

How Opteq Compares with its Competitors

Opteq specializes in high-end solutions for ISP's who, are dissatisfied with the limited scalability, complexity and functionality of current performance management solutions offered on the market. Our product is an integrated modular resolution that provides unsurpassed control over network quality management. Unlike competitors we have assembled all the tools required for ISP's to optimally manage their network.

Listed below are a few examples of how Opteq might help your ISP business

- Network Visibility: Full layer 7 visibility of network traffic
- Real-time Monitoring and alerting which helps identify issues on the network before they become potential problems
- Bandwidth management allocating bandwidth to services based on current average demand
- High Performance ISP Caching (WCCP & Transparent)
- Intelligent load balancing
- Complete bandwidth management & subscriber provisioning
- Top Ratio link compression
- Server acceleration / reverse proxy
- Content filtering and SMTP scanning
- VPN and VLAN creation / termination

Scalability is the biggest challenge in bandwidth management, and we've addressed that with a highly efficient multiprocessor solution that leverages the latest in computing technology. Opteq quad-core systems can manage multiple gigabit links without concerns of overrunning the CPU at peak times.

We are often asked how our traffic management differs primarily from competitors Let us tell you how:

Opteq removes all congestion, drops and retransmission from the network. It does this by effectively pacing traffic of all types to the link speed and unlike the competitors it does this without the need to queue traffic.

Opteq's host pacing technology emulates separate links per user, controlling throughput bidirectionally to the bit and window shaping / ack pacing across these to further control congestion. Most competitors do not do use TCP window manipulation at all! This is the only natural way to reduce the number of packets on the network at any given time and therefore reduce congestion allowing higher priority traffic a free passage through the network.

The fastest way for an IP packet to traverse any link is when there is zero congestion between the client and host. I.e.: no packet queues in the router or anywhere along the route from client to server. Applications then travel at wire speed.

By entering the WAN link speed into the Opteq device its fair allocation of bandwidth algorithms will automatically and fairly split bandwidth amongst the users. This is often all that is required to greatly increase the performance of the network and improve the end user experience.

One can bias performance and create tiered services per user by using the rulesets, dynamically pacing non latency sensitive protocols further and allowing high priority traffic to always have a free lane through the network.

Almost all Internet communications have a client/server model where the client is sending requests and the server is sending data. This is true for ftp transfers, streaming video and streaming audio. Even if the client and server are sending UDP packets there is always a client/server relationship of some form.

Opteq paces hosts by controlling client requests so that higher priority traffic always has a free path through the network and never gets caught in a queue or buffer shuffle. It is proactive in that it paces all traffic to the speed of the WAN link so that there is no congestion at all along the traffic path. It monitors the flows of traffic and dynamically ensures high priority traffic has the bandwidth it requires.

This is radically different from the queuing methods employed by Cisco's p-cube , Allot's per flow queuing – and other so called "packet shaping tools". It is the only method that allows you control over Internet traffic coming into your network.

The biggest advantage to pacing client requests is that you get at the source of traffic so that traffic flows at the given wire speed without forming queues or buffers. This results in much happier clients and much higher contention ratios are achievable without degradation to quality.

Queuing methods used by competitors are reactionary. Once the link "fills up" and reaches the given wan speed , then packets can no longer traverse the WAN at link speed and therefore enter the routers / shapers buffers and are Queued in order to be prioritised. Only Outbound traffic is queued. By queuing you are adding latency as you have to delay the packet in a buffer in order to prioritise it and let it back onto the wire. Additionally queuing methods can do nothing about the rate of incoming traffic. Within an ISP this is the majority of traffic.

Competitors will argue, and it is true that TCP and UDP will react to this congestion (given time) themselves but the point is that there is already congestion which is principally what you are trying to avoid if you are concerned about the traffic quality. Buffers overflow and traffic is dropped on the floor and in the case of UDP packets are lost forever or retransmitted adding to congestion if its a TCP stream.

The problem with DPI alone

To control user activity competitors require many rules and DPI for application recognition. However policies based on explicitly having to identify the application are problematic as there is always going to be unidentified traffic as signatures change or worse still traffic becomes encrypted. This traffic is then thrown by competitors into an "all other " classification and managed in a single umbrella rule. Multiple traffic types some good some bad having to compete for restricted bandwidth. There are many legal forms of p2p downloading as well which get restricted by these general catchall shaping rules.

Although Opteq has powerful layer 7 filtering capabilities it does not rely on application filtering to control network hosts. Through Opteq's fair share technology it does not matter what the traffic is or if its encrypted or anything as the algorithms at this level fundamentally work per host. So should there be any unidentified traffic (and there always will be) Opteq's unique host pacing and fair allocation of bandwidth automatically ensures fairness across the network and ensures no one user gets an unfair share of bandwidth.

In the end a user using all his bandwidth on encrypted p2p for example will get the same amount of bandwidth as an elderly person trying to do his or her internet banking.

While "other" vendors promote micro-managing protocols on your network, our solutions promote net neutrality with a more efficient and less labour-intensive solution. The Opteq solution gives you the control you need without the hassles.

Unlike most competitors Opteq can:

- Scale significantly without hardware upgrades
- Manage traffic in both directions from a single appliance
- Eliminate queues and therefore congestion
- Manage users fairly by default – i.e.: Opteq's algorithms manage users fairly no matter what application they may be using
- Scale to include other core services (DNS, WCCP cache, firewall, VPN)
- Compress traffic between POP's
- Much more as above ...

Internet caching

More than two thirds of the traffic on the Internet today is generated by the Web. Web caching can play a valuable role in improving service quality for a large range of Internet users.

In terms of the ability to improve the service performance of delivery of content to a global network of clients, and in terms of the ability to improve the carriage efficiency of the network, caching of content makes sense to the content provider, to the ISP, and to the end client.

Page hit rates of somewhere between 40 to 55 percent are achievable for a well configured cache resulting in a reduced byte transmission of 15% or more.

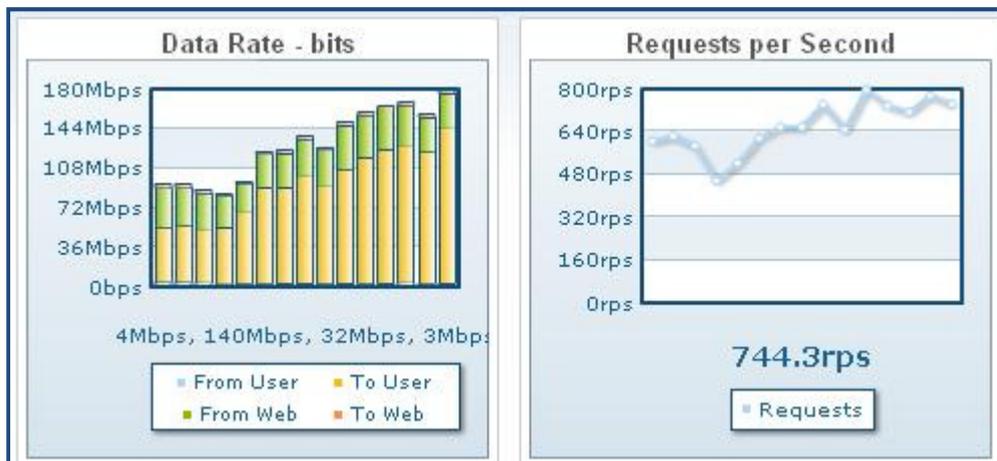
Networks utilising cache appliances see a radical improvement in latency. Typical results reduce latency from 300ms to < 10ms for cached items appearing as improved link speed to the end user.

In cable networks where there is a high speed copper or fibre by caching closer to the users cached objects may be delivered at many multiples of Mbps as opposed to 100Kb for example on non cached data across the world wide web.

Opteq provides scalable caching performance for extreme workload capacity for mid range and larger network environments. The Opteq Cache is the ideal choice for large enterprise or service provider networks with dedicated Internet or WAN uplinks from 25 Mb/s to 155 Mb/s with heavy Internet, intranet or web use.

The option to provide high performance "on-appliance" anti-virus and Web content filtering further extends the value of the appliance as an edge Internet or deeply distributed Intranet accelerator.

An example of the cache at work



The cache above is serving in excess of 700 requests per second at around 170 Mbps throughput. Of that traffic the Opteq is currently sending 140Meg to clients served from the cache and downloading only 32Meg from the internet at 750RPS – that is an ENORMOUS saving and the client is extremely happy.