Written by Paul Stalvig | Technical Marketing Manager

RTSP: One of the Foundations of Convergence

Overview

The growth of the Internet has provided the ability to share, communicate, and preserve more communications than ever. Real-Time Streaming Protocol (RTSP) was developed in the late 1990s, and is a foundation for the current convergence of communications. RTSP provides applications the ability to stream voice and video using standard VCR-like controls such as "play" or "pause." A quick look at QuickTime, Windows Media Player, RealPlayer, MPEG4IP, and Skype show how widespread RTSP has become.

As with all IP-based telecommunications, RTSP is susceptible to jitter, packet loss, feedback, and timeliness. RTSP typically uses Real-Time Transport Protocol (RTP) as the actual data stream, and this is often referred to as the data-channel of RTSP. RealNetworks' RealPlayer uses a proprietary Real Data Transport protocol. RTP has a sister protocol; the Real-Time Transport Control Protocol (RTCP) that collects and communicates the Quality of Service (QoS) information, such as bytes sent, packets sent, lost packets, jitter, feedback, and round-trip delay, for the data being transported by the RTP. Applications can use the information from the RTCP to improve quality by slowing down the data flow, using smaller packet sizes, or using a lower compression of the data.

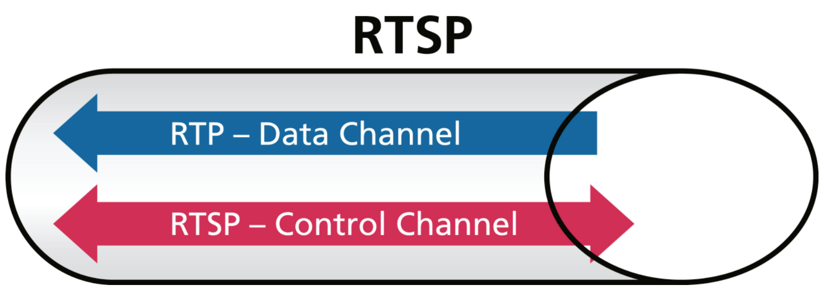

Figure 1: *RTSP with Data and Control Channels*

It is sometimes easier to think of RTSP in terms of a television with a VCR. The RTSP acts as the remote control to the VCR and the RTP is the television sending the display information you watch. The only part missing is RTCP, which would allow you and your television to communicate about the quality of the television's display information.

Many organizations have built RTSP proxy servers to offload the overhead from control channel communications due to the large amount of information that implementations cause on a stand-alone streaming server. This enables the streaming servers to provide more streams of information and the RTSP proxy servers to perform more auditing, billing, and reporting functions. The RTSP proxy servers are also in a prime location, providing a point in the communications where another application can analyze the control channel information, which improves the overall QoS of the streamed telecommunications.

RTSP proxy servers are not the same as BIG-IP® Local Traffic Manager™ (LTM) RTSP proxy service. The RTSP proxy servers focus on only the data channel communications of the IP-based real-time streaming telecommunications. LTM RTSP proxy service is a function of how BIG-IP LTM handles the RTSP communications within BIG-IP LTM.

RTSP launches IP-based communications by establishing a connection on port 554 to the RTSP server. The real-time control protocol then sends the command to the server to begin playing. The server chooses a free port, typically between 16384 and 32767, and begins to send the information directly to the client, not necessarily using the same route on which the information came to the server.

It is interesting to note that RTSP is not confined to either TCP or UDP, but can use one or the other, or even both. When using TCP, the packets are sent through a connection and delivery is somewhat guaranteed. The nature of TCP makes it prone to causing delays and jitter in real-time communications. So, why not use UDP? Using UDP has its own set of issues. For example, there is no guarantee of packet delivery or smaller packet size. Also, routing of UDP is sometimes difficult when using older networks. Over time, and through trial and error, UDP has become a quasi-standard for streaming communications.

Challenge

When implementing RTSP, the multiple port and protocol nature of RTSP can make streaming audio/video difficult to manage. Many have tried tunneling the audio/video through HTTP protocol, but that compounds the affects of jitter, timeliness, and protocol overhead and adds to the complexity of the solution.

Others have used direct server return (N-path routing, see figure 2) to enable a server direct routing access to the clients. This method can create issues when checking the health of the servers. It can also force persistence based on the source IP address, which will cause all requests from a carrier or Internet Service Provider (ISP) to be routed to a single server.

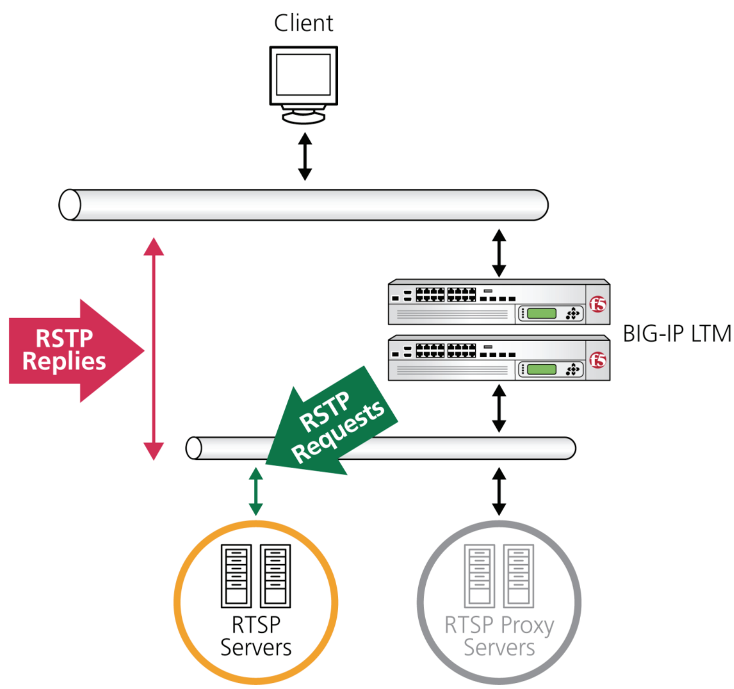

Figure 2: *N-Path Routing of RTSP*

RTSP proxies and servers can deal with numerous connections and un-tunable connection parameters. However, this creates significant overhead, including costs for servers and bandwidth used.

Solution

In a standard scenario, a session for the control port (usually 554) from the client is load balanced to the RTSP proxy server. BIG-IP LTM monitors the session as the RTSP proxy server negotiates the port and protocol of the actual stream(s) on behalf of the RTSP server. The stream is then initiated outbound from the RTSP server directly to the client via BIG-IP LTM, without using N-path routing. BIG-IP LTM can also appropriately NAT/SNAT the new connection between the client and RTSP server. Additionally, BIG-IP LTM is able to perform health checks on the RTSP servers and the individual streams of information.

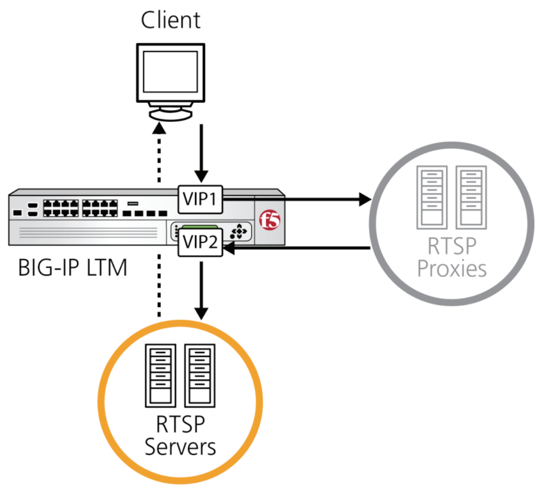

Figure 3: *New Routing of RTSP with BIG-IP LTM*

By having all the RTSP data streams available, BIG-IP LTM can effectively monitor the streams of data and load that each server is providing. By analyzing this information, BIG-IP LTM can make efficient load balancing and availability decisions. The RTSP proxy also allows HTTP tunneled RealPlayer communications (RDT) to be set back to the standard RTSP ports.

By using a BIG-IP LTM RTSP proxy, LTM is fully aware of the RTSP communications. The RTSP profile in LTM can provide additional configurations which enable an administrator to fine-tune how the RSTP connections are handled (Figure 4). The RTSP profile helps reduce unnecessary overhead on RTSP servers by ensuring that connections are not open too long or held open by the amount of unacknowledged data that can exist before a connection is considered closed.

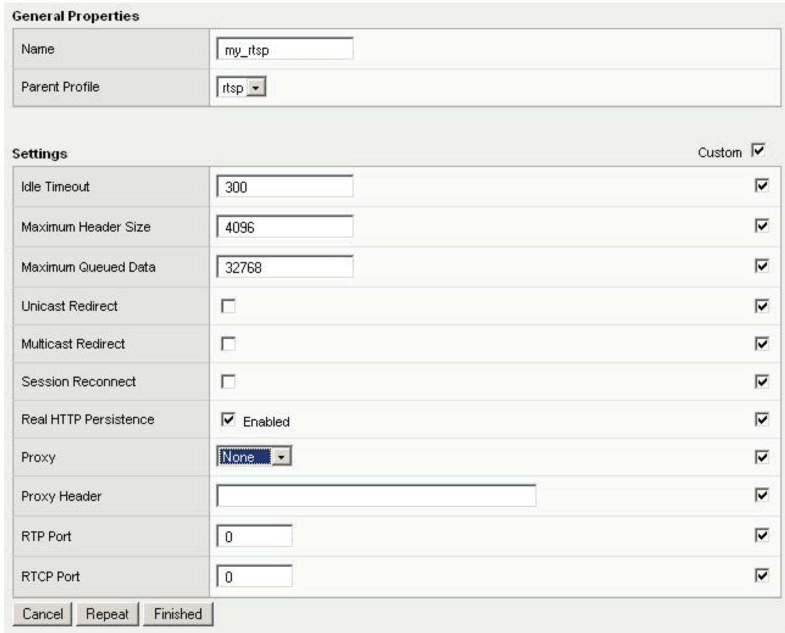

Figure 4: *Configuration of RTSP Profile*

With the capabilities of F5, using IP-based telecommunications is no longer a difficult process. RTSP servers and proxy servers can become part of the standard web site infrastructure. Operational expenses for configuration, support, and maintenance can be reduced. Additionally, this technology can introduce the concept of Web 2.0 into an organization with minimal expenditures.