



Building the Next Generation Network Bridge: From Today's Network to the Future

Introduction

While the “buzz” on the street seems to be all about the next generation network (NGN) and the IP Multimedia Subsystem (IMS), it is perhaps more important to talk about how today's existing service provider (SP) networks can start making the slow migration toward the NGN. Very few SPs will simply be able to build a whole new network, and even those that can still need to support legacy devices and services for the foreseeable future. To make matters worse, many of the standards for IMS are still evolving or are implemented differently by different vendors.

The modern SP needs tools and technology to help them start deploying NGN services today on their existing networks with a mindful-eye toward tomorrow. F5 Networks provides these tools and technologies and is already a major contributor to many existing SP networks providing cutting-edge technology.

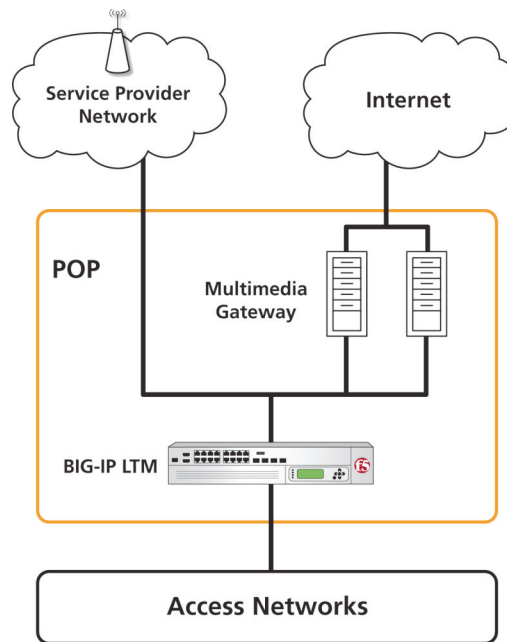
Teaching Old Networks New Tricks

Customer Case: A traditional mobile carrier in North America was looking to expand their services as part of their 3G network to deploy Internet access to all their subscribers. Once IP traffic was received by the local Point-of-Presence (POP), traffic not destined for inter-network service needed to be routed to the Multimedia Gateway (MMG) for interface to the Internet. The MMG provided for optimized access to the Internet for various classes of traffic (web, video and so on).

Customer Challenge: The customer faced several challenges. First, not all IP traffic was destined for the Internet, some was intra-network traffic that the customer didn't want going to the MMGs. Second, they wanted a way to have subscriber traffic dynamically go to the regional MMG instead of being hard-coded to their home MMG or having to back-haul all Internet-bound traffic to their data center. Third, particularly since the service was new and being offered in response to competitive offerings, they needed the MMG to be as fast and reliable as possible.

F5 Solution: The same F5 BIG-IP® Local Traffic Manager™ (LTM) that has provided high availability (HA) and scalability for the enterprise for more than a decade can provide the same benefits to the modern SP who is migrating to or otherwise implementing IP-based services. By placing BIG-IP LTM devices in each POP directly in the default data path, BIG-IP LTM can add intelligence to traffic routing; provide unique network translation services; and ensure consistent, reliable access to application services like the MMG.

The BIG-IP LTM device uses the concept of a virtual server to represent different services for which it provides traffic management. In the enterprise, this might be a virtual server to handle inbound web requests which it then simply load balances across a farm of web servers. In this case, because BIG-IP LTM can handle any IP-based traffic and differentiate them based on IP, port, or even elements in the payload, it could easily differentiate intra-network traffic from Internet traffic, intelligently routing subscriber requests either to the SP network or the MMG for Internet delivery. In addition, as the SP continues to grow their service offerings, BIG-IP LTM can intelligently differentiate individual services and route them appropriately, for example streaming video to one set of dedicated MMGs and web traffic to a set of basic web proxies or content delivery network (CDN).



Another unique aspect of BIG-IP LTM is that it can capture and manage traffic that isn't even specifically intended for it. Originally, the SP proposed to hard-code the MMG IP addresses into the subscriber handsets; unfortunately, each POP had a different IP meaning, requiring handsets in each region to be specifically configured. Not only was this a provisioning burden, it also could not account for users who roam or move from one region to another. Unless their devices were reconfigured (which is difficult to do with roaming users) all their Internet-bound traffic would be back-hauled to their home POP before it would make it to the Internet. Since BIG-IP LTM was already in the default data path, this provided a much more elegant solution. The SP simply configured all handsets to have the same—non-existent—IP address listed for the MMG (in this case 10.0.0.1 and 10.0.0.2). Then, each BIG-IP LTM device was configured to look for traffic destined to those addresses and redirect them to their local MMG. This effectively provided dynamic reconfiguration of the handsets to always use the local MMG.

Finally, in the face of stiff competition, the SP needed to ensure that users always reached a working and well performing MMG; slow, inconsistent or non-working services would not garner much uptake. Here, BIG-IP LTM provided a classic enterprise solution for the SP. Multiple MMGs were installed in each POP and subsequently load balanced by BIG-IP LTM. This provided increased capability to scale as subscriber use amplified (simply adding another MMG as needed) as well as the fault tolerance needed to ensure that the subscriber requests always reached a working and well performing MMG. In addition, BIG-IP LTM was configured with health monitors to consistently maintain awareness of the availability and performance characteristics of each MMG, which was also right out of the enterprise handbook. This even included the capability to monitor the local POP Internet connection and re-route requests to another POP in the event that none of the local MMGs were functioning or capable of reaching the Internet.

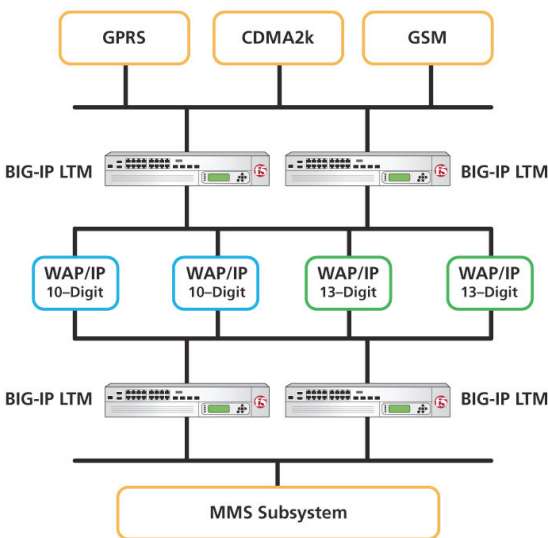


Bridging the Divide

Customer Case: A large SP was planning migrations from their initial service deployments to newer versions that were better positioned for the IMS infrastructure. As with many SPs, they needed a way to integrate the new systems with the old to provide seamless migration and upgrade, all without any impact to the subscriber base. Their initial target service was an upgrade of their original Multimedia Messaging Service (MMS).

Customer Challenge: Although the base back-end multimedia controllers were provided by the same vendor, changes in the software required the use of a 13-digit ID instead of the 10-digit ID previously used. This meant that in order to transition all subscribers to the new system, all devices would need to be reconfigured (or replaced) simultaneously to implement the new system. The other option was to deploy an entirely new architecture and manually migrate users on an ad hoc basis. Neither solution was optimal.

F5 Solution: The SP had already deployed F5 BIG-IP LTM devices within the legacy system to load balance the WAP/IP gateways which provided a connection between the subscribers and the MMS back-end infrastructure; furthermore, they planned the exact same solution for the new WAP/IP gateways for the upgraded system. This provided the perfect point of integration between the two systems.



BIG-IP LTM, like all of the F5 BIG-IP family of products based on the TMOS™ software platform, has unique programmatic capabilities built into it: iControl and iRules. Using iRules on BIG-IP LTM devices in front of the WAP/IP gateways (the side closest to the subscriber), BIG-IP LTM could differentiate those devices using 10-digit IDs from those using the new 13-digit IDs and dynamically deliver traffic to the appropriate set of WAP/IP gateways without reconfiguring the handsets. As new, upgraded handsets (using the 13-digit ID) became available to the subscribers, they automatically migrated to the new system. Legacy users still received the same service they always had.

Another set of BIG-IP LTM devices behind the WAP/IP gateways (the side closest to the MMS subsystem) provided additional critical services. They provided the load-balancing and high-availability checking for the MMS subsystem servers—but more importantly, they provided last-

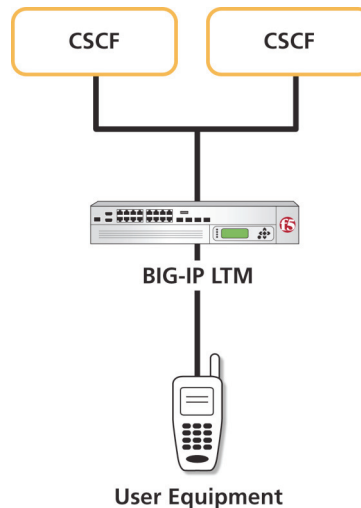


hop capability. Last-hop guarantees that traffic back to the handsets is returned to the WAP/IP gateway that initiated the request, which ensures the end-to-end transaction.

When Standards Aren't Standard

Customer Case: A mobile carrier was initiating the first IMS services deployment within their network. While IMS (and its underlying protocols of SIP, RTSP, and SCTP) are all standards-based, the implementations of IMS are so nascent that not all providers implement them in the same way. As new implementations are rolled out, this uncharted territory can cause significant challenges. This particular customer was trying to deploy Caller ID Blocking (or Per Call Restrict, otherwise known as *67) services within their IMS infrastructure.

Customer Challenge: As they started their first pilot roll-out, they discovered an issue with their Call Session Control Function (CSCF) equipment that caused the CSCF to crash and reboot after running for some time without any difficulties. Further examination found that it was directly related to the way the CSCF handled SIP messages in relation with the *67 service as implemented by the SP. The simple method of providing the *67 service was to simply delete the CALL-ID Header within the Invite SIP message when the FROM header was set to “anonymous”—completely blocking the recipient from seeing the information. The problem occurred when the call was completed and the SIP CANCEL message was sent. Because the message did not have the CALL-ID information, the CSCF was unable to determine which session to terminate and therefore didn't terminate any sessions. As time progressed, the CSCF became resource constrained (filling up with stale connections never terminated) and failed.



Since the IMS and SIP standards aren't completely ratified or exhaustive, it was difficult to determine who was at fault for the failed interaction between the devices. The SP felt that the issue was with the CSCF, but without any firm basis to make the claim, they were at the vendor's mercy. The CSCF vendor claimed that the issue was with the SP implementation, not their equipment; however, they believed they could solve the problem within nine months. Both parties claimed to have followed the necessary standards.

F5 Solution: For this customer, delaying the deployment for nine months (and an estimated \$1 million in lost revenue and expenses) was simply not a viable solution; fortunately, they already had BIG-IP LTM as part of their solution set, providing high availability for the call proxy servers. In this case, iRules saved the day.



Because the TMOS software platform is based on full-proxy architecture, BIG-IP LTM is capable of inspecting and interacting with all content that it manages and delivers, in both directions, no less. iRules gave the customer the ability to intelligently manage their issue. BIG-IP LTM simply inspected the inbound INVITE request, made note of the CALL-ID information, and kept its own tracking record of the session. When an anonymous CANCEL message came in, BIG-IP LTM re-inserted the appropriate information to make it possible for the CSCF to correctly identify the session to cancel. The entire process took less than four hours to solve and implement.

Conclusion

Very few, if any, SPs are going to be capable of simply activating a complete IMS infrastructure and all of them are going to need to service legacy services simultaneously with any new services. Fortunately, most of the challenges facing SPs as they move forward are not all that different from the challenges enterprise organizations have faced for many years when dealing with IP-based application delivery. With the increased capability of F5 Networks to handle the ubiquitous protocols of the IMS world (SIP, RTSP, SCTP) in addition to the many other TCP/IP protocols used in IP networks, the BIG-IP family of Application Delivery Networking equipment is perfectly poised to help the service provider of today—and tomorrow.