Unlocking the power of context to manage network growth

A shift to better subscriber experience and network utilization

By Monica Paolini, Senza Fili

Sponsored by

1. Introduction.

Identify context and act in real time to provide the best QoE

With the growth in mobile data, operators face more than an increase in the volume of traffic. The complexity of managing that traffic has also intensified, driven by subscribers expecting an ever-higher quality of experience (QoE), and by a traffic and device mix, usage patterns and advanced policy-based services that are more diverse and continue to evolve. This complexity creates an additional set of challenges. Operators not only have to expand capacity, they need to decide how to use that capacity in a way that maximizes the return on their 4G investment, in terms of improving both network utilization (i.e., lowering their per-bit costs) and the subscriber's QoE (i.e., increasing revenues and customer retention).

Use cases
Managing network growth
Geolocation
Roaming
Content optimization
VoLTE and RCS

This is a difficult task because network usage naturally fluctuates in the short term, and it evolves in the long term in a mix of predictable and unpredictable ways. Predictable changes depend on recurring or gradually changing factors, such as location and time-of-day usage peaks, or the growth in video traffic as a percentage of the overall traffic. Operators can respond to these changes using traffic forecasts extrapolated from historical data. Unpredictable changes depend on where subscribers are and what applications they are using, or on some local event, the launch of a new app, or simply some breaking news. They have a major impact on network load and subscriber QoE, and managing them requires a finer-grained, real-time knowledge of traffic distribution across the network.

To manage predictable and unpredictable changes in the traffic load, operators have to collect and analyze the real-time data about network load and performance, as well as subscriber activity, and relate this data to policy – i.e., they have to identify the context in which they provide service. Contextual information is the foundation for providing active subscribers with the best QoE for the applications they are using, given the capacity constraints of their cell sites. How mobile operators will achieve this will vary through time and by location: optimizing traffic during the evening peak hours in a residential neighborhood is likely to require different policy tools and solutions than at midday in a dense urban location.

Although the higher degree of complexity in traffic management imposes a burden, it also presents an opportunity for mobile operators to overhaul their service platform. They can offer services that are richer and more personalized, and they can utilize network resources more efficiently while doing so. By managing this complexity, operators can transform their best-effort networks, which treat all traffic in the same way, into dynamic, context-aware, intelligent networks that deliver a wider choice of differentiated services.

In this paper, we present five use cases that illustrate the challenges that mobile operators face in managing traffic, the tools they can use to address them, and the benefits they can reap from real-time, context-aware traffic management.

© 2014 Senza Fili Consulting • www.senzafiliconsulting.com

2. Managing network growth.

More devices, apps and services increase pressure on the control plane

With the rapid growth of mobile devices and use of mobile services, much of the attention has been on the increase in traffic in the user plane. Yet the implications of network growth go beyond the user plane. The control plane not only has to manage larger volumes of traffic, it also has to use more advanced capabilities to increase the efficiency of RAN resource utilization and to avoid bottlenecks that could bring the network to a halt.

As networks become more dynamic and complex, the relationship between the mobile core and the RAN gets tighter. Core functions have to become context aware, and linked to real-time RAN performance and load. At the same time, the RAN is moving beyond best-effort transmission; it has to handle traffic flows according to policy, as well as optimization and prioritization requirements, set by the operator.

The result is an exponential increase in signaling traffic using Diameter – the protocol that enables network elements to communicate with each other to coordinate and optimize end-to-end transmission, from the internet to mobile devices. Scalability has become a major challenge in managing larger signaling loads, and operators have to address it to avoid service disruption. At the same time, we should abandon the view that signaling is a burden. It provides the foundation on which to establish context, implement charging, and leverage traffic and services management. Effective signaling processing is a prerequisite to many of the advanced services and functionality described in the following use cases.

Traffic growth also leads to an increase in Domain Name Server (DNS) traffic and a higher risk of network disruption from distributed denial of service (DDoS) attacks caused by either malicious external action or by internal issues such as application or network misconfiguration. When the DNS server can no longer process requests, data and voice over LTE (VoLTE) services in the covered part of the network may become slow or unavailable for all subscribers.

Taking control of network complexity

Defining context with Diameter. Diameter routing connects policy-related elements to real-time network performance, subscriber location and usage to provide rich, highly targeted services, to enhance quality of service (QoS), and to optimize content delivery where and when appropriate.

M2M and Diameter. As machine-to-machine (M2M) services become more widely adopted, their varied but specific requirements will require advanced policy and traffic management capabilities. Diameter will be a crucial enabler to ensure the accurate routing of millions of signaling messages.

DNS protection. Effective management of DNS traffic is the first step to ensuring smooth data services and support of high QoE. To protect the network against DDoS attacks and other sudden surges in DNS traffic, operators have to be able to rapidly identify threats and fight them decisively to minimize network disruption. To address new challenges that come with traffic growth and increased malicious activity targeting the mobile infrastructure, operators can use tools like IP-client rate limiting to control traffic surges, and DNS caching can reduce the load on the DNS server. DNS processing complements Diameter functionality to optimize message routing to provide scalability.

A step further with network functions virtualization (NFV). NFV will expand the value of context because it makes it easier for operators to pull together information from multiple sources and to establish a platform that provides the scalability needed to leverage context-aware data.

© 2014 Senza Fili Consulting • www.senzafiliconsulting.com

3. Geolocation.

Matching subscriber location with destination services

Knowing where subscribers are, where they have come from, and where they spend their time is powerful information that only operators have access to across the entire subscriber base. Operators know not only the subscribers' physical location, but also their IP geolocation within the network, which allows operators to route services and apps to the closest and best network resources to optimize QoE.

Over-the-top (OTT) players also have some location-based data, but not for all subscribers and typically not for use of the device outside their app. Subscribers are also increasingly sensitive to the privacy and security implications of OTTs' access to personal data. Because of their neutral position, mobile operators can address many of these concerns by filtering access to personal data based on subscriber preferences.

Basic location-based services – e.g., driving directions – need only limited contextual information. To drive to a nearby address, everybody gets the same directions. When inserted in the wider subscriber and network context, however, location or destination information opens the door to a wider-reaching set of services that rest on a more accurate – and valuable – personalization. You are more likely to be interested in a weekly roaming bundle when you arrive at an airport outside your country, than after you have been in the country for a few weeks and are back at the airport ready to head home. Your location is the same, but your needs are different.

Destination geolocation can also improve traffic management efficiency. Operators can route traffic differently depending on where subscribers are in relationship with core network resources, or use locally cached content. They can also offer specific services or additional traffic allowances to traveling subscribers, with proposed features and prices that can be dependent on network load in the country. To improve the performance and cost of providing these services, they can direct traffic to use local resources instead of transporting traffic to and from their home network.

New service creation using geolocation
Offering location-based services that use location information only (e.g., coupons, traffic directions, location-based information and advertisements).
Pooling contextual information on location, time and subscriber preferences to fine-tune content delivery and optimize traffic routing (local content or home content based on subscriber history and preferences; locally cached content when available). Operators can also offer a choice of local versus home content. Spotify subscribers, for instance, may want to be able to use the service when abroad, but may prefer their home-country music selection and may not be a good target for local ads in a language they does not understand.
Matching location and network information to define real-time service offerings and their pricing: limiting offers during time of network congestion, offering discounts at off-peak times. Cost of providing services and QoE, especially while roaming, may drop when using local network resources.
Using recent geolocation data to target specific services (e.g., offer a transportation pass or roaming bundle upon arrival at a travel destination, in the home language and currency).
Developing services in partnership with OTTs or mediating service delivery on behalf of OTTs using geolocation information. The context of subscriber and network location is extremely valuable to OTTs and content partners. With it, they can more accurately target their services to the operator's subscribers while protecting subscriber privacy.
Complying with emergency-call and legal-intercept requirements that require information on subscriber location.

© 2014 Senza Fili Consulting • www.senzafiliconsulting.com

4. Roaming.

Removing the fear of data roaming with better-targeted, relevant services

LTE networks require a new approach to voice and data roaming, with a move from circuit-switched to packet-based IP voice services and Diameter signaling. With the advanced LTE roaming functionality, operators can leverage Diameter to establish context to support more-advanced roaming services. They can also gain back revenues that today are lost to Wi-Fi and OTT services, such as Skype or Whatsapp, that use IP data connectivity to provide a substitute for more-expensive roaming voice services.

In some regions, regulation requires operators to offer advanced roaming services. The EU Roaming Regulation in the European Union has attracted much attention because it poses tough challenges to operators. The new rules allow subscribers to choose an alternative roaming provider (ARP) for data services outside the home country that can be different from their domestic service provider (DSP). They can do so ahead of traveling, by selecting an ARP using the International Mobile Subscriber Identity [IMSI] method, or from within the visited country by selecting an ARP from a portal (local breakout [LBO] method). To support either method, mobile operators have to establish a relationship with any ARP that requests it and exchange roaming data in real time, act as resellers of the data roaming services, and provision them.

The increased number of roaming providers and the need to decouple service provisioning from charging add complexity to the roaming platform. Not only does the mobile operator have to manage the roaming traffic, it has to exchange real-time information to know what services are available to the subscriber (e.g., for prepaid subscribers when they have no more credit available) and have to allow the ARP to charge for them appropriately. In addition, the mobile operator may want to support policy-based services for roaming traffic or monitor usage to prevent fraud. This function can be introduced at a later stage, but from the onset, roaming traffic has to be integrated within the operator policy to decide how to manage roaming versus non-roaming traffic.

A richer, more flexible roaming experience

- Diameter is crucial to enabling LTE roaming for both basic and advanced services, because it links the policy and charging rules function (PCRF) to the profiles from the home subscriber service (HSS) database that are necessary for authentication and mobility management.
- Operators can tailor LTE roaming as a function of policy, subscriber preferences, location, security requirements, and availability of network resources. This dynamic approach gives operators more flexibility in the use of network resources, increasing efficiency, reliability and coverage. Tools such as load balancing provide the additional scalability needed to deal with the increase in signaling driven by the increase in roaming traffic and policy-driven services.
- Better-targeted services are more attractive to subscribers and can lead to higher revenues and QoE. For instance, roaming can be expanded to include tiering (i.e., premium access for top-tier subscribers), app-based options (e.g., social networking, video streaming), or rich communication services (RCS; e.g., video calling).
- Operators can use network load information to fine-tune service prices for roaming as they do with on-network services, as long as they can establish tight coordination with their roaming partners.
- Security is a bigger threat for roaming in LTE than in 2G or 3G. Through the S/Gi interface, the Diameter edge agent (DEA) provides secure connectivity, protecting the network both from mobile devices, and from roaming partners and the internet. In addition to Diameter, Internet Protocol security (IPsec) secures LTE roaming traffic as it crosses the mobile core borders.
- With VoLTE adoption and EU roaming regulation, LBO will become a widely used service, and it will drive growth in data roaming usage.

© 2014 Senza Fili Consulting • www.senzafiliconsulting.com

5. Content optimization.

A light touch to improve experience, fine-tuned by context

We are coming to the end of the era of pushing content through the network as is, on a best-effort basis that takes into account neither who the subscribers are and what they are doing with their devices, nor network availability and network load. There is a growing realization that content has to be optimized to allow operators to use their network resources more efficiently and provide a better subscriber experience. However, establishing and implementing the right tradeoffs in content optimization entails a very complex framework that depends on the interaction of multiple factors (see table at right). Once a mobile operator selects some content optimization tools from among the wide range available, it has to choose which tools to use when and, where relevant, how they should interact with each other.

Optimizing all content may have negative effects, such as unnecessarily degrading the subscriber experience or reducing revenues. As an example, video optimization in a high-traffic location at peak time reduces the overall traffic load and is likely to translate into a better experience for all subscribers – not only for those using video. When the network is congested, individual subscribers prefer a lower resolution to a stuttering video feed or no video at all. When there is no congestion, the same subscribers have a better experience with higher-quality video, and there is virtually no marginal cost to providing uncompressed video in a network that has unused capacity. At the same time, though, compressed video generates lower traffic volume, and hence may reduce the overall per-subscriber traffic and the revenues from the subscriber.

The initial approach of deploying content optimization techniques such as throttling, compression, transcoding or transrating across all traffic or with some basic limitations (e.g., optimize video only) is proving to be unnecessarily strict, and in some cases expensive. Operators increasingly want to be in control of when and where they optimize content – or transmit it in its original form. To do so, they need easy access to fine-grained, real-time contextual data to enforce policy.

How to choose which content to optimize?
▪ Subscriber:
– Plan features and tier (e.g., higher-quality video for top-tier subscribers) and subscriber preferences (e.g., allow video compression to reduce impact of video on data allowance).
– Monthly traffic limit (e.g., compress content when approaching limit).
▪ Device:
– Optimize content based on device capabilities (e.g., screen size).
▪ Traffic type:
– Focus on content optimization where scope for optimization is larger (e.g., exclude social networking traffic, manage video traffic more extensively, and give voice priority over other traffic).
▪ Network selection:
– Manage traffic depending on network (e.g., compress content on 3G, but not on Wi-Fi).
– Select network based on service (e.g., voice over 3G, video over Wi-Fi).
▪ Network utilization:
– When network operates at capacity, increase efforts to optimize traffic and, if needed, restrict access to services to low-tier users.
▪ Roaming status:
– Use separate policy for roaming traffic when requested by subscriber (e.g., compress all traffic when roaming).
▪ OTT or third-party traffic:
– Follow specific policy for OTT/third-party traffic (e.g., sponsored data).

© 2014 Senza Fili Consulting • www.senzafiliconsulting.com

6. VoLTE and RCS.

Extending context awareness to voice and video calls in LTE

The expanding role of data traffic has reduced voice traffic to less than 10% in many mobile networks, but voice performance remains crucial for all operators. With VoLTE and rich communication services (RCS), mobile operators have the opportunity both to improve the quality of voice services and to bring voice into the fold of the wider set of data services, offering integrated voice and data and richer applications.

To take full advantage of this opportunity, mobile operators have to embrace the same context-aware approach used for mobile services in defining and provisioning voice services. The challenge for operators is to leverage the context information within the session information protocol (SIP) and the IP Multimedia Subsystem (IMS). A central role in meeting this challenge goes to Diameter, for its ability to interface with SIP and IMS elements in the core, and because it enables policy enforcement charging and traffic management for voice and other RCS traffic.

With a successful integration of Diameter with SIP and IMS, mobile operators can apply their policy and QoS tools to voice traffic over VoLTE and RCS video calls, allow for real-time charging, and manage this traffic on the basis of real-time network load or roaming status.

Consistency and interoperability are added advantages that operators gain. Increasingly, service plans include both voice and data as a single unit (i.e., data is no longer an add-on option, selected and charged for separately), and applying the same subscriber policy makes plan definition easier to communicate and enforce. For instance, the operator can provide the same high-priority access to top-tier subscribers, regardless of whether they use data services or voice/video calls.

Beyond basic voice services

- Consistent use of policy and QoS for voice, video calls and data:
  - Give high priority and better voice quality to top-tier subscribers.
  - Manage voice and video-call traffic on the basis of policy and real-time network load.
- Establish how to manage voice and video-call traffic in relation to data traffic (e.g., assign voice traffic highest priority regardless of subscriber tier).
- Optimize video-call traffic depending on context information, such as subscriber plan, device, network load, or network availability.
- Specify preferred networks for voice traffic (e.g., avoid Wi-Fi for voice traffic).
- Improve QoE by optimizing voice and video calls depending on network conditions (e.g., providing high-quality voice and video only when sufficient network capacity is available), or by using QoS to allocate sufficient network resources to voice and video calls.
- Extend LBO and other advanced roaming services to voice and video calls, to increase roaming traffic and revenues.
- Manage voice and video-call traffic from partner OTTs or as part of sponsored plans:
  - Depending on partnership status, operators may treat OTT and sponsored traffic as high priority or best effort.
  - OTT traffic may be treated differently depending on the subscriber plan (e.g., paying Spotify subscribers may get higher priority than free-service Spotify subscribers).

© 2014 Senza Fili Consulting • www.senzafiliconsulting.com

7. Implications.
Taking control of network growth.

In most mobile networks today, data traffic is still best effort: every subscriber, every application, and every device is treated in the same way. When the network becomes congested, the QoE quickly deteriorates for everyone. As adoption of mobile data services expands, this approach to traffic management is fundamentally inadequate. It fails to provide the experience that subscribers expect, and it fails to utilize network resources effectively. As a result, operators fail to get the return on their infrastructure investment and the subscriber revenues that their network could support.

In this paper, we argued that to take control of network growth, operators have to actively manage traffic using context and real-time network information to enforce policy. With this approach, operators can use network resources more efficiently, support new services, and improve QoE. We explored how they can do so in five use cases: managing network growth, geolocation, roaming, content optimization, and VoLTE and RCS.

Use case	What is needed
Managing network growth	Support policy enforcement using real-time, contextual information with Diameter Avoid service disruption from malicious action or from application and network misconfiguration with DNS protection
Geolocation	Develop new rich services that depend on physical location and IP geolocation for functions like content selection and traffic routing.
Roaming	Leverage Diameter and advanced LTE roaming functionality to promote adoption of data roaming, a way to increase revenues and support features such as LBO in the European Union.
Content optimization	Selectively apply content optimization solutions, depending on network load conditions, operator policy, and subscriber preferences.
VoLTE and RCS	Extend use of policy and QoE to VoLTE, video calls and other RCSs, to improve QoE, generate additional revenues, and expand the range of voice and video services.

© 2014 Senza Fili Consulting • www.senzafiliconsulting.com

About Senza Fili

Senza Fili provides advisory support on wireless data technologies and services. At Senza Fili we have in-depth expertise in financial modelling, market forecasts and research, white paper preparation, business plan support, RFP preparation and management, due diligence, and training. Our client base is international and spans the entire value chain: clients include wireline, fixed wireless, and mobile operators, enterprises and other vertical players, vendors, system integrators, investors, regulators, and industry associations. We provide a bridge between technologies and services, helping our clients assess established and emerging technologies, leverage these technologies to support new or existing services, and build solid, profitable business models. Independent advice, a strong quantitative orientation, and an international perspective are the hallmarks of our work. For additional information, visit www.senzafiliconsulting.com or contact us at info@senzafiliconsulting.com or +1 425 657 4991.

About the author

Monica Paolini, PhD, is the founder and president of Senza Fili. She is an expert in wireless technologies and has helped clients worldwide to understand technology and customer requirements, evaluate business plan opportunities, market their services and products, and estimate the market size and revenue opportunity of new and established wireless technologies. She has frequently been invited to give presentations at conferences and has written several reports and articles on wireless broadband technologies. She has a PhD in cognitive science from the University of California, San Diego (US), an MBA from the University of Oxford (UK), and a BA/MA in philosophy from the University of Bologna (Italy). She can be contacted at **monica.paolini@senzafiliconsulting.com.**

© 2014 Senza Fili Consulting, LLC. All rights reserved. This white paper was prepared on behalf of F5 Networks Inc. The views and statements expressed in this document are those of Senza Fili Consulting LLC, and they should not be inferred to reflect the position of F5 Networks. The document can be distributed only in its integral form and acknowledging the source. No selection of this material may be copied, photocopied, or duplicated in any form or by any means, or redistributed without express written permission from Senza Fili Consulting. While the document is based upon information that we consider accurate and reliable, Senza Fili Consulting makes no warranty, express or implied, as to the accuracy of the information in this document. Senza Fili Consulting assumes no liability for any damage or loss arising from reliance on this information. Trademarks mentioned in this document are property of their respective owners. Cover page photo by Slavoljub Pantelic/Shutterstock.

F5 Service Provider Portfolio: Solutions to Unlock the Power of Context to Manage Network Growth

F5 is dedicated to enabling service providers to leverage LTE to provide a superior customer experience. Context aware technologies play a primary role in our approach in solution design for service providers to stay on top of the telecom value chain for mobile consumers. Context aware functionality allows service providers to maintain top quality, network performance, while expanding their service and product offerings in a customized fashion to specific audiences. Here are some examples of F5 context aware solutions:

Diameter Signaling Management Solutions

Diameter signaling messages serve as an excellent source of information on network operations and subscribers, which when extrapolated, may be used to differentiate service offerings and improve the customer experience. F5's industry-leading Diameter Traffix Signaling Delivery Controller™ (SDC™) solution gives operators the required network visibility into the control plane. It routes Diameter messages according to an unlimited combination of AVP values for optimal routing flexibility. It uses the granular information found in Diameter messages for network management, scale, and optimization. As the market's most mature Diameter solution, the SDC consolidates a Diameter Routing Agent (DRA), a Diameter Edge Agent (DEA), a Diameter load balancer, and a Diameter interworking function (IWF) on a single platform. Operators benefit from context-aware intelligent routing, reliable load balancing, and flexible, seamless connectivity for fast to market roaming solutions and many other use cases such as billing optimization.

Intelligent Traffic Management and Policy Enforcement Solutions

F5 offers intelligent traffic management solutions on a unified platform that simplifies delivery of network services such as dynamic service chaining. Using context and subscriber aware technology, BIG-IP® Policy Enforcement Manager™ (PEM) offers a full proxy architecture and rich IP capabilities for critical traffic visibility and analytics and sophisticated traffic steering capabilities, including the ability to inspect and route traffic based on data type and subscribers' profiles. By leveraging this intelligent information, PEM enables operators to implement bandwidth controls per subscriber or application, along with dynamic traffic steering that simplifies delivery of network services.

DNS Services to Manage Network Growth

F5 offers comprehensive control plane solutions that optimize, intelligently scale, and securely manage messaging interfaces such as RADIUS, DNS, and SIP. F5 BIG-IP® Global Traffic Manager™ (GTM) is our DNS solution for service providers to optimize their LDNS infrastructure and deliver a higher subscriber quality of experience, resulting in increased revenues and reduced churn. BIG-IP delivers a high performance DNS authoritative solution scaling to surpass demand and securing your DNS infrastructure rom distributed denial-of-service (DDoS) attacks. With destination geolocation, operators route traffic differently depending on where subscribers are in

relationship with core network resources, or use locally cached content. GTM delivers faster DNS responses, provides optimized access to mobile services, and enables an enhanced subscriber experience.

Application Delivery Firewall Security Solutions

F5 offers integrated, high-performance ICSA security solutions, such as the F5 Advanced Firewall Manager that protects the entire network infrastructure, and scales to perform under the most demanding conditions. Operators benefit from the solution's intelligence and flexibility for enhancement and simplification of network security in the increasingly threatening landscape, with a common platform to deliver applications and improve responsiveness. F5 security solutions offer a stateful, full-proxy network firewall with unparalleled session scale, throughput, and connections per second, and defend against DDoS attacks across all layers including network, session, and application, including SYN, floods and IP port scan attacks, DNS floods, and SQL injections.

In a roaming scenario, the SDC's Diameter edge agent (DEA) provides secure connectivity, protecting the network from connectivity with roaming partners. In addition, Internet Protocol security (IPsec) secures LTE roaming traffic as it crosses the mobile core borders.

SDN and NFV Solutions

F5 solutions provide service providers with the ability to move to software-defined networking (SDN) and Network Function Virtualization (NFV) architectures. F5 products are available across a variety of platforms, from highly scalable physical devices with virtualization options to virtual editions, ready-to-run F5 solutions on all major hypervisors. The flexible and programmatic F5 APIs give service providers the tools to tie F5 solutions into virtualization management and orchestration systems for an agile and dynamic network. This includes the ability to monitor service utilization, proactively add resources as demand increases, and deprovision resources when demand decreases.

For more information on F5 solutions for communications service providers, visit the Service Provider section of f5.com or contact: info@f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com
F5 Networks, Inc.
401 Elliott Avenue West,
Seattle, WA 98119
888-882-4447
www.f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com
F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.c

©2014 F5 Networks, Inc. All rights reserved. F5, F5 Networks, and the F5 logo are trademarks of F5 Networks, Inc. in the U.S. and in certain other countries. Other F5 trademarks are identified at f5.com. Any other products, services, or company names referenced herein may be trademarks of their respective owners with no endorsement or affiliation, express or implied, claimed by F5.