Simwood

QoS for VoIP

Over the Internet or WAN

Introduction

QoS, or Quality of Service, is a term used in VoIP circles but often misunderstood and as a consequence rarely implemented, despite the best of intentions. Depending on the network(s) involved in handling your VoIP traffic and in particular the RTP audio element of calls it can make the difference between a better-than-PSTN experience and VoIP being unusable.

In this guide we attempt to clarify some misunderstanding in technically minimal terms and give you a practical guide to ensure you are doing what you can to implement QoS, or at least understanding where your efforts are fruitless and why. We concern ourselves with SIP and RTP here rather than other flavours of VoIP but the principles are fundamentally the same.

We conclude with an Action Plan to give you end-to-end QoS by configuring the elements you control and shining a light on those you can't. If you're serious about quality you will know what to look for in replacing the elements you can't control to optimise the experience for your customers.

As Simwood has developed we have progressively replaced elements higher up our supply chain in pursuit of the levels of service our customers demand. Because we were free from both a legacy telco network investment and an ISP network carrying typical web access/hosting or torrent traffic we were able to engineer a network fit for purpose with no imposed constraints. The result therefore differs to any network which did have that background. We aim to be 'the' network for VoIP business.

What is QoS

If you imagine two ethernet switches connected together, each will have a port connected to the other. Traffic will enter switch A and be forwarded to switch B via the port that connects them. That port has a finite capacity expressed in terms of volume over time, e.g. Gb/s.

When the traffic leaving the port is within that capacity, it will be forwarded by the switch on a first-in-first out basis. If it exceeds the capacity of the port, the switch will attempt to queue the excess packets and then send them on a first-in-first-out basis. When the queue is full, the switch will have no choice but to discard packets causing, as the term suggests, ‘packet loss’.

QoS describes a multitude of technologies to manage that scenario and give the switch some hints as to what is important, in two respects:

1. What should I forward first
2. What can I discard first

If properly configured, those two settings can see a fully congested port behaving as if it is not congested for the traffic that matters.

Our Brocade hardware for example achieves this by having 8 hardware queues on each port, each representing a different priority. Queue 7 is fully emptied before queue 6 and so on. Queues can be configured to be first-in-first-out or employ software algorithms to discard some packets in order to avoid congestion occurring in the first place. If a queue is configured in such a way then each packet is assigned a discard precedence. If not, nothing is preemptively discarded.

Over the years QoS has evolved from a technology used in LANs (i.e. at layer 2) through to one used in WANs and on the open Internet. There are therefore many terms and standards such as CoS, ToS and DSCP (a.k.a DiffServ) all types of which attempt to specify the above in different ways, to suit different applications.

For our purposes we’ll be interested in DSCP, which is the most relevant for wide-area networking and encompasses ToS. CoS is more relevant for LANs.

DSCP is itself very complicated and simplified here. In essence, it allows each packet to be classified with one of 64 different values that devices along the way can use to define one or both of the above values, i.e. map a packet to a queue and set the discard precedence if relevant. This gives network admins potentially 64 different behaviours across the network to suit varying traffic, but of those 64 a few are specifically defined in varying RFCs and given names for simplicity.

A table of the common names and values is below:

DSCP Name	DS Field Value	
	Binary	**Decimal**
CS0	0	0
CS1	1000	8
AF11	1010	10
AF12	1100	12
AF13	1110	14
CS2	10000	16
AF21	10010	18
AF22	10100	20
AF23	10110	22
CS3	11000	24
AF31	11010	26
AF32	11100	28
AF33	11110	30
CS4	100000	32
AF41	100010	34
AF42	100100	36
AF43	100110	38
CS5	101000	40
EF	101110	46
CS6	110000	48
CS7	111000	56

Those named CSx (where x is 0-7 and CS abbreviates ‘Class Selector’) are backwardly compatible with the IP Precedence value. They are defined in RFC2474 if you’d like to read more.

Others are named AFxy (where x represents the Class as above, y indicates the drop precedence where relevant and AF abbreviates ‘Assured Forwarding’) and are defined in RFC2597.

Finally EF abbreviates ‘Expedited Forwarding’ and is specifically defined in RFC 3246.

We have highlighted a few values since they are the most relevant to us here for VoIP purposes and will be discussed now, in descending order of relevance.

EF

EF is recommended for RTP and will typically be mapped to a high priority queue with no preemptive dropping, i.e. the queue will work first-in-first-out and not drop packets in an attempt to avoid congestion. It is commonly recommended that EF traffic does not represent more than 20% of a link's capacity as beyond this level it risks causing the problem it seeks to avoid.

CS3

CS3 is used for SIP signalling, i.e. it isn't as sensitive to delay as EF and has some tolerance for packet-loss as dropped packets will be re-transmitted.

CS0

CS0 is used in two ways. In networks which do not have any specific QoS configuration, packets will be considered to be CS0. In this scenario all traffic will be in a single queue and generally treated as first-in-first-out. This is the default behaviour for most hardware.

CS0 is also used for 'best efforts' traffic such as HTTP and will commonly be configured for queuing with a software queue managing queue size by discarding packets pre-emptively to avoid congestion.

CS1

Where QoS is configured on a network, CS1 can be used for below best-efforts, commonly called Scavenger Class. It is used for bulk transfers to ensure they do not disrupt best-efforts use of the network as well as for managing attack traffic in certain scenarios.

This is one area that can be confusing since it is the only case where a 'higher' DSCP value is used to mark a lower priority and then only if the network is configured as such - if the network admin hasn't included a Scavenger Class then CS1 will be treated as higher priority than CS0.

There are two critical things to understand now:

- **Marking packets is not QoS.** Assigning a DSCP value to certain traffic is fairly trivial to do either in your software or firewall. However, this will in itself achieve precisely nothing and is the most common mistake made by people who 'think' they have QoS. By default your switch or router will probably ignore DSCP and either needs to be told to trust it, trust an alternative marking such as ToS, or to do something else altogether.

- **RTP is UDP not TCP.** TCP behaves quite nicely in that when packets are dropped, it will send subsequent ones at a lower rate. This is useful for general Internet usage and even better where implemented QoS drops the unimportant packets first. RTP is UDP which is fire and forget. In other words even if your QoS policy drops it, it'll keep coming. QoS should therefore be considered as constraining the TCP to make room for the UDP; if your link is full of UDP QoS will not help.

Why it matters

One of the great things about convergence is that by definition one network is carrying multiple types of traffic, rather than multiple circuits being required. Some networks (such as the Simwood network) will be built for carrying VoIP traffic and have over-provisioned links in order to avoid congestion. Others, such as commodity Internet access for example, may routinely run with congested ports for economic reasons. Many are somewhere in between and will see congestion some or all of the time. There are probably few VoIP pioneers who haven't experienced degraded audio when checking e-mail for example. QoS fully implemented could have prevented that given that without the e-mail audio was uninterrupted.

QoS cannot make a lossy connection non-lossy, nor remove jitter, but if packet-loss and jitter arise from modest congestion or packets being delayed respectively, QoS can help.

A common misconception nowadays is that "the cloud" provides consistent performance between any two points connecting to it. That is a gross simplification and simply not representative of reality. Your ISP will have multiple routes in and out of their network, through other providers all of whom similarly have multiple routes. Some may be congested, some may not. Another provider may be better or worse and equally your experience across 'the cloud' to the same end-point may be completely different to someone else using a different ISP. QoS can help here but the real solution is to understand and control the paths your traffic is taking.

There is little to be lost through implementing QoS. Side-effects are usually a consequence of over-implementation involving too many service levels. Our strong recommendation is that most traffic should be best efforts, bad traffic should be Scavenger Class, and SIP/RTP should be marked CS3 and EF respectively. Trying to prioritise ordinary traffic such as HTTP will lead to over-complication and consequent problems, e.g. which is more important a software download over HTTP, video over HTTP or a Skype call which also runs over HTTP ports? Specialist traffic shaping appliances can differentiate between different applications but trying to do so without them is fraught with potential problems. We suggest keeping it simple and prioritising the really important, i.e. voice.

Finally, in VoIP terms QoS is no replacement for capacity and our strong recommendation is that all network links in between both end-points of a call should be over-specified such that QoS is only used to handle micro-bursts of traffic rather than be routinely congested.

Why it might be irrelevant

As mentioned, QoS is usually off by default on most equipment even if packets are successfully marked. A single device without it configured will render it useless if there is congestion on that device. If you use commodity co-location or hosting the network devices will probably not be under your control.

Even if you do control local switches or are an ISP running your own network, traffic will generally leave that network on its journey. If it leaves over IP Transit it will likely be treated as 'best efforts' by the transit network. Quality transit providers will run un-congested networks, others will have congestion as a necessary part of their business model - it is cheap for a reason you know. It is generally exceptional for transit networks to honour your DSCP flags as doing so potentially exposes their network to abuse but some providers offer specialist QoS products.

Whilst your QoS flags may be ignored by transit networks they will probably not be removed. For this reason you have no way of knowing whether QoS is present on the entire path but to look at it more positively, you have nothing to lose through marking packets as they may be applied at some point in the path.

Simwood approach to QoS

We indicated above that EF traffic should not amount to more than 20% of a link's capacity and indeed that follows - how does one prioritise a priority? Whilst 90% of the traffic on the Simwood network is RTP we run the network with no link seeing more than 20% usage. If it does, we upgrade it. This means 80% of each link is unused or available for the minority of our traffic. As a result the network performs superbly for both VoIP and best efforts traffic and QoS is not essential. Contrast this with a commodity network where links would be sized to suit the majority of commodity traffic and QoS becomes very important for a workable VoIP experience.

We ignore DSCP values set on traffic coming into the network, but do not strip them. Instead, we identify the intent of the traffic and mark it accordingly such that RTP is always given priority over signalling with everything else treated as best-efforts. We reserve a Scavenger Class for the lowest priority traffic such as attacks or known bulk transfers (e.g. downloads from our mirror server).

Not all traffic on the network is destined for the network as we provide high quality IP Transit to discerning ISPs. By default traffic is treated as best-efforts (CS0), as is usual for IP Transit, but we also look for three other DSCP values - CS1, CS3 and EF. This enables customers to mark Scavenger Class, SIP signalling and RTP and they will be given appropriate priority over the network. Unmarked traffic will not be identified or marked. We expect transit EF and CS3 traffic to be admission controlled and monitor levels of it for abuse - excessive traffic will result in all being marked CS0. Admission controlled means the marking has taken place on a soft-switch or SBC which has authenticated the user. VoIP traffic originating from a soft-phone should be marked CS0.

This can be summarised as below:

Priority	DSCP Marking	Description
Lowest	CS1	Scavenger Class
	CS0	Best efforts. Default.
	CS3	Transit SIP (admission controlled)
	n/a	SIP to Simwood
	EF	Transit RTP (admission controlled)
Highest	n/a	RTP to Simwood

We always handle the RTP on calls to and from the Simwood network. This enables us to offer transcoding and other features to customers but also means that traffic from/to our customers will follow a consistent path that we can control. Over 70% of our traffic flows to peers ensuring the path is as short as possible (since we connect directly) and links are un-congested. As some traffic will leave our network over IP Transit, we use the best tier 1 networks to try and ensure that whilst handled as best-efforts, it will do so over un-congested paths. Our 20% rule applies on both transit and peering ports.

Knowing what you now know, contrast this with commodity hosting or IP Transit where the majority of traffic is not VoIP, QoS is not applied or is applied against more congested ports. Similarly, contrast us handling all RTP with wholesale VoIP providers who either do not own a network and rely on commodity hosting themselves and/or do not handle the RTP making each call take a potentially different path over non-optimised networks beyond your or their control.

Action plan

In order to make your VoIP reliability the best it can be there're a few steps you can take:

1. How much do you care? It is very important to understand your and your customers' expectations and the value you add. For some it is about saving money with little regard for quality, for others quality is everything and cost doesn't matter. You wouldn't expect the price-driven to be operating their own network whilst you wouldn't expect the quality-driven to be using commodity service providers. We think it is really important therefore to be clear about where on that scale your business sits as whilst some of the steps that follow are free, to be really in the game you will probably need to change service provider or commit to expenditure.

2. Ensure you're appropriately marking SIP and RTP on your equipment. If you're running closed-source solutions you will need to consult your documentation. If using open source solutions such as Asterisk or Freeswitch the easiest solution is to use IP Tables on the server concerned to mark outgoing packets. Instructions to do this can be found at: http://wiki.freeswitch.org/wiki/QoS. For Asterisk you can alternatively set it in the configuration: https://wiki.asterisk.org/wiki/display/AST/IP+Quality+of+Service. You can of course also mark packets on your firewall, router or switch if supported.

3. Your VoIP equipment will be connected to a switch or possibly directly to a router. Work through these to the extent that they're under your control. Setting them to follow your DSCP settings can be configured per-port and is very simple. Be careful doing this and consider your steps. Do you really want to give a third party the ability to mark EF packets coming in to your equipment and potentially depriving your other services of bandwidth?

4. Learn more. Your equipment will have many more options than simply trusting DSCP. You may wish to limit the bandwidth available to certain classes of traffic, or downgrade the class if bandwidth usage exceeds a certain level. By contrast you may wish to additionally reserve specific bandwidth for a given class. If and how you do this will vary by vendor and there's lots to learn.

5. Follow steps 2-4 from your customer's end of the call. Most SIP phones and ATAs enable DSCP marking and some do it by default. If your customers are on ethernet circuits (such as Simwood Carrier Ethernet) you may

have more QoS options but be careful if using xDSL products. It is highly likely that xDSL services will apply traffic shaping which may or may not benefit VoIP. Further, xDSL is built on several more layers each of which may be subject to their own priorities and characteristics. For example, IP over DSL is not on the wire it is on PPP which itself is on L2TP which could be on further layers of L2TP. Irrespective of the priority given at the IP layer, the underlying layers will each have their own priorities and they will not be visible or controllable. Beyond all this the bandwidth available is dynamic! VoIP over xDSL is not recommended but can be made to work most of the time with effort.

6. Having considered the elements of the VoIP path under your control, you now need to assess the paths that aren't. These break into two types:

    1. **VoIP paths.** Is your VoIP wholesaler handling RTP? If so your path to them should be consistent but whether it is QoS enabled will depend on their own network (if any) and the IP Path in between. If they are not handling RTP you will have a different path for each of their suppliers and they will generally be longer; you may wish to monitor the media IP address presented in SDP to assess and will also find out who they're buying from!

    2. **IP Paths.** Your path to your VoIP wholesaler or your customers may be dependent on your colocation or IP Transit provider and may involve several transit providers. Try examining some trace-routes and approach the parties involved to establish how traffic is handled. Is it all best-efforts or do they apply QoS? What is their approach to optimising VoIP traffic and avoiding congestion? Some filtering of answers may be required or you may need to ask more searching questions based on this document - 'our network is QoS enabled' is often heard, usually from people who don't have a network. In a similar vein, you may wish to enquire as to the equipment involved as it will impact how the network performs under load - you may be better traversing a few ISPs routing in hardware than a shorter path across a "network" routing in software.

7. Your answers to 6 should complete your view of the QoS capability of a path but will need repeating for each VoIP supplier and IP end-point if you use multiple suppliers, or in the case of VoIP supply they do not handle RTP. You should end up with your own elements under your control and a view of how well the bit in the middle fares or changes you can make to improve it.

8. You are now in a position to consider the value to be gained through using a specialist such as Simwood in a number of areas:

    1. **Access.** Would your customers benefit from dedicated optical capacity direct to their premises with 100% bandwidth reserved and QoS capable? Simwood Carrier Ethernet may be of interest.

    2. **IP Connectivity.** If you are in colocation, consider co-locating with Simwood or taking connectivity from us. Your traffic to us will be 'on-net', your traffic to other IP end-points will benefit from our network's QoS, extensive peering relationships and quality IP Transit providers. If you are an ISP you may wish to consider Simwood's IP Transit, either full or Partial Transit for our extensive peers alone.

3. **VoIP service.** Simwood can provide a consistent SIP end-point. It will not change according to how we route your call and we control it. It will generally be short as the majority of our traffic comes through peers, and we peer with other willing networks who we exchange traffic with. Similarly, when traffic leaves our network to our own suppliers we manage the path and often peer with them. Of course, if it leaves us (or comes in to us) by SS7 then we will be the end-of-the-line as far as VoIP is concerned.

Conclusion

QoS can help improve the stability of your VoIP service but it cannot make the bad good. Configuring it is more than simply marking packets and equipment from end-to-end needs to be configured or at least enabled to take account of your preferences. If that equipment is not under your control, or the path varies due to your supplier preferences, then your options are limited. Good suppliers will be applying QoS and giving stable routes between end-points but QoS will deliver little benefit if weak-links remain in the chain. Consider the action plan and review your suppliers to ensure that as much of your call's paths are QoS enabled, either under your own control or using a trusted supplier such as Simwood. As always, if you need any help or advice just get in touch.