QoS Bandwidth Management

QoS Bandwidth Management

Introduction

As Voice over IP (VoIP) emerges as the future of voice communications, questions remain about its quality and security. Put simply, the actual capacity of an IP network to carry voice is poorly defined and not well understood. Congestion, packet loss and delays in an IP network can adversely affect the Quality of Service (QoS) for the subscribers. Therefore, mechanisms to manage bandwidth usage must be utilised to maximise the service quality.

The bandwidth calculation 'per call' is simple; however, meshed, router networks are not. It is not always clear how to use these figures to produce a deterministic and consistent VoIP service. Compounding this calculation is the differing 'per call' bandwidths resulting from a variety of codecs that subscribers may use. When all of these factors are mixed together in a real network, often only experimental techniques can really determine actual network capacity.

However, session border controller technology can provide a pragmatic approach to this quality-related problem through the use of advanced bandwidth management techniques. The Newport Networks 1460 session border controller addresses these issues by limiting the amount of traffic in the network, policing existing calls and rejecting new sessions as appropriate, thus underpinning call quality within the network.

Bandwidth Management

There are three main components within bandwidth management:

1. *Session Admission Control (SAC).* A model within the session border controller holds the details of resources and capacities within the network. Whenever a new call is attempted, the model is examined, the appropriate resources reserved for that call and the total available resource decremented appropriately. If there are no resources available in the model, the call is rejected. Typically, this aspect of the bandwidth management capability is termed Session Admission Control.
2. *Policing.* The call establishment signalling protocols enable the subscribers' terminals to automatically negotiate the codec type that will be used for the duration of the call, and hence, the call's data rate. Each call is policed individually against the data rate of the negotiated codec. Any traffic significantly above the negotiated rate will be discarded. This, in turn, underpins the accurate and consistent view of network resources offered by the SAC function.
3. *Anti-Tromboning.* This feature optimises the use of the access network. For example, in an IP-Centrex service environment, the call may be most effectively transported entirely within the subscribers' private network. Therefore, calls that are to remain within the subscribers' own private network should not be included in the access network accounting mechanisms.

Session Admission Control (SAC) and the 1460 Session Border Controller

The approach taken to guarantee quality is to model the network being used and then limit the resources used for calls. The limits are set in order to meet commercial objectives, or to meet network limits, such as access network bandwidth.

The following example is intended to illustrate the network model and show how resources are accounted for as calls are established. The customer is a national supermarket chain with a corporate IP-Centrex service. The traffic is identified by a VLAN tag in a 'Metro Ethernet' environment. Each store is a different size and, therefore, may have different traffic requirements.

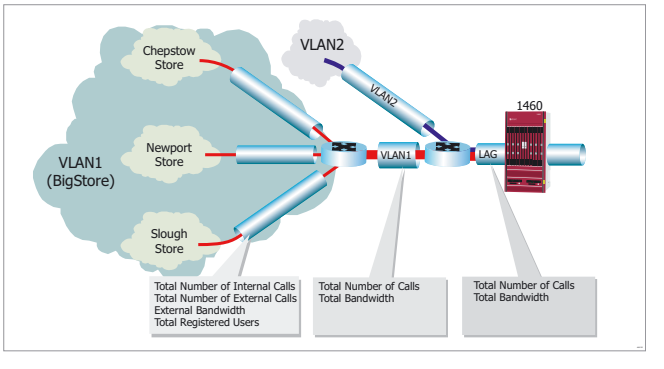

Figure 1 - Multi-Level Session Admission Control Model

In this scenario, the 1460 session border controller can limit the total number of in-store VoIP calls by controlling the service-signalling to and from the Service Provider; this is analogous to buying a PABX with a 20 line capability. Each store can also be limited to a maximum number of external calls placed and registered subscribers connected to the service; this is analogous to the number of external lines that the customer has purchased.

Since the access network between the subscriber and the Service Provider's core network is often a point of congestion, its usage needs be carefully controlled for technical reasons as opposed to commercial ones. By applying limits to the servicing VLAN, as identified by the VLAN tag, the overall usage purchased by the subscribers can also be limited. In this way, a Service Provider can tailor the subscribers' package as a whole and on a site-by-site basis.

Bandwidth capacities through the 1460 session border controller are modelled within a Link Aggregation Group (LAG) - an aggregation of physical links that connect to a single network. LAG Session Admission Control restricts the number of calls and the bandwidths used in order to remain within the physical capacity limitations of the 1460. This is essential in enabling physical resilience within the LAG. For example, the LAG could be made up of three gigabit Ethernet connections, where the capacity of one of these connections is reserved for resilience. In this case, the SAC limits traffic to be equivalent to two gigabit connections (out of the three available).

Thus, the 1460 session border controller's network model has three stages: subscriber site, total subscriber usage and physical network capacities; at each stage, resource utilisation is accounted for independently. Any over use (or potential over use) results in the new call request being gracefully rejected through call signalling protocol.

Session Admission Control (SAC) and SIP Signalling

SIP signalling uses an 'offer/acceptance' mechanism to negotiate the codec type to be used for the call. The SIP 'INVITE' message contains a prioritised list of codecs that the calling party offers to the called party. The called party responds with the preferred codec choice and the call proceeds using this codec.

However, there is no mechanism in SIP to reject a call after responding to the initial 'INVITE'. Therefore, the 1460 performs the resource calculations in two phases:

1. When the session border controller proxies the initial 'INVITE' message, the most bandwidth-intensive codec from the offered list is used for the first SAC calculation. If this SAC calculation indicates that the system cannot accept the new call, it is rejected.
2. If the call can be accepted using this 'worst case' codec, when the called party makes its response, the network model is updated with the bandwidth used by the accepted codec.

Media Policing

Once a call has been admitted, the 1460 session border controller uses media policing to underpin the SAC calculation and protect the quality by limiting the bandwidth use per call. This feature also reduces service theft by limiting media bandwidths to that authorised in the signalling setup. In a carrier-class session border controller, such as the Newport Networks 1460, over 100,000 calls can be policed independently, using sophisticated hardware acceleration.

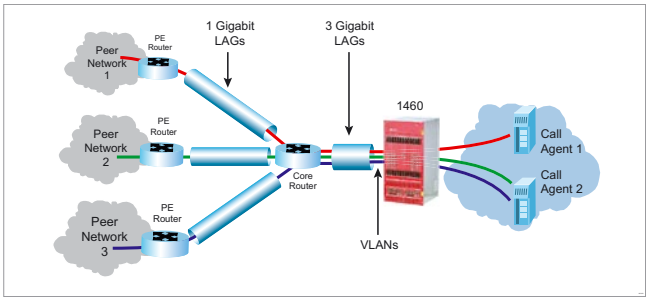

Figure 2 - Media Policing

The policing mechanism restricts data and packet rates based on the negotiated codec type with excessive data/packets being discarded to control the consumption of network resources. This is an extreme, but effective, action where the discarding of data will substantially disrupt the media of any call that exceeds its negotiated bandwidth.

Additionally, the 1460's policing mechanism will police packet sizes to protect against some forms of malicious attack. For example, if the agreed codec type is G.711 with a 10 ms sample, larger packets, such as G.711 with a 20 ms sample, will be discarded.

The 1460's policing policy is independently set for each call from each subscriber, even if the subscribers are behind the same firewall or using the same IP-PABX.

Anti-Tromboning

If a group of subscribers have their own private network, it is often desirable to keep bandwidth-intensive media flows within the subscribers' network, yet maintain signalling control of the call by the Service Provider. Enabling only the media to flow locally is termed 'anti-tromboning'. IP-PABX installations naturally achieve this, but the media traffic of IP-Centrex subscribers is normally 'tromboned' from the private network, through the access network and back to the private network, unless the session border controller uses precautionary measures to avoid this.

The 1460 maintains a 'context' for all registered IP-Centrex subscribers that might require the anti-tromboning feature. When a new call is established, the calling and called parties' registration contexts are examined to determine if the call should be anti-tromboned. If the 1460 detects that both parties reside on the same private network, anti-tromboning is enabled, causing the media to flow directly between the terminals and freeing up capacity on the Service Provider's access network.

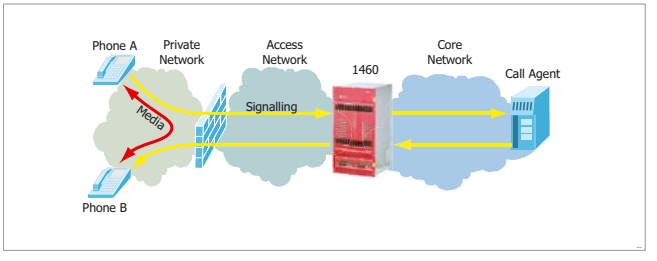

Figure 3 - Anti-Tromboning in Operation

Conclusion

Session border controllers occupy a unique position in the network architecture enabling them to protect the quality of communications. The carrier-class 1460 session border controller, with its advanced hardware accelerated architecture, performs the three aspects of effective bandwidth management; Session Admission Control, policing and anti-tromboning, in a coordinated and effective way with little impact on total performance.

Glossary

LAG	Link Aggregation Group
PABX	Private Automatic Branch Exchange
QoS	Quality of Service
SAC	Session Admission Control
SIP	Session Initiation Protocol
VLAN	Virtual Local Area Network
VPN	Virtual Private Network

91-0066-01-0002-C