GUIDE

F5 and the Growing Role of GTP

Key use cases including traffic shaping, network slicing, iot, and security

F5 and GTP—for years, F5 BIG-IP solutions have managed GPRS Tunneling Protocol (GTP) traffic. Now, in part due to EU regulations, GTP traffic growth is much higher than ever before. International trends show that mobile data is increasing, as is the demand for smart and secure GTP traffic handling. In this document we will explain GTP and some market-proven use cases where BIG-IP solutions provide significant value.

GTP carries General Packet Radio Service (GPRS) within GSM, UMTS, and LTE networks through a group of IP-based communications protocols. GTP and Proxy Mobile IPv6-based interfaces are specified on various interface points within 3GPP architectures.

GTP was originally used in GPRS (2.5G networks), later developing a similar role in 3G and 4G networks. For 4G, the key nodes have different names and, to a certain extent, are comparable to nodes used in 3G networks. GTP is staying with us, and it's vital that service providers grasp how to optimize its capabilities, both now and in the era of 5G.

While the 5G architecture was defined in Release 15 of the 3GPP specifications, Release 14 defined the separation of Control and User Plane Separation (CUPS) of EPC nodes. CUPS provides the architecture enhancements for the separation of functionality in the Evolved Packet Core's SGW, PGW, and TDF. This enables flexible network deployment and operation by distributed or centralized deployment and the independent scaling between control plane and user plane functions, without affecting the functionality of the existing nodes subject to this split.

Key Use Cases for GTP

USE CASE 1

Load Balance To The Best Packet Data Network Gateway

BIG-IP solutions help the network select the best available Packet Data Network Gateway (PGW), and guarantee that if there is usable capacity available, traffic will route to it. Without this smart load balancing function, traffic is sent to a SGW through a static table provisioned in DNS, which can't manage the dynamic behavior of the network or its traffic. The result is, at best, an inefficient use of scarce and expensive network resources, and at worst, reduced service capability and quality.

GTP Traffic Management: load balance towards PGWs

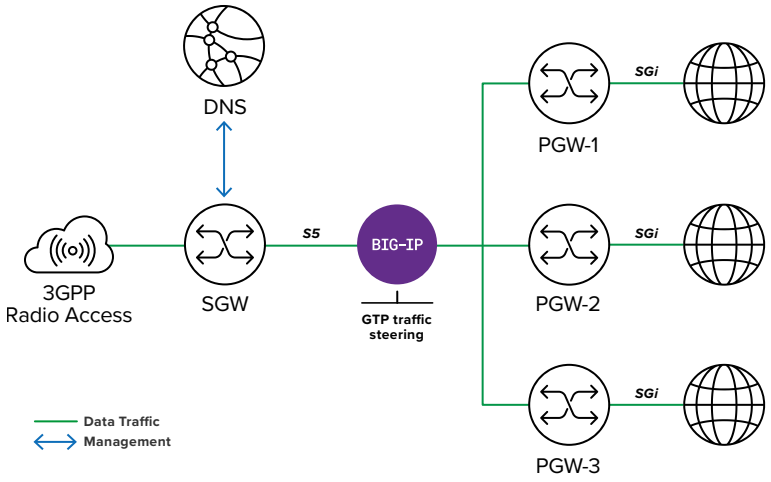

Problem Statement

A mobile operator wants to scale the PGW capacity in the most effective way. The PGW address gets selected based on a DNS query, but if the PGW has limited availability the traffic could be impacted.

Solution

An Access Point Name (APN)-based DNS resolution points all APNs to a BIG-IP solution with a load balancing algorithm that selects a PGW. A health monitor constantly verifies the health/performance of the PGWs so the BIG-IP solution can steer incoming signaling data (GTP-C) messages to the best one available. If a PGW is degraded or unavailable, it will automatically be skipped until the monitor marks it "healthy."

By provisioning the DNS with the address(es) of only BIG-IP devices, all S5 GTP traffic targeted at PGWs is routed to them (noted in the illustrations as "GTP Traffic Steering"). The health monitor determines whether individual ones should be used or skipped based on availability. Traffic can be rerouted in the event of unforeseen errors or during regularly scheduled maintenance or modifications.

USE CASE 2
Route PER MVNO

BIG-IP solutions help mobile operators route traffic for customers on a particular Mobile Virtual Network Operator (MVNO) service to specified PGWs.

GTP Traffic Management: route on APN and IMSI

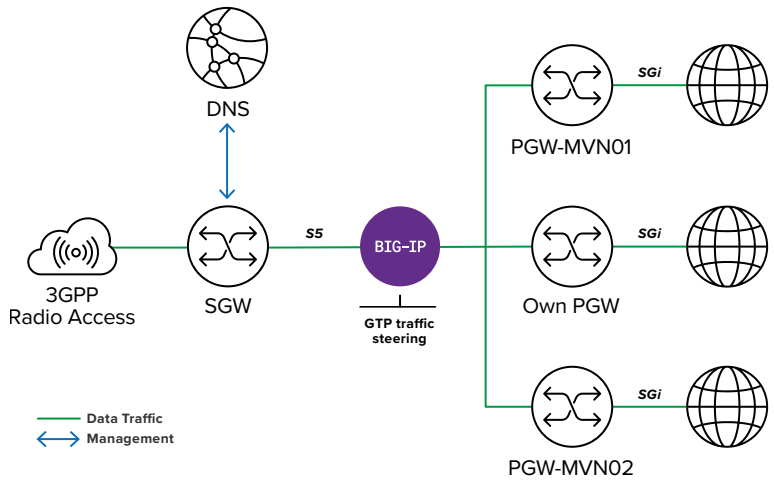

Problem Statement

Network operators want to use dedicated PGWs for an MVNO, but the same APN ID is used for both the home network and MVNO traffic. Therefore, a DNS-based APN resolution procedure to find the proper PGW will not work.

Solution

An APN-based DNS resolution points all traffic to a BIG-IP solution provisioned with a table mapping International Mobile Subscriber Identification (IMSI) ranges to the corresponding MVNO PGW. Incoming GTP-C messages are steered to the right PGW based on IMSI ranges by inspecting GTP-C IE attributes (see 3GPP 29.274).

BIG-IP solutions are provisioned with the IMSI range for each MVNO. Subscribers using the services of the hosting mobile network have an APN that looks like the mobile operator's own subscribers. But the MVNO should have the traffic of its customers managed by its own PGW, where additional services can be applied and charged for under full control of the MVNO. For instance, one MVNO would go through PGW-MVNO1 while another would be routed via PGW-MVNO2.

When a BIG-IP device receives GTP traffic, it evaluates the GTP-C information and determines both the selected APN and the IMSI of the subscriber. The pre-provisioned table shows the IMSI-ranges belonging to each specific MVNO, as well as who the MVNO is or its PGW address(es), then routes the traffic accordingly. If there is no match for the IMSI to an MVNO, the traffic defaults to the hosting mobile operator's PGW(s).

It's worth mentioning that an MVNO can add an additional IMSI range simply and without affecting service. The hosting network can also add a table for an additional MVNO at any time, without impacting service for the others.

USE CASE 3

Dial-IMSI Routing

When MVNO customers have dial-IMSIs, they use the IMSI of their MVNO while in-country and the IMSI of a hosting operator while roaming. For providers to benefit from their commercial roaming relationships, a routing function determines whether traffic is from a roaming MVNO subscriber or a carrier's own customer, then sends it to the appropriate PGW.

GTP Traffic Management: Dial-IMSI routing

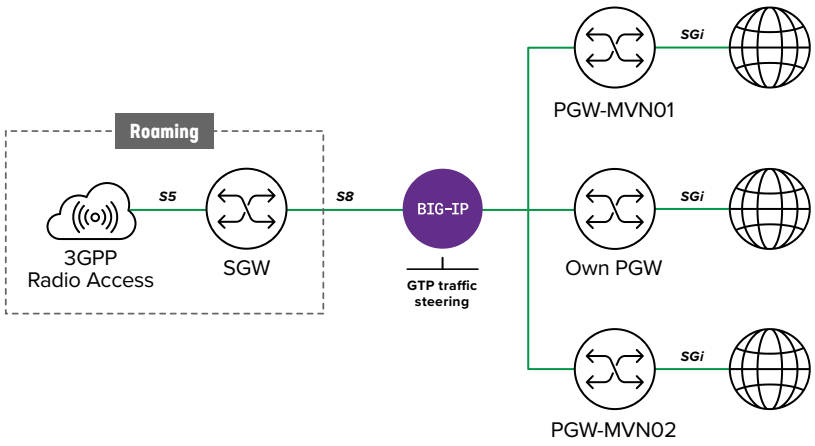

Problem Statement

MVNOs have agreements in place to use the IMSI ranges of host operators for their roaming subscribers. The host operator wants to send MVNO roaming traffic to the MVNO PGW and their own subscriber traffic to dedicated PGWs. Their DNS-based resolution points to the BIG-IP GTP-Proxy for all IMSIs.

Solution

BIG-IP solutions are provisioned with table mapping to route traffic with specific IMSI ranges to their corresponding MVNO PGWs.

When customers access the Internet and other data services while roaming, traffic is routed from the network via the S8 interface to the BIG-IP solution. Typically, this is done via the services of one or more IPX carriers. The IMSI ranges of each MVNO are listed, via pre-provisioning. When an MVNO customer with dual-IMSI service is roaming, the customer's device selects the IMSI belonging to the host network. The network verifies the IMSI and, based on the Operator Code in the IMSI, it determines where the home network is located and offers the traffic to its IPX carrier. The GTP traffic arrives at the home network, for example, a BIG-IP device fronting the home network's PGWs. If the IMSI matches an MVNO IMSI-range in the device routing table, the GTP traffic goes to the PGW(s) of that host. If there is no match, the traffic is load balanced to the host's own PGWs.

USE CASE 4

Modify/Correct APN

In some cases, the MVNO requests specific APNs, to which it has assigned special services or conditions. Once a routing decision is made, the APN shifts to a pre-configured value, then sends the GTP traffic to the MVNO PGW(s).

GTP Proxy: correct APN and route on IMSI range

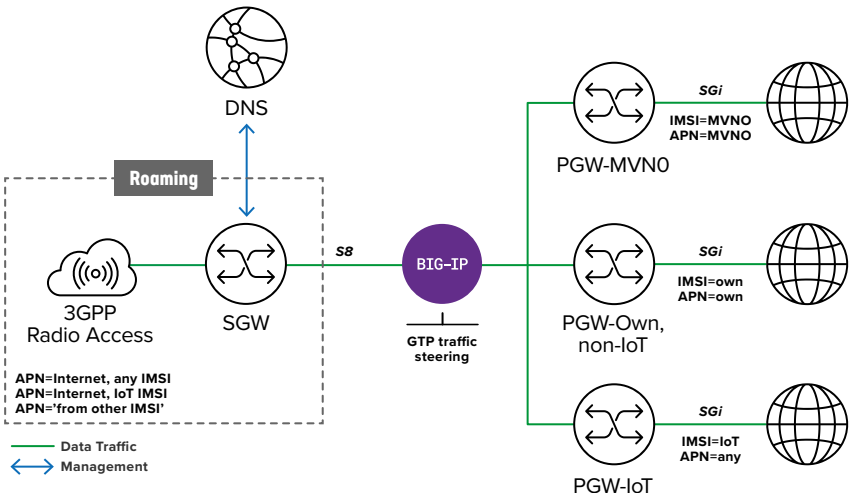

Problem Statement

When an MVNO customer uses the wrong APN due to dual-IMSI behavior, the result is a lack of services. Likewise, IMSIs belonging to IoT devices should be routed to specific PGWs (PGW-IoT).

Solution

BIG-IP solutions are provisioned with IMSI-MVNO and IMSI-IoT lists. If the IMSI matches the IMSI-MVNO list, the APN changes to “MVNO” and routes to PGW-MVNO. However, if the IMSI matches the IMSI-IoT list, it routes to IoT without changing the APN. When the IMSI doesn’t match either list, it routes to default PGWs without changing APNs.

When modifying and correcting APNs, the routing considerations and solutions are the same as with dual-IMSI routing. The difference is that the BIG-IP device also checks the submitted APN and whether it came from an MVNO customer or specific IoT IMSI ranges. If so, a special rewrite for the APN value is necessary.

BIG-IP devices check IMSI ranges against a pre-provisioned list of APNs. If the range matches, an new APN is listed and either inserted in the original GTP-C message or it replaces the original APN. It’s then routed to a MVNO PGW or a PGW handling IoT traffic.

USE CASE 5

Route on GTP IEs at S5 Level (GTP Packet Broker)

Maximize your investment in third-party solutions by improving management of GTP traffic. For example, rather than routing all GTP traffic through VoLTE inspection, only send the traffic for which the APN is different than the IMSI. Because we're talking about user data (the GTP-U), the volumes require sizable capacity extensions.

And as these networks move to NFV environments, the third-party solutions need to be available as NFV-based solutions as well.

GTP packet broker: route on GTP IEs at S5 level

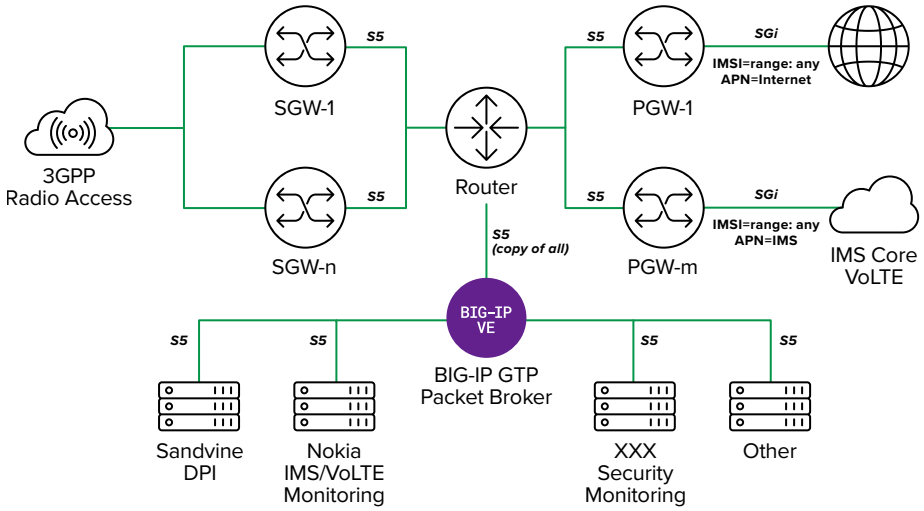

Problem Statement

A customer wants to save on DPI and monitoring resources by processing only required GTP traffic. Because monitoring on S5 provides more detail, the solution needs to be virtualized.

Solution

The BIG-IP device's GTP Director serves as a packet broker for GTP-C and GTP-U traffic, routing a copy of ALL GTP traffic to the BIG-IP device. It then steers all traffic matching certain criteria, like IMSI=X, or APN=Y, or GTP IE=Z.

The router sends a copy of ALL GTP messages to the BIG-IP device. BIG-IP Virtual Editions (VEs) have been provisioned with selection criteria for each third-party destination. If the criteria match what needs to be investigated, that part of the traffic will be forwarded to those nodes. Because the traffic is a copy, unmatched portions can be dropped, and the information will continue to flow between SGW and PGW.

USE CASE 6

IoT Traffic Routing

Connectivity demands will continue to grow exponentially as the number of connected devices in homes and businesses of all types continues to explode. Service providers can take steps now to effectively optimize and secure their network for incredible growth in IoT traffic.

When an operator offers IMS services, which is typically the APN=IMS agreed upon between GSMA operators, it allows roaming and offers the best possible QoS (Quality of Service). When a limited number of services, like VoLTE, are served by the same IMS network this all works fine; however, when some of those services are managed by a completely different IMS network, possibly with a different technology (e.g., virtualized vs. bare metal), you need a smarter solution to serve as the routing differentiator.

GTP Traffic direction for IoT: route on APN=IMS and IMSI

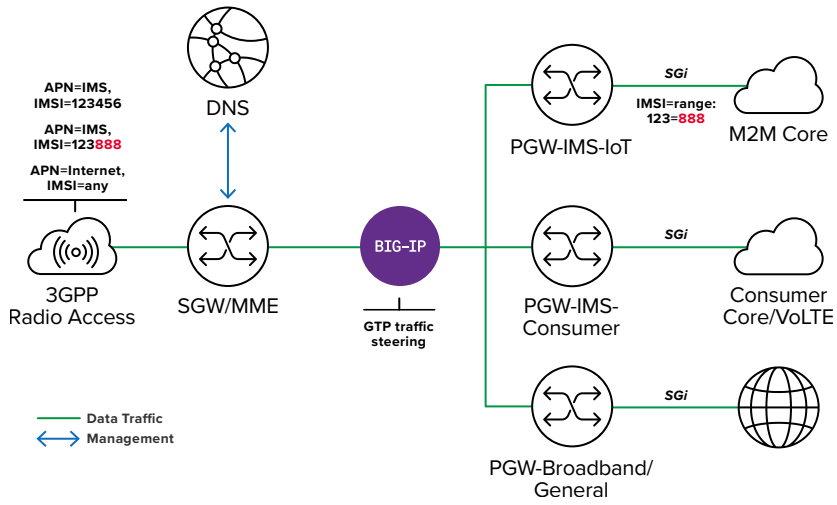

Problem Statement

Sometimes, the same APN=IMS is used for VoLTE/IMS-Consumer services and IMS-IoT, but they need to route to different IMS cores, as IoT traffic is generated from a different IMSI range than the VoLTE/consumer traffic. So a differentiator in addition to APN=IMS is required.

Solution

GTP Traffic Steering can identify IoT traffic, Consumer traffic, and broadband/general traffic and direct it appropriately, enabling different policies for each traffic type. The APN-based DNS resolution points to the BIG-IP solution, provisioned with a table mapping IoT IMSI ranges to the corresponding PGW IMS-IoT. All non-IoT APN=IMS and IMSI ranges are routed to PGW IMS-Consumer. The device steers all APN<>IMS to PGWs for broadband/general traffic.

To determine which IMS platform and PGW APN=IMS traffic should go to, the BIG-IP device must be provisioned with the IMSI ranges of the IoT devices that route to the specialized IMS network, which can only be reached by the assigned PGWs. After defining the IoT IMSI ranges, an APN=IMS in the GTP-C information triggers a lookup to establish whether the IMSI matches the IoT IMSI ranges, after which it's routed to the appropriate PWG. If they don't match, the traffic goes to the PGWs fronting the regular IMS network.

All IMSI ranges are pre-provisioned on BIG-IP devices. It is also possible to verify external databases, after which a routing decision is made and GTP information elements can be updated.

USE CASE 7
Network Slicing

Networks are experiencing massive increases in video traffic and an inherent susceptibility to latency. The same physical network must contend with a virtually infinite number of online devices, largely propelled by the booming Internet of Things (IoT). Network slicing isolates network functions that support the requirements of a particular use case, helping manage new, differentiated services and scale the allocation of high-throughput traffic bandwidth and control plane communications. Each network slice can be optimized to provide the resources and network topology for a specific service and the traffic within it. Functions such as capacity, connectivity, coverage, and speed can be allocated to meet specific use case requirements, while functional components can be shared across different network slices.

In this case, operators shouldn't rely on the pre-provisioned information on BIG-IP devices: rather, they should stay flexible by querying external databases to make routing decisions and change information inside the GTP messages, as needed. This also applies when the right network slice needs to be selected based on information that is in other databases or maintained by other teams, customers, or third parties.

GTP traffic management: GTP Director

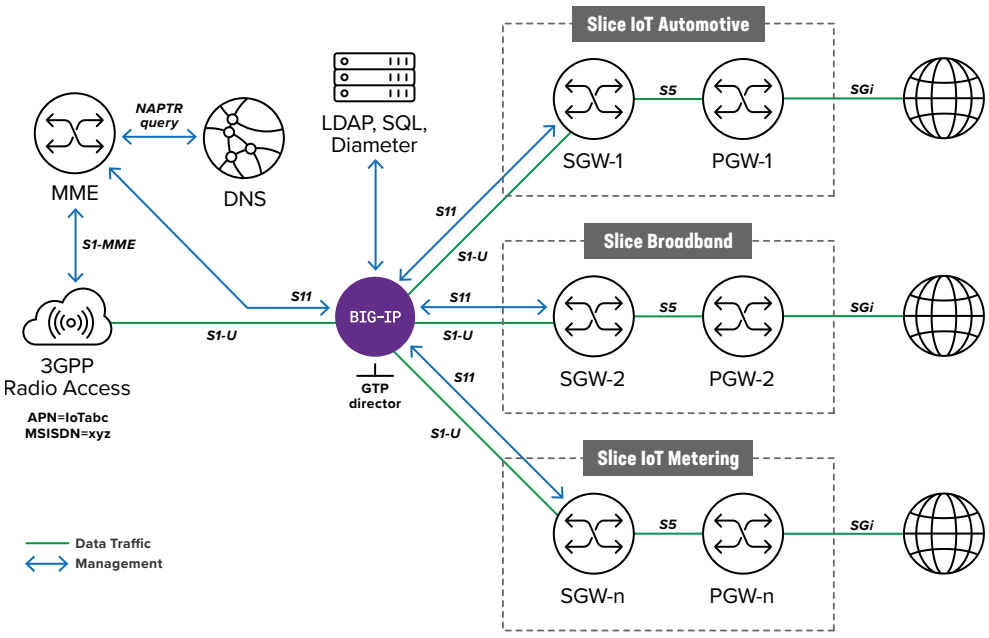

Problem Statement

When Mobile Packet Core selection is based on APN name, traffic direction changes are slow when network slicing for specific services or customers with multiple APNs.

Solution

The BIG-IP device selects the SGW based on DNS, availability, MSISDN, and service using preconfigured/on-board rules to select the right S-GW/slice. Alternatively, the device queries the LDAP database about the subscriber and routes to the SGW/slice based on parameters received from the database.

The BIG-IP device allows verification with external sources to make the right routing decision and change GTP information, if and where needed, using the GTP Director. It's a customer-friendly interface where flow-based decisions can be made, including dropping traffic; LDAP, SQL, or Diameter-based environment queries; collecting routing information; and enriching or changing GTP content, if required by the specific implementation. GTP Director has REST-API based mechanisms to query the defined types of databases while still allowing additional sources to be consulted. Those databases can have fully-customized data structures, with only the relevant information for correct GTP routing or message modification exposed.

USE CASE 8
GTP Firewall

The GSMA has raised awareness of si p and the resulting FS.20 document, GTP Security, serves as a guide for implementing security measures for these vulnerabilities.

Common GTP security issues include confidential data disclosures, denial of service, network overloads, and a range of fraud activities. While security strategies vary from case to case, all organizations should implement solutio ns that provide full traffic visibility and comprehensive distributed denial-of-service (DDoS) protection.

GTP Firewall

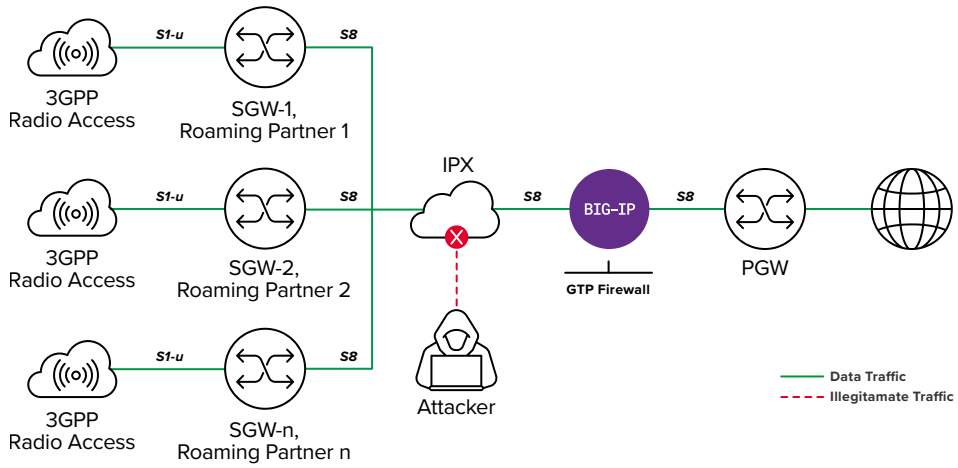

Problem Statement

When the SGW is in a visited network, both GTP-C and GTP-U are transported via the S8 interface. The home operator has no control over the GTP traffic entering its network.

Solution

With BIG-IP solutions in place, GTP-C signaling is checked on protocol conformance and against security rules. Meanwhile, GTP-U user plane traffic is only allowed if a pinhole exists, based on prior receipt of the Tunnel Endpoint Identifier (TEID). In addition, a general check of GTP can be performed for a specific roaming partner.

BIG-IP devices can serve as a GTP firewall to confirm protocol conformance, check against specific security rules, and allow controlled access of user data with pre-aligned security values. The GTP-U stream can also be verified, for instance “GTP-over-GTP.”

With the GTP Protocol conformance, the firewall is confirming whether the GTP-C messages are constructed according to the rules as defined in the 3GPP specifications. For example, GTP information elements are checked for mandatory versus optional, whether the length is follows specification, if the correct variable is used, etc.

GTP security rules ensure various checks are performed to detect and, where possible, act on vulnerabilities described in GTP Security and other cases for which specific firewall rules are written. Organizations are also using machine learning to detect abnormal patterns and generate new firewall rules automatically. Information detected on one firewall can be used to update other firewalls before the same problem occurs.

GTP-U traffic only passes through the GTP application layer gateway if the TEID was initially signaled via GTP-C. If the TEID is unknown, the traffic is not allowed through.

Additional GTP-U security checks include whether the traffic is using the right port, verifying the relevant 3GPP specifications, or if there is a case of “GTP-over-GTP.” When a security check is flagged, an alarm is generated or traffic can be stopped.

Conclusion

GTP is a vital protocol for signaling and transporting mobile data, whether in the initial GRPS networks via 3G and 4G, or the developing 5G network. Strong mobile user data growth requires better scaling of the networks while operators respond to the GSMA-driven initiative to be more aware of the GTP security risks and become better able to respond to vulnerabilities. Through GTP Session Director and other GTP security mechanisms, F5 provides strong GTP security for 3G, 4G, and 5G networks.