



# QoDSL Technical Introduction

Date: March. 2005

Author: Sonia Hanson

Version: 1.1





## 1. Introduction

### Voice and data convergence

Many large enterprises use IP virtual private networks (VPNs) because they offer a cost effective method for data and applications to be shared securely between multiple sites.

The introduction of Multi Protocol Label Switching (MPLS), means IP networks can now differentiate between different types of data. Voice or live video packets, for example, can be designated “high priority,” whilst less time-critical data such as email can be given a lower priority and transmitted whenever sufficient bandwidth is available.

The significance of MPLS may not, at first glance, be obvious, but it is the key to “Class of Service” performance prioritisation which can give voice traffic the guaranteed throughput needed to provide clear, uninterrupted, toll-quality calls.

## 2. Quality of Service

### Packet prioritisation

The ability to classify and prioritise data packets from particular applications such as voice or video is commonly referred to as Quality of Service (QoS), and MPLS-enabled networks are able to deliver QoS

IP VPNs which can deliver QoS are often used by larger enterprises to link large sites. But because until now the technology required a relatively expensive leased line to operate over, it has not proved cost effective for smaller sites or organisations

### Drawback of DSL

Broadband network connectivity based on DSL technology is now commonplace and the number of enterprises using low cost DSL-based IP VPN connections to link smaller sites - and even home workers - to the corporate LAN infrastructure has grown dramatically over the last three or so years.

But while standard DSL based IP VPNs are ideal for standard data applications such as email and web browsing, they are not suited for voice applications because the necessary level of network performance cannot be guaranteed. If data intensive applications are being used, or there is network congestion, performance can degrade to the point where the voice is rendered unusable.






### QoDSL: bringing QoS to DSL

Tiscali can provide a solution to this problem with QoDSL, which combines the low cost of a DSL-based managed VPN with the QoS features normally associated with expensive leased lines.

With a QoDSL VPN, customer premises equipment (CPE) and the provider core network identify and route packets based on an assigned priority, ensuring that voice and other time-sensitive packets are transmitted with the minimum delay.

## 3. End-to-end QoS

How it works:



### LAN QoS

Data from the customer LAN arrives at the CPE with ToS settings (Layer 3) or, when not possible, Layer 2 which Tiscali converts to Layer 3. The CPE then reads the ToS bit classification and places the packets in one of two queues: the high priority Low Latency Queue (LLQ) for voice (and video) applications, or the normal priority queue for other data packets.





## QoS Over DSL

Tiscali perform Egress Queuing (Layer 3), again with two queues.

## Tiscali Core Network QoS

Cisco 1000 series switches sit at the edge of the Tiscali core network - known as the Provider Edge (PE). These switches determine how to label MPLS packets based on their ToS bits and queue the labeled packets accordingly.

**NOTE:** MPLS labels, not ToS bits, are used to provide QoS across the Tiscali core. Significant ToS bits such as those indicating IP Precedence are incorporated into the experimental bits of the IP packet and these experimental bits are then read to assign packets to the correct queues.

As each packet leaves the PE, Layer 2 CoS applied.

## QoS Over DSL

After traversing the Tiscali core network, the PE at the other end removes the MPLS label from each packet and assigns it a new ToS bit. Packets are once again placed in either low latency or normal queues.

## LAN QoS

The customer CPE classifies inbound packets based on their ToS settings. There is the option to write layer 3 ToS (our preference) or Layer 2 CoS.

## 4. Why QoS is vital for voice applications

### Factors affecting quality

Four factors affect the quality of voice applications run over a network:

- Latency
- Jitter
- Packet loss
- Minimum guaranteed bandwidth

### Latency:

This is defined as the time needed to:

- A data packet to move across the network to an endpoint
- Encoding and packetization at the transmitting end
- De-jittering and de-coding at the receiving end

For voice traffic to remain intelligible it can not be subject to any significant degree of delay. If there is too much traffic on the network, or if a voice





packet gets stuck behind a large data packet (such as an email attachment), the voice packet will be delayed to the point that the quality of the call is critically compromised. The maximum amount of latency that voice packets can tolerate whilst maintaining toll-quality is 150 milliseconds.

**Jitter:**

When a packet stream travels over an IP network, there is no guarantee that each packet in the stream will travel over the same path, as there is in a circuit-switched network. As they do not take the same path, intervals between packet arrival times vary since one packet may take more “hops” than others, delaying its arrival considerably and causing it to have a much higher latency. Even when packets take the same path, the amount of time taken to cross a given router can vary considerable from packet to packet because of changes in congestion levels, preceding packet sizes and other stochastic processes.

In order for voice to be intelligible, consecutive voice packets must arrive at regular intervals. Jitter describes the degree of variability in packet arrival intervals. Voice packets can tolerate about only about 75 milliseconds of jitter delay.

To lessen the effects of jitter, packets are gathered in a jitter buffer at the receiving end. The jitter buffer must be sized to capture an optimal proportion of the data packets while keeping the effective latency as low as possible.

**Packet loss:**

Packet loss is a common occurrence in data networks, and conventional devices and applications simply request the retransmission of lost packets. However, since time is critical in voice applications, lost voice packets are not retransmitted.

Voice traffic can tolerate less than a 3 percent loss of packets before callers experience disconcerting gaps in conversation.

**Minimum Guaranteed Bandwidth:** Bandwidth is the raw data transmission capacity of a network, and inadequate bandwidth causes both delay and packet loss. Since IP traffic is irregular, packets will often suffer delay in the absence of some method of prioritisation.






Each service must get a guaranteed minimum throughput. That minimum is determined by its service class and must be met, irrespective of the Virtual Path (VP) congestion or other service classes using that VP

Latency must be minimised for all customers operating below their guaranteed throughput. This is achieved by ensuring there is no queuing of traffic into the VP for customers operating at or below their guaranteed throughput levels. This is particularly important for the successful implementation of voice applications.

The diagram below shows how this is achieved on an ADSL circuit. Effectively, there are two queues:



**High Priority:** This bandwidth is always available to traffic marked as high priority

**Other Traffic :** Other traffic on the ADSL connection can use the remaining bandwidth. If the High Priority bandwidth if it is not being used it can be used by other traffic when required

#### Packet shaping

To guarantee minimum bandwidths Tiscali use a method called Hierarchical VP Shaping (HVS), which shapes both the VP and Virtual Circuit (VC).

VP shaping is achieved by slowing down individual VCs when VP congestion starts to occur. When a VP becomes congested a feedback mechanism instructs the VCs to stop transmitting until the congestion dissipates.

Within the HVS each VC is weighted, and the weights are used as follows:

- During off-peak times when VPs are not congested, all VCs have their traffic demands met. All cells presented by the VCs are de-





queued and transmitted down the VP. The weighting does not influence traffic flows.

- ❑ During peak times when VPs are congested, the available network resources are shared between the active VCs according to their weights: the higher the weight of a VC the more network resources are allocated to it.
- ❑ Each VC is guaranteed a minimum level of throughput. That minimum value is calculated based on its weight.
- ❑ No VC can get more than its share of network resources.

### **QoDSL controls these factors**

Since QoDSL provides a way to constrain latency, jitter and packet loss to acceptable levels and to guarantee set minimum bandwidths, voice applications can be made to deliver the same toll-quality calls that customers are accustomed to from dedicated voice networks - even over a contended access technology such as ADSL.

## **5. Optimising QoDSL**

### **Controlling delay**

Controlling delay is key to optimising QoDSL, and delay should in all cases be kept below 150ms. Since jitter increases effective latency, it is important to take steps to control both latency and jitter, and because the size of a jitter buffer directly affects the latency perceived by users of voice applications, networks must be provisioned both for low latency and low jitter at both ends of the network.

### **Reducing End-To-End Delay Using Classification**

To reduce end-to-end delay, low latency packets can be given a higher priority at Layer 2 and Layer 3. To achieve this, action needs to be taken at every hop during transmission:

- **Customer LAN**

Customers must ensure that the QoS extends across their own network infrastructure - defined as every hop between the Tiscali CPE and the network end point. This can be achieved by ensuring that all routers and switches on the customer's LAN are configured for Low Latency Queuing (LLQ) and CoS priority switching. If a customer is unable to do this then end-to-end latency, jitter and packet loss may compromise call quality.





- **Customer LAN To CPE**  
Data entering the CPE is classified as low latency traffic based on the presence of either diff serve explicit forwarding (DSCP EF) or IP Precedence 5 marking in the ToS section of each data packet. Both match statements will be active simultaneously and invoke the low latency (priority) queuing of that data packet.
- **CPE to PE and across core network**  
These two methods of data classification are used on the same data in the same way at the PE entry point to ensure correct prioritisation of the data across the Tiscali MPLS VPN core network.
- **CPE to customer LAN**  
Finally, the same two methods of classification is used in the reverse direction to prioritize and Layer 2 tag (802.p, CoS) data from the CPE into the customer's network.

## Reducing End-To-End Delay Using Prioritisation

- **Priority queuing**  
Dual output data queues are implemented with strict priority de-queuing of high priority data traffic. The de-queued data will then be transmitted onto the DSL connection. Any packet marked "high priority" entering a network element will immediately be placed in front of normal-priority data so that it gets first call on the network resources and will not be delayed by normal-priority packets.
- **Reducing delay at an endpoint**  
Several techniques are used to reduce delay at an endpoint:
  - Optimisation of jitter buffering – configurable from an IP phone
  - Optimisation of packet size – configurable from an IP phone
  - Use of a low compression codec. Tiscali recommends G729
  - Ensuring network protocol stacks are efficient and correctly prioritised for low latency traffic
  - QoS-enabling customer LAN
- **Compensating For packet loss**  
Tiscali can guard against packet loss by guaranteeing precise throughput levels to customers. This is achieved using VC weighting and Hierarchical Shaping as explained above.





## CONCLUSION

Voice and data have traditionally been kept apart, with separate networks, different staff to support those networks, and different budgets to pay for them. But today voice can be regarded as nothing more than another data network application – albeit one which has certain special characteristics.

Voice over IP (VoIP) implementations are becoming increasingly common in large organisations, both at single sites (over a LAN) and between sites linked with dedicated leased line wide area networks (WANs).

Until now it has been impossible to use VoIP applications reliably over lower cost DSL based connections because of the absence of suitable QoS capabilities

But by understanding the characteristics of voice traffic and applying the appropriate Tiscali QoDSL mechanisms, small and medium sized companies can now successfully implement VoIP systems over low cost DSL connections, unifying their communications infrastructure, reducing their communications costs and harnessing the business benefits of VoIP without the high costs associated with leased line connections.

