5G and CUPS Correlation

Challenges with Traffic Visibility with 5G and 4G LTE CUPS

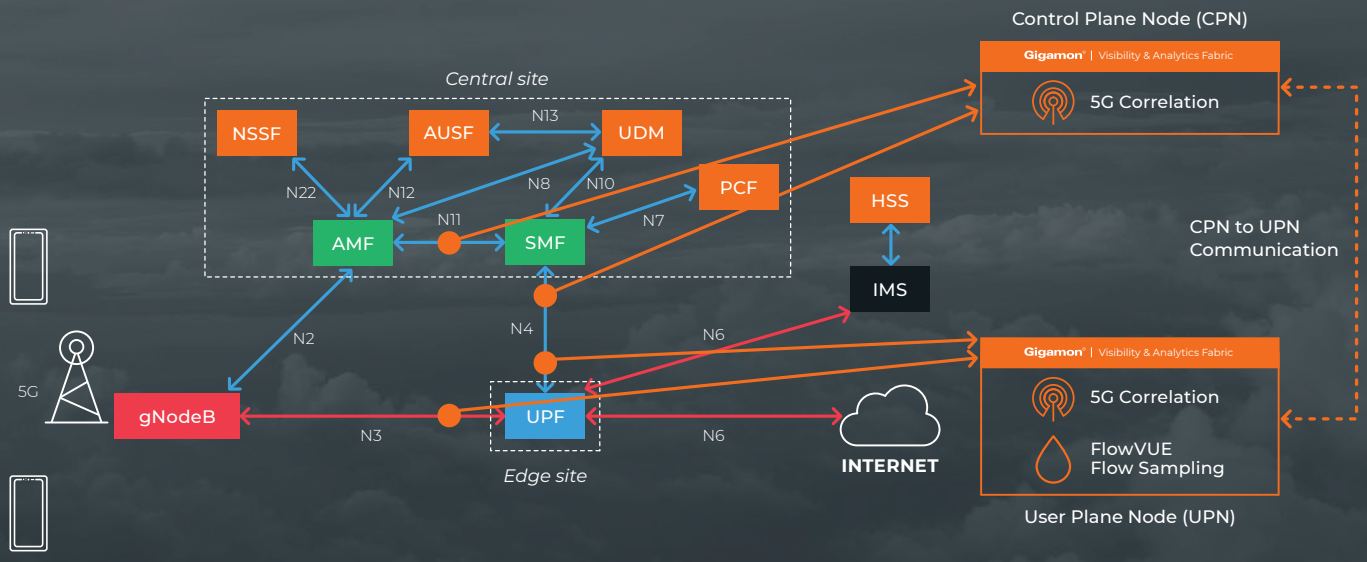

Figure 1. 5G Session correlation for feeding monitoring and security tools

5G CORRELATION FEATURES

- Stateful correlation of HTTP/2 or GTP-C messages with GTP-U sessions, which:
  - Correlates subscriber, user device, RAN and network slice identifiers with corresponding tunnel endpoint IDs (TEIDs)
  - Forwards subscriber-specific control and user plane sessions to a tool or group of tools
- Stateful filtering based on subscriber and equipment IDs (IMSI/SUPI, IMEI/PEI, MSISDN/GPSI and so on)
- Stateful whitelisting based on subscriber ID (IMSI/SUPI), user device ID (IMEI/PEI), RAN ID (ECGI/NCGI) and/or Network Slice ID (NSSAI)
- Stateful user session sampling based on subscriber ID (IMSI/SUPI), user device ID (IMEI/PEI), RAN ID (ECGI/NCGI) and/or Network Slice ID (NSSAI), including allocation of separate samples for each tool port or tool port group from a common pool of correlated control and user plane data
- Stateful load-balancing based on subscriber ID (IMSI/SUPI) or user device ID (IMEI/PEI)

5G networks are poised for exponential growth in traffic and number of connected devices, and mobile carriers are searching for ways to efficiently and effectively monitor performance, Quality of Experience (QoE) and security for their services and subscribers, as well as identify and monetize new offerings. Furthermore, the disaggregation of the control plane and physical separation of the control and user planes (referred to as Control and User Plane Separation, or CUPS) adds additional complexity.

To ensure accurate, cost-effective analytics from their tools infrastructure, service providers are dependent on two critical components:

- Ability to correlate traffic flows on a per-subscriber, user device, RAN or network slice basis
- Visibility across all segments of the mobile network

5G with CUPS

The core of a mobile operator's 5G network uses three main protocols:

- Hypertext Transfer Protocol version 2 (HTTP/2) for the control plane to register, establish, manage and tear down user devices/sessions, which appears only at the control plane locations
- GPRS Tunneling Protocol (GTP) for tunneling the user plane traffic to carry subscriber traffic from the subscriber device to the internet or subscriber services, which appears only at the user plane locations
- Packet Forwarding Control Protocol (PFCP) for communicating between the control plane and the user plane gateway functions

Visibility into a subscriber's traffic requires understanding the subscriber attributes and stateful information contained within HTTP/2 and PFCP, in order to correlate subscriber-specific user sessions carried within GTP and give monitoring tools an accurate view of each subscriber's traffic on the network.

4G LTE with CUPS

The core of a mobile operator's 4G network that employs control and user plane separation uses three main protocols:

- GPRS Tunneling Protocol (GTP) version 2 for the control plane (GTP-C) to register, establish, manage and tear down user devices/sessions, which appears only at the control plane locations
- GTP version 1 for tunneling the user plane (GTP-U) traffic to carry subscriber traffic from the subscriber device to the internet or subscriber services, which appears only at the user plane locations
- Packet Forwarding Control Protocol (PFCP) for communicating between the control plane and the user plane gateway functions

Similar to 5G, 4G LTE with CUPS visibility requires understanding the subscriber attributes and stateful information contained within GTP-C and PFCP in order to correlate subscriber-specific user sessions carried within GTP-U.

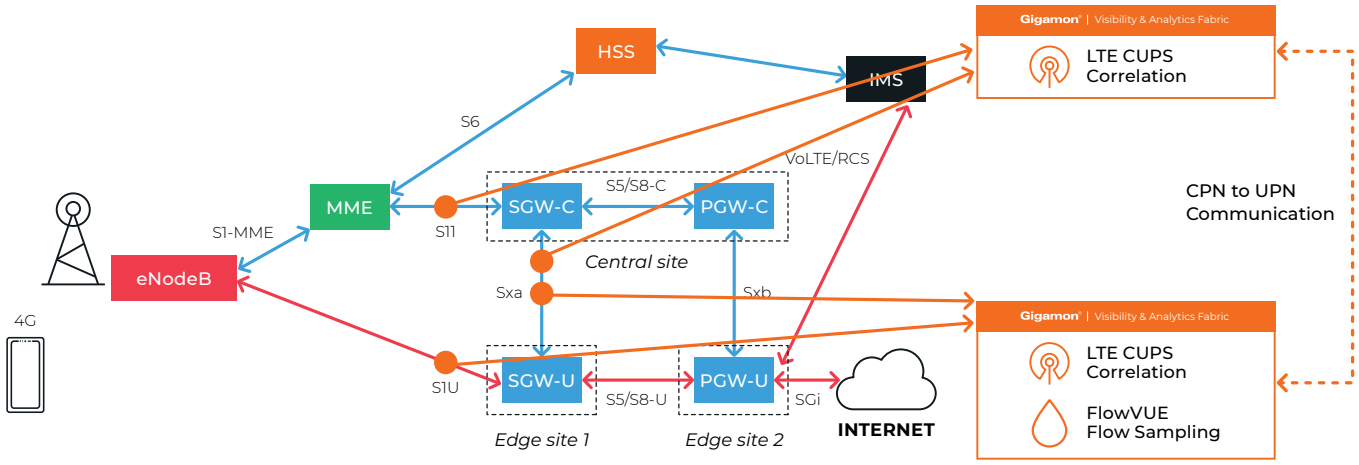

Figure 2. 4G LTE CUPS session correlation for feeding monitoring and security tools

5G and 4G LTE Correlation

The Gigamon 5G Correlation application helps carriers gain access to the subscriber's traffic in the GTP user plane tunnels by reliably correlating and passing all of the identified subscriber's control and data sessions to the analytics/monitoring probes and/or billing subsystems to ensure an accurate view of the subscriber sessions (see Figure 1).

Further, given the rate of increase in the volume of information traversing the mobile service provider network and the expected explosion in the number of connected devices, the tools infrastructure will be unable to scale accordingly — and proliferating tools across the network to monitor millions of subscribers can be expensive and cost prohibitive for both CapEx and OpEx.

With the Gigamon 5G Correlation application, mobile operators can install a monitoring policy within the Gigamon Visibility and Analytics Fabric™ (VAF) that will intelligently deliver subscriber traffic to specific tools. It does this by correlating the subscriber-specific attributes, including the subscriber ID, also known as the international mobile subscriber identity (IMSI) or subscriber permanent identifier (SUPI); device ID, also known as the international mobile equipment identity (IMEI) or permanent equipment identifier (PEI); mobile station intelligent subscriber digital network (MSISDN) number; E-UTRAN cell global identifier (ECGI) or NR cell global identifier (NCGI); network slice selection assistance information (NSSAI); mobile core logical interfaces (for example, Gn/Gp, S11/S1U, S5/S8, N11/N4, N4/N3) carried in HTTP/2 or GTP-C messages. All subscriber-device user plane traffic is carried within GTP-U tunnels, which are identified by tunnel endpoint ID (TEID). Once the GTP-U TEIDs are known and correlated to the subscriber-level attributes, then subscriber traffic can be processed in a subscriber-aware manner.

With this 5G Correlation capability, the Gigamon VAF leverages a subscriber-aware monitoring policy, and can optimize current tool infrastructure investments by forwarding only relevant data to the tool while increasing visibility into subscriber traffic that can help improve QoE, performance and security.

5G Correlation enables important subscriber-aware features that optimize traffic monitoring:

1. Subscriber-aware filtering, whereby subscriber traffic can be selected and forwarded to monitoring tools utilizing attributes, including subscriber identifier, equipment identifier, RAN technology type or mobile core logical interface.
2. Subscriber-aware load-balancing, whereby all control and user plane sessions are load-balanced to as many as 16 tool ports using subscriber or equipment identifier-based load-balancing criteria.
3. Gigamon FlowVUE® flow sampling traffic-scaling application. Correlated FlowVUE ensures that all sessions for a percentage sampled subset of subscribers can be forwarded to monitoring tools, allowing the traffic to be scaled to fit the existing tools. FlowVUE also allows the operator to allocate to each tool separate (and possibly overlapping) samples from a common pool of traffic data.
4. Whitelisting, which allows a defined list of high-value subscribers (over two million) to be chosen to receive full-time monitoring, even with FlowVUE processing enabled.

5G Correlation can be combined with other GigaSMART® traffic intelligence applications, such as Application Filtering Intelligence, which enables monitoring tools to perform more efficiently by eliminating entire sessions of unwanted traffic. The correlated traffic that is destined for servers hosted in an operator's infrastructure can also be decrypted by the GigaSMART SSL Decryption function for malware analysis or inspection by other security functions. The 5G Correlation application also has the ability to ensure that all IP fragments within the subscriber session are sent to the same tool through sampling or IMSI/SUPI, IMEI/PEI, ECGI/NCGI or NSSAI whitelisting.

When used with FlowVUE, our intelligent scaling of active subscribers, carriers can have a representative view of a subscriber's usage patterns. Armed with these subscriber-level insights, Gigamon products can help operators identify roaming subscribers across peered networks through IMSI or SUPI filtering.

By gaining a complete activity view of high-value subscribers, carriers can look to:

- Optimize ARPU by improving operational efficiency, which allows the carrier to better compete with their in-market competitors on an expenses basis
- Ensure business continuity, security and subscriber experience
- Identify and monetize new offerings

Key Benefits

- Optimize the tool infrastructure
  - Optimize tools processing by accurately filtering, replicating and forwarding monitored subscriber sessions
  - Reliably correlate subscriber sessions (control and data) to increase analytics accuracy
- Pervasive visibility into subscriber traffic
  - Extend visibility and proactively identify service issues impacting (and frustrating) subscribers
  - Facilitate drilldowns into roaming users across peer networks
- Maximize QoE and monetize services
  - Gain pervasive subscriber-level visibility and empower monitoring tools to gauge end-user QoE
  - Utilize real-time stateful visibility to enable reliable accounting, billing and subscription management
- Make the transformation to a unified tool rail possible
  - Cost-efficiently correlate user plane traffic at the Gigamon VAF layer instead of needing attached tools to support correlation across multiple probe instances for each tool
  - Ensure that each monitoring and security tool coherently receives control and user plane traffic to help decrease MTTR and reduce ambiguities in the attached tools' results
  - Improve insights by increasing accuracy and precision
  - Distribute correlated traffic samples to reduce demand on each performance monitoring tool

Rethinking Operator Infrastructure Monitoring with Subscriber-Aware Visibility

5G Correlation enables user and data plane correlation. After correlation, both the user and data plane traffic can be directly sent to the tools when the primary objective is offloading tools from the overhead of 5G correlation. The solution goes much further than just tool optimization. It is one of the core building blocks for operators looking to build a best-in-class, modern, subscriber-aware visibility platform, as shown in Figure 1.

By combining 5G Correlation with other traffic intelligence capabilities in the VAF, operators can gain deep insights into their networks and both:

- Optimize per-subscriber monitoring cost
- Offer new services that increase the average revenue per user (ARPU)

This is done with tiered monitoring strategies that separate higher-ARPU subscribers from lower-ARPU subscribers.

Such an architecture enables operators to scale their traffic to meet their tools' processing throughput. Whitelisting allows all traffic from specific IMSIs to be sent to the tools, whereas sampling selects a configurable set of user sessions for analysis. Both whitelisting and sampling are part of the FlowVUE application in the GigaSMART suite of traffic intelligence applications, and can be used by mobile operators in a variety of ways to implement highly scalable and efficient monitoring methodologies.

Some examples are:

- Prevent entire application sessions corresponding to voluminous over-the-top (OTT) traffic such as YouTube, Netflix and other video sites from reaching the tools, eliminating expensive and unnecessary upgrades to the tool infrastructure
- Decrypt SSL traffic destined for servers hosted by the operator and feed it to a security tool for malware inspection
- Send only a sample of non-premium sessions to the monitoring tools for analysis
- Send different (possibly overlapping) samples to different monitoring tools
- Sample a set of sessions to analyze the quality of service at a particular cell site
- Send traffic to tools based on requested network connection (for example, IMS for VoLTE by using APN and QCI as selection criteria)

Pervasive Visibility with a Visibility and Analytics Platform

In this era of Big Data, mobile carriers have searched for a way to efficiently and effectively monitor performance and QoE for their subscribers, as well as identify and monetize new offerings.

Gigamon allows convergence on a single visibility platform that not only simplifies and automates network traffic visibility, but also provides built-in intelligence to address Big Data. This can shape how mobile carriers choose to monitor and manage their networks to provide better, faster connections and new services, while increasing operational efficiency and network uptime.

Legacy approaches to monitoring have offered limited traffic visibility with limited filtering capabilities, are difficult and costly to scale and manage, and often require change orders or network downtime in order to adapt to the evolving network.

Gigamon products provide the architecture and intelligence for mobile operators to create a monitoring infrastructure designed for the era of Big Data and 5G, and deliver pervasive visibility, awareness and control from the converged edge to the cloud. Sitting between the IT infrastructure and the tools that need access to data, the VAF provides a holistic approach to traffic visibility that includes:

- **Architecture advantages:** The GigaVUE® family of visibility nodes offers the volume, port-density and scale needed to connect the right analytical tools to the appropriate large or bonded pipes. Tool trials are streamlined, new tools can easily be added or removed, and uptime is protected while downtime is prevented with a solution that is outside the production network and provides pervasive visibility.
- **Feature advantages:** Advanced filtering, packet manipulation and session-aware traffic identification reduce the amount of data arriving at each tool, while ensuring that the data is formatted precisely for the tool's consumption. Because each tool doesn't need to parse the incoming stream or waste processor cycles on nonrelevant data, it can be optimized and focus on the more important task of data analysis.

- **GigaSMART applications:** Traffic Intelligence, Application Intelligence and Subscriber Intelligence (which includes 5G Correlation and FlowVUE) provide effective monitoring of large-to-massive traffic volumes through the logical reduction of traffic, which is more suitable for connecting to existing tools with lower speed ports or limited capacity. Gigamon 5G Correlation application enables visibility at the subscriber/session level to coherently monitor subscriber QoE and network security, and thus maximally monetize services. Through subscriber-aware flow sampling, the FlowVUE application enables intelligent management of huge traffic volumes while keeping user sessions intact.

For more information, visit the GigaSMART and Visibility and Analytics Fabric pages.

Gigamon®

Worldwide Headquarters
3300 Olcott Street, Santa Clara, CA 95054 USA
+1 (408) 831-4000 | www.gigamon.com

07.20_03