate Class for each type of application you want to control, and then use the iRule to assign each Rate Class to the appropriate type of traffic.

The following table lists different policies that highlight the flexibility of BIG-IP Rate Shaping, giving you the ability to specify an infinite number of policies to manage traffic and optimize network resources.

Time Span	Policy
18:30:14 - 18:45:14	None (peak usage starts at around 19:00)
18:45:14 - 18:49:27	Limit HTTP to 5 Mb/s. Reject all BitTorrent traffic
18:49:27 - 19:00:27	Reject all BitTorrent traffic
19:00:27 - 19:11:27	Reject all UDP traffic Limit BitTorrent and Other P2P traffic to 5 Mb/s with Other P2P having priority
19:11:27 – 19:22:40	Limit BitTorrent to 3 Mb/s Limit UDP to 1 Mb/s Limit Other P2P to 1 Mb/s
19:22:40 – 19:29:52	Limit total BW to 17 Mb/s HTTP is given highest priority Limit BitTorrent to 3 Mb/s Limit UDP to 1 Mb/s Limit Other P2P to 1 Mb/s
19:29:52 – 19:37:28	Limit total BW to 17 Mb/s Limit UDP to 1 Mb/s Limit Other P2P to 1 Mb/s
19:37:28 – 19:41:04	Limit total BW to 17 Mb/s Limit Other P2P to 1 Mb/s
19:41:04 – 19:44:41	Limit total BW to 17 Mb/s
19:44:41 Onwards	None

Figure 7 illustrates the effect of four different policies used to manage four different types of traffic.

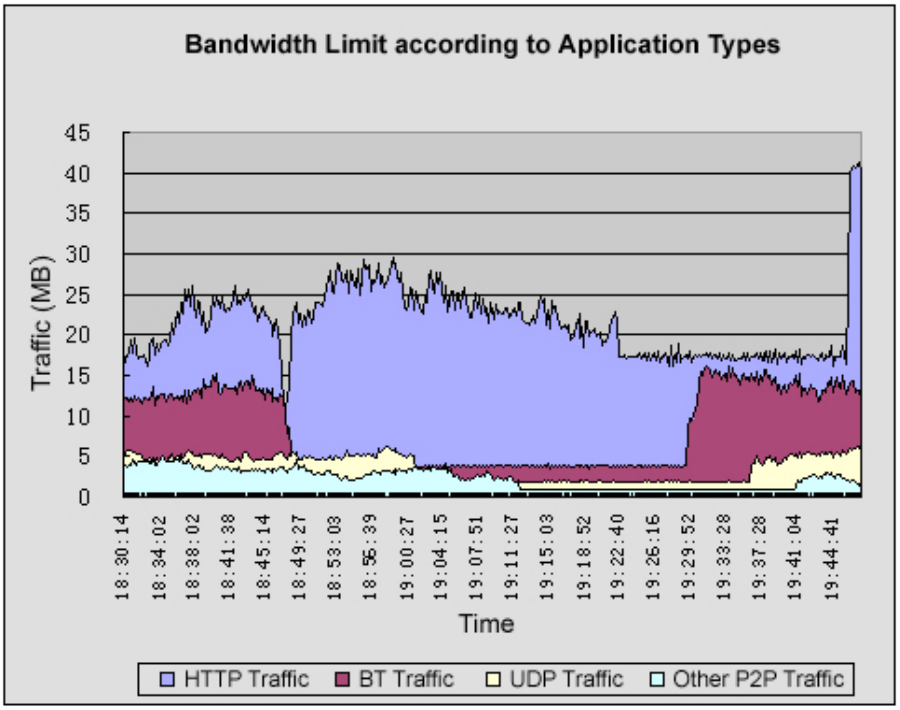

Figure 7: Effects of Bandwidth Limit Policies in Test Cases

Customer Implementation

The customer that was faced with the broadband challenge configured BIG-IP LTM to limit bandwidth by application between the MAN and the backbone. They used separate Rate Classes to manage traffic limiting:

- P2P traffic at 4MB/s
- All other HTTP traffic at 20MB/s
- eMule users at 1MB/s (detection of eMule's traffic is similar to detecting BitTorrent traffic whereas the first character of each payload packet is 0xE3, (*Source: AT&T Labs - Research*)
- All other P2P traffic to 1 MB/s

Figure 8 shows the traffic throughput before and after implementing BIG-IP LTM policies for four different types of traffic.

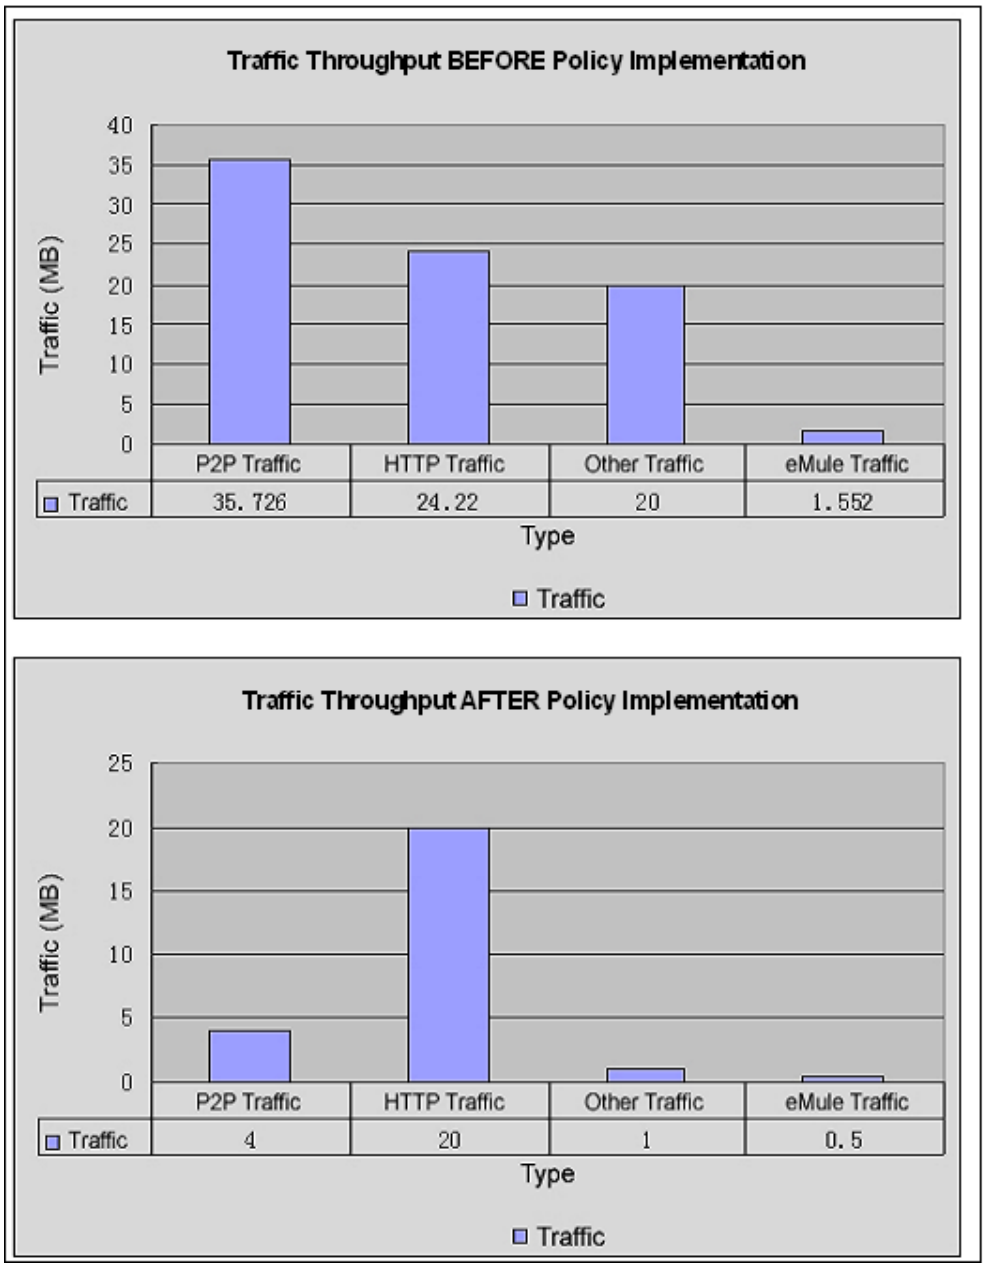

Figure 8: Effects of Traffic Management Policy Implementation

Conclusion

F5 gives service providers the ability to deftly manage the impact of otherwise uncontrollable user applications. The key capabilities that enable service providers to regain network control include:

- **BIG-IP iRules** – identify specific types of traffic for precise control. iRules enable BIG-IP LTM to read packet contents, identify traffic signatures within the packet, and assign all traffic with that application signature to a unique Rate Class. With iRules, you can identify the type of traffic and assign a rate class to control that type of traffic on any traffic flow variable.
- **Rate Shaping** – BIG-IP Rate Shaping gives you the power and flexibility to manage specific types of traffic in a variety of different ways. Because Rate Shaping is built on F5's TMOS full application proxy architecture, you control throughput in any direction (inbound, outbound). With Rate Shaping, you can create different traffic policies for each individual Rate Class to control and prioritize bandwidth usage for different types of traffic.

Although this paper focused on broadband issues, Rate Shaping capabilities also include:

- Traffic limiting, prioritization, and borrowing
- Maintaining enough bandwidth for high-priority applications and traffic
- Defining traffic and application limits
- Controlling the rate at which those resources are allowed to spike or burst
- Full support for bandwidth borrowing
- Traffic queuing (stochastic fair queue, FIFO ToS priority queue) to prioritize traffic types
- Granular traffic classification L2 through L7

About F5 Networks

F5 Networks is the global leader in Application Delivery Networking. F5 provides solutions that make applications secure, fast, and available for everyone, helping organizations get the most out of their investment. By adding intelligence and manageability into the network to offload applications, F5 optimizes applications and allows them to work faster and consume fewer resources. F5's extensible architecture intelligently integrates application optimization, protection for the application and the network, and delivers application reliability – all on one universal platform. Over 16,000 organizations and service providers worldwide trust F5 to keep their applications running. The company is headquartered in Seattle, Washington with offices worldwide. For more information, go to www.f5.com.