Managing and Troubleshooting VoIP

Sung Moon

Director of Technology

Avaya Global Services

INTEROP®

THE LEADING BUSINESS TECHNOLOGY EVENT

Topics

- Design and Implementation
- Monitoring
- Troubleshooting

Decide With Your Service Provider…

- What goes in?
- How does it get treated?
- What comes out?

What Goes In?

What Goes In? VoIP Audio

- How will this traffic be classified?
- Recommend marking at L3 (IP Header) using DSCP for end to end integrity of marking.
  - L2 tagging is not preferred because it must be converted to L3 and back to L2 at various points.
- Expedited Forwarding PHB (RFC 3246) is the industry standard.
- Simply put, mark your audio with DSCP 46, and coordinate w/ your SP.

What is Expedited Forwarding PHB?

- Expedited Forwarding PHB (RFC 3246) describes a forwarding mechanism that is essentially a priority queue.
  - "EF is intended to provide a building block for low delay, low jitter and low loss services…"
- Various equipment vendors and service providers might have different terms for this type of queuing, but it is likely the best queuing mechanism or service they offer.
- At the enterprise edge, assign audio traffic to this queue; typically called the priority queue (PQ).

What Goes In? VoIP Signaling

- How will this traffic be classified?
- Again, mark at L3.
- Use one of the Assured Forwarding PHB (RFC 2597) values.
  - DSCP 26 (AF31) is a common one for signaling.
  - AF31 means class 3, drop precedence 1 (low).
  - Cisco now uses DSCP 24, aka CS3 or IP Precedence 3.
- Simply put, mark your signaling with DSCP 24 or 26, and coordinate w/ your SP.

What is Assured Forwarding PHB?

- Assured Forwarding PHB (RFC 2597) describes a forwarding mechanism that calls for…
  - 4 classes of traffic, or effectively 4 queues.
  - 3 drop precedences (low, med, hi) per class, or per queue.
  - AF31 is class 3, drop precedence 1 (low) = DSCP 26.
- At the enterprise edge, put signaling in its own class-based weighted fair queue (CBWFQ) and engineer the bandwidth so that none of the traffic in this queue is dropped.

Assured Forwarding	Class 1	Class 2	Class 3	Class 4
Low Drop Prec (1)	DSCP 10	DSCP 18	DSCP 26	DSCP 34
Med Drop Prec (2)	DSCP 12	DSCP 20	DSCP 28	DSCP 36
High Drop Prec (3)	DSCP 14	DSCP 22	DSCP 30	DSCP 38

What is CS3?

Type of Service (ToS) Field in the IP Header							
IP Precedence or Class Selector			Type of Service				Rsrvd
0	1	1	0	0	0	0	0
Differentiated Services Code Point (DSCP)						Unused	
0	1	1	0	0	0		

- CS3 is Class Selector 3.
- same as IP Precedence 3.
- same as DSCP 24.
- What does this mean? Depends, but irrespective of the marking…
- At the enterprise edge put signaling in its own class-based weighted fair queue (CBWFQ) and engineer the bandwidth so that none of the traffic in this queue is dropped.

Do You Have To Use These DSCP Values?

- Not really. These DSCP values were identified because they conform to IETF specifications and industry norms.
- Router and switch vendors often assign specific DSCP values to specific queues, and they're not all uniform in implementation.
- Which DSCP values you decide to use should depend on two factors.
  - What's easiest for you, depending on your networking equipment.
  - What your SP can make work, depending on their networking equipment.
  - You and your SP must be in synch!

Do You Have To Use Two Queues?

- Not if the bandwidth is 1Mbps or greater.
  - For large bandwidth links the serialization delay is minimal, so a large signaling packet will not slow down a small audio packet, which is a primary reason for using two queues.
- The amount of signaling relative to audio is negligible, so putting both audio and signaling in the same priority queue is feasible.
- The simplicity in configuration and potential efficiency in processor utilization can offset the lack of boundary for troubleshooting.

How Does It Get Treated?

Match the Service with the Queuing at Enterprise Edge

- We've stated what queuing mechanisms should be used at the enterprise edge, where you hand off your traffic to the SP.
  - Priority queuing for audio.
  - Class-based weighted fair queuing for signaling.
- The SP cloud, especially MPLS, can be complex.
- Depending on the SP, they may have services and SLAs that parallel these queuing mechanism.
- SP must properly classify audio, per the agreed upon DSCP, and give it priority service.
- SP must properly classify signaling, per the agreed upon DSCP, and give it dedicated service with very low loss.
- In all cases, specify SLA parameters such as delay, jitter, and loss.

Match Allotted Bandwidth with Call Admission Control

- BANDWIDTH – the scarce resource; the petroleum of networking.
- You and the SP must agree on the bandwidth allotted for the PQ service and CBWFQ service.
- You must apply Call Admission Control (CAC) so as not to exceed the provisioned bandwidth.
  - Administer the VoIP system, and allot the bandwidth, to allow N calls.
  - The N+1st call must be blocked or routed over other facilities.
  - If this is not done, the N+1st call will potentially disrupt all N+1 calls, not just the +1.

What Comes Out?

What Comes Out?

- Should be exactly what went in!
  - DSCP marking is the same.
  - Packet order is the same.
  - Jitter is relatively unchanged.
  - Loss is zero or minimal.

- Let’s hear experiences from the audience.

Bringing It All Together

Topics

- Design and Implementation
- Monitoring
- Troubleshooting

Monitoring the SP Network

- Wouldn't it be nice if there were a tool that told you..
  - If what you got out is NOT what you sent in?
  - If the delay, jitter, and loss thresholds were breached?
  - If the allotted bandwidth was exceeded?
- Avaya Converged Network Analyzer is in this space.
- Fluke Visual UpTime is in this space.
- Anybody in the audience know of others?

Monitoring the Enterprise Network

- Was the DSCP applied at the user endpoint preserved across the enterprise to the edge?
- Based on anecdotal evidence within Avaya Global Services, unintentional re-marking in the enterprise is more common than it should be.
  - Results in poor audio quality and/or poor system performance.

Topics

- Design and Implementation
- Monitoring
- Troubleshooting

End to End Troubleshooting

- Most effective (though not most efficient) troubleshooting is to capture at both ends.
- Compare extracted audio at source to audio at destination.
- Compare RTP stream at source to stream at destination.
- Merge source and destination captures to see end to end view.
- ClearSight Analyzer is a good tool in this space.
- Anyone know of other tools?

End to End Troubleshooting: Questions Answered

- Was the audio poor when it exited the source? Problem in media gateway or trunk.
- Or was it poor when it reached the destination? Problem in network.
- Was the RTP stream dirty at the source? Possible duplex mismatch.
- Or was it dirty when it reached the destination? Problem in network.
- Were packets dropped along the path? Signaling disruption delays call setup or audio setup.

General Troubleshooting: Find Patterns and Commonalities

- Problem occurs at same time day after day.
  - SP network congestion; high user traffic period.
- Problem affects only a specific device or specific types of devices.
- Problem affects only a specific group of users.
  - Common application.
  - Common IP subnet or WAN link (Philippines case).
- Problem is always audio related.
  - Priority queue problem.
  - RTP header compression problem (hanger case).
  - CAC not working properly.
- Call features (connect, hold, transfer, etc) are delayed.
  - Problem with signaling queue.

General Troubleshooting: Trending

- Very tedious but sometimes necessary.
- Having a trouble reporting template is very useful.
  - Problem description, day & time, incoming or outgoing, which side heard the problem, etc.
- Must be meticulous about reporting troubles, using the template.
- Takes time – days or weeks or longer.
- After some time, a trend emerges.
- Mitchell, SD case.