# HSS RESILIENCY

## SURVIVING ATTACH STORMS IN THE NETWORKS OF TODAY AND TOMORROW

STRATEGIC WHITE PAPER

Mobile service providers are taking advantage of the lower operational costs associated with Long Term Evolution (LTE) networks by deploying network equipment with the largest possible capacity, even if they don't need all that capacity immediately. One side effect of this concentration of capacity is that a sudden, unexpected failure by one of these elements will ripple through the network and subject other network elements to a much larger than normal load. This can cause service outages that disrupt end user services and generate attach storms. To eliminate potential problems caused by attach storms, network elements must be able to shed excessive data traffic load gracefully. This will lower the impact on the network and enable it to return to a steady state condition automatically without causing a long term outage.

# TABLE OF CONTENTS

# MANAGING THE MOBILE DATA TRAFFIC TSUNAMI

Today’s mobile networks must support an ever-increasing number of advanced devices executing an increasing number of functions and generating an overwhelming volume of data traffic. To keep up with this trend and prepare for future requirements, mobile service providers are deploying Long Term Evolution (LTE) networks. LTE is the clear choice for a fourth generation wireless technology that can support the volume of mobile data traffic generated by feature phones, smartphones, super phones, tablets, notebook computers, and mobile modems. It offers the superior data speeds and bandwidth end users are demanding today, and will require tomorrow. It also offers the lower operational costs per megabyte that service providers need.

Unfortunately, by developing the network to provide more bandwidth and higher speeds, service providers are also creating the need for even more capacity. The end user applications that will use the extra bandwidth and speed will eventually generate more traffic that will require even more bandwidth. This ongoing cycle can only be addressed by additional network enhancements.

As a result, mobile service providers are planning ahead. They are taking advantage of the lower operational costs associated with LTE networks by deploying network equipment with the largest possible capacity, even if they don’t need all that capacity immediately. One side effect of this concentration of capacity is that a sudden unexpected failure by one of these elements will ripple through the network and subject other network elements to a much larger than normal load. This can cause service outages that disrupt end user services. Given the power of today’s social media, a service provider can quickly gain a reputation for poor service. This will make it difficult to maintain average revenue per user (ARPU), retain current subscribers and attract new ones.

To eliminate potential problems, network elements must be able to shed excessive data traffic load gracefully. This will lower the impact on the network and enable it to return to a steady state condition automatically without causing a long term outage.

The Home Subscriber Server (HSS) is one of the elements in a LTE network and its associated Evolved Packet Core (EPC) that is most affected by a network element failure. The HSS keeps track of LTE subscribers and the elements in the network serving their data connections. If a failure of a component, such as a gateway, occurs the HSS is subjected to short duration loads potentially two orders of magnitude greater than what it normally processes. This is followed by longer duration loads of an order of magnitude or more while mobile devices attempt to re-attach to the network.

If the HSS is not able to process its part of the procedure required to re-attach a mobile device to the network then HSS resources are wasted and the mobile device continues to attempt to attach. This cycle continues to increase overall load and leads to an “attach storm” that places extreme loads on the entire LTE network. If the HSS does not weather the attach storm, then a protracted outage occurs. Therefore, the HSS must have the ability to efficiently and successfully shed high loads and continue processing to minimize disruptions to LTE subscribers.

# ATTACH STORMS AND THEIR IMPACT ON THE HSS

Every LTE network is built with elements that support a large number of user devices. A simple box level view of an LTE network can be created with six components (Figure 1):

- The **user equipment (UE)**, which represents the mobile devices connected to the network
- The **eNodeB**, which is the radio node
- The **Serving Gateway (SGW)**, which is the router that serves as the local anchor point and interface to the radio network
- The **PDN Gateway (PGW)**, which is the router that serves as the global anchor point and interface to the Internet
- The **Mobility Management Entity (MME)**, which is the signaling element that coordinates authentication and resource allocation
- The **HSS**, which manages subscriber profiles, authentication and high level state

**Figure 1. A simple box level view of an LTE network**

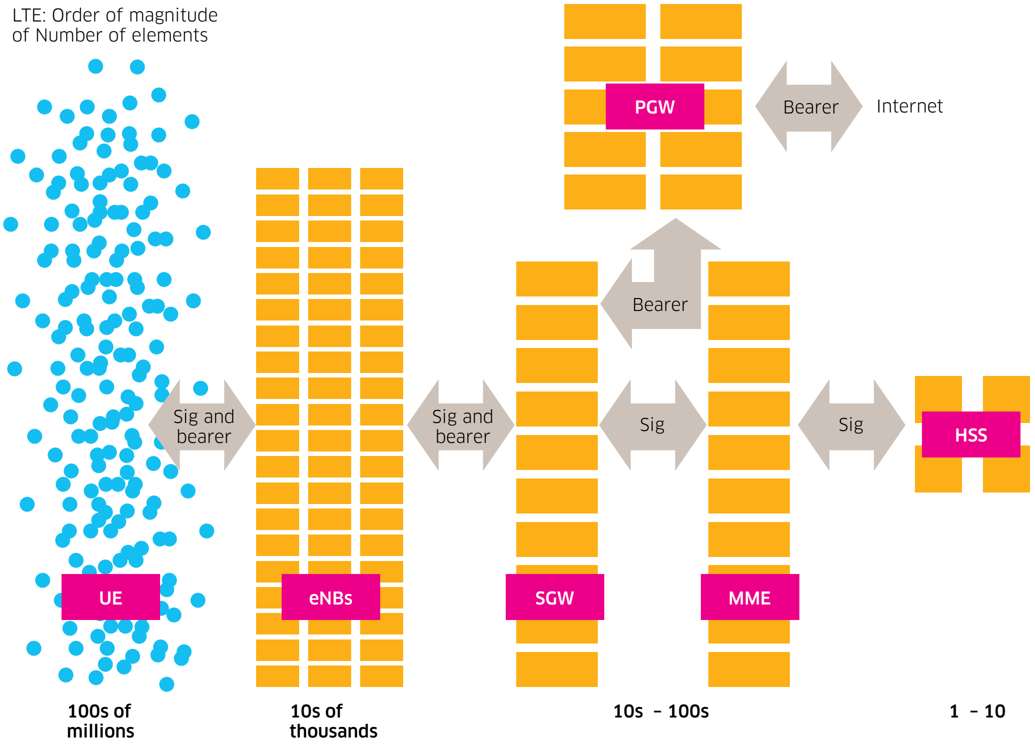

In every LTE network, there is an increasing concentration of data traffic from UE into fewer elements as it moves from the UE to the HSS, with the HSS having the fewest elements. This makes the HSS a critical player during unexpected failures of other elements, such as the PGW, SGW, and MME.

## Attach storm conditions

An attach storm is created when a failure of one of these elements causes a large number of UE to attempt to re-attach to the network in a short interval. For example, every mobile “bearer” must traverse a PGW. If that PGW fails, then the SGW detects the failure, removes that bearer and indicates its removal to the MME. The MME then updates the HSS with that information. Once all the bearers are removed, then the SGW indicates that the UE needs to re-attach to the network.

Typically networks are configured such that all bearers for UE go through one PGW making this situation likely. The impact on the HSS of the failure is twofold.

First, the changing of state of all the bearers served by that PGW causes a large number of Notification Requests (NORs) from the MME to the HSS — one for every Policy Decision Point (PDP) Context served by that PGW. These NOR requests arrive at the HSS spread over two to three seconds. Current experience with LTE networks indicates that the relationship between a normal attach load, which contains two to three transactions, on the HSS and PDP Contexts served by a PGW is approximately 1:1000. Thus, during a failure, the HSS (absent any throttling by the MMEs) will see a spike of a load whose magnitude can possibly be over 100 times normal transaction load. In practical terms, this translates into hundreds of thousands of NOR transactions in a very short period in a large LTE network.

It should be noted that recent standards changes offer the option for the NOR message to not be sent to the HSS during failures. If this option is deployed in the EPC it eliminates the NOR spike and thereby reduces or eliminates the extreme load condition from the changing of state of the bearers. However, it does not reduce the load from the UE re-attaching to the network.

The second impact on the HSS is the result of the load created by the subsequent re-attaches of the UE to the network over time. There is some variability in the mechanism and timing of the re-attach (per standards, an immediate re-attach is specified but implementations appear to deviate from that). But, typically, all UE will attempt to re-attach within approximately 120 seconds. Any device that cannot attach initially will then retry after a period of time. Current network experience indicates that this is approximately every minute (note that exponential back-off would help the situation, but does not appear to be implemented in mobile networks). Depending on the number of devices affected by the failure, this can almost immediately cause a plateau that is five to 10 times a normal mobile attach load on the HSS.

## Attach storms and load

Ideally the attach storm is dealt with at this point. The HSS systems should quickly work through the backlog of UE trying to attach and return to steady state. If they do not do this in short order, then the normal process of UE coming onto and off of the network increases the overall load — especially on the HSS, which is serving multiple MMEs.

Figure 2 shows the offered load to an HSS where the HSS fails due to the initial spike of traffic or due to an inability to process given the high load. Within a little over two minutes of a PGW failure the load on the HSS is 16 times normal and increasing at a rate of 1x normal load per minute. This is a situation where the HSS can never catch up and drastic manual action is needed to restore the network to a steady state.

Figure 2. Example of an offered load to an HSS during a failure

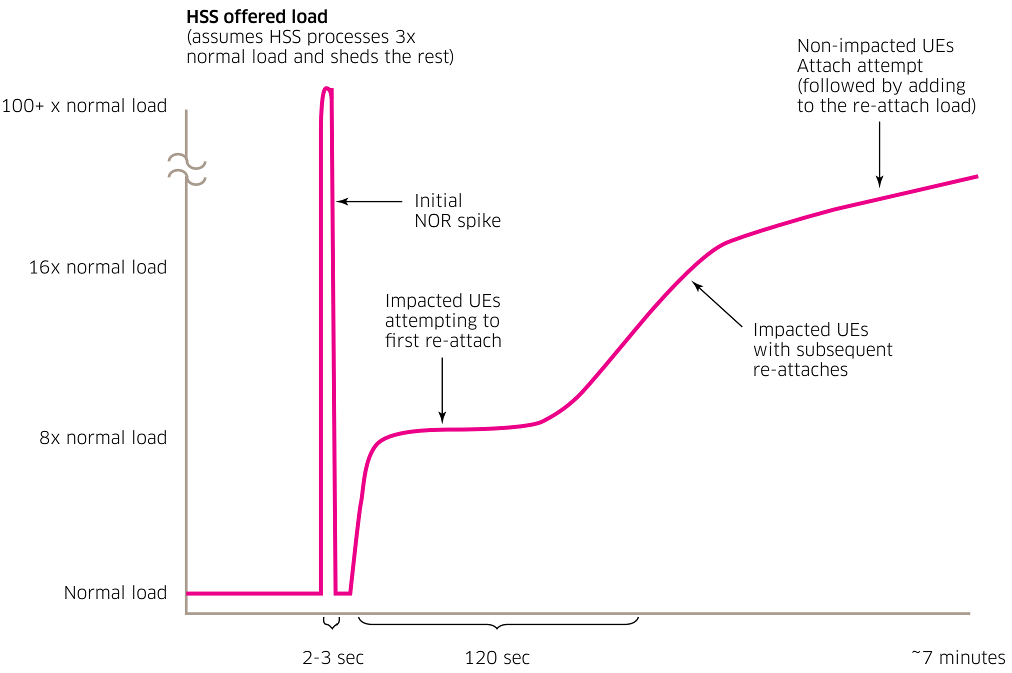

# EFFECTIVE HSS ATTACH STORM COUNTERMEASURES

There are several options available to resolve an HSS attach storm.

The first option is to install HSS capacity at 10 or 20 times normal load so that the attach storm is smothered quickly. The drawback with this approach is that the capital and operational expenses for that many HSS systems is quite high.

Another is to partition or regionalize the network such that an overload situation does not spread beyond that partition. Unfortunately, to achieve effective operations and capacity planning with this approach the partitions/regions must still be large. Plus an overload situation within the regions would still be necessary.

A third approach is to adjust mobile terminal re-attach timing (for example, exponential back-off). However, this does not address the NOR spike created by the closing of a large number of PDP Contexts.

The ideal approach is to ensure the HSS can shed excessive load gracefully and continue processing at (or near) capacity.

## Shedding load

A graceful shedding of load can be achieved by enabling the HSS to manage the attach requests more efficiently. When the initial spike occurs, the system can be configured to process only a fraction of the load because it is more important for the HSS to stay in operation than to try and cover an offered load more than 100 times normal. The attach storm coming after the initial spike is the most important.

For example, if the HSS is sized at three times normal capacity (to take into account geo-redundancy and some headroom) then when offered eight times normal load in the 120 seconds after the spike, it will serve three times the normal load and shed the remaining five times (Figure 3). After the initial 120 seconds, when the period on re-attaches is a minute, the HSS processes approximately 10 times the normal load (reduced from 16 times since the mobile devices that have had successful attaches do not try to re-attach). From that point on, the HSS processes the normal load from elsewhere in the network (approximately one time normal load) and has two times the normal load to apply to the backlog that has built up. At two times capacity to apply, it takes approximately five minutes for the HSS to return to a steady state (seven minutes after the initial incident).

**Figure 3. A graceful shedding of load by enabling the HSS to manage the attach requests more efficiently**

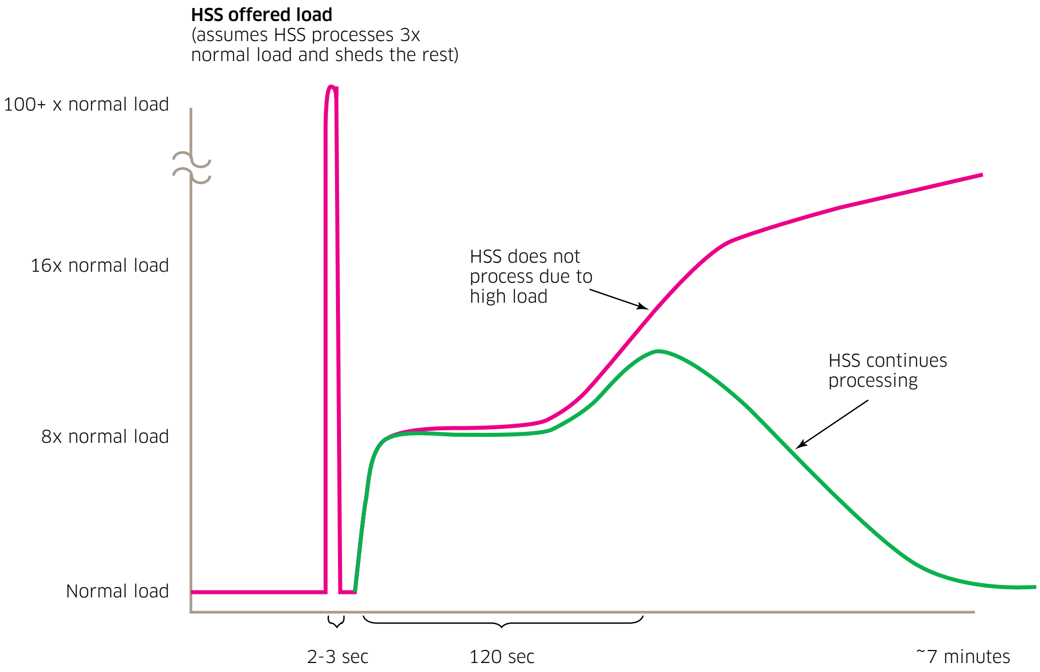

## Mechanism for shedding load

To protect its integrity and be able to continue processing, the HSS must be able to shed excess load that it cannot handle effectively in the application layer and the Diameter stack.

The critical elements to consider in the application layer are the CPU load and application queues. If an overload situation is detected in these measures, then a pre-determined number of various message types are discarded. The flexibility on a per message basis allows service providers to adjust percentages so that sequential operations can be better supported (when two sequential operations are required then it is better to discard the first operation rather than the second operation).

A similar approach is required in the Diameter I/O queues. Once various levels are reached a transaction rejection occurs, independent of the type of operation. For example:

- Overload1Threshold = 60 crossed then reject 10 percent of all messages
- Overload2Threshold = 70 crossed then reject 30 percent of all messages
- Overload3Threshold = 80 crossed then reject 60 percent of all messages
- Overload4Threshold = 90 crossed then reject 100 percent of all messages

Thus if the offered number of messages is very high then essentially all of the messages are discarded. This is exactly the behavior needed for an HSS to weather a NOR spike, and then start processing the re-attach load.

### Examples of overload in other networks

This overload type of situation is not unique to LTE networks in particular or telecom networks in general. For example, the Amazon® cloud service recently underwent an outage[1] triggered by a router configuration error. Servers in the network created a "replication storm" that required manual interaction and a large amount of additional capacity to be added to the system to bring it back to a steady state.

Another recent example is a "registration storm" in the Skype™ network resulting in a long outage.[2] In this case a software fault in the client caused the PC clients to crash in certain circumstances. When the clients re-started they attempted to register and connect to Supernodes. The resulting load caused the Supernodes to fail when faced with 100 times the normal load. The network was only able to return to a steady state with the addition of massive amounts of temporary capacity.

## ALCATEL-LUCENT APPROACH TO ATTACH STORM MANAGEMENT

Alcatel-Lucent provides effective management of attach storms as an integral feature of the Alcatel-Lucent 8650 Subscriber Data Manager.

### Alcatel-Lucent 8650 Subscriber Data Manager (SDM)

The Alcatel-Lucent 8650 SDM is an efficient, next-generation converged database product that consolidates Home Location Register/Authentication Center (HLR/AuC), Authentication, Authorization and Accounting (AAA), and Home Subscriber Server (HSS) data from multiple GSM/UMTS HLRs, AAAs and HSSs into a single, virtual data store with centralized administration, management and reporting.

A key component of the Alcatel-Lucent Subscriber Data Management solution, the Alcatel-Lucent 8650 SDM data (profile) centralization capabilities help service providers gain a real-time, 360-degree view of their customers. The centralized subscriber data immediately becomes more relevant to a service provider and can be used to create the highly personalized services that customers now expect. Centralized subscriber data is also a key requirement for service delivery environment (SDE) and service delivery platform (SDP) strategies, and for the transformation to all-IP networks and applications.

[1] http://aws.amazon.com/message/65648/
[2] http://blogs.skype.com/en/2010/12/cio_update.html

## Attach storm mitigation

To mitigate against attach storms, the 8650 SDM is designed to shed load, as required.

The physical architecture of the 8650 SDM is partitioned into front end nodes and back end nodes. The front end nodes are stateless and provide protocol processing, while the back end stores the subscriber data. The back end nodes are partitioned into a geographic redundant mated pair arrangement, the Network Redundancy Group (NRG), where each node is a MySQL cluster with replication between them. The front end nodes are stateless so that any front end node can process incoming requests from the MME.

The front end nodes use a dynamic index server to determine which back end node contains a particular subscriber's data. When a subscriber index is not found in the cache, the index server is populated with a broadcast request sent to one node in each of the NRGs. If a back end node does not serve the specified subscriber's data, then it silently discards the message. If the back end node has the subscriber data, then it responds to the front end node query.

Since a failure in the network, such as a MME failure, may result in excess traffic to front end nodes, the dynamic index server in those nodes may not have those index entries. This will place extra load on the back end nodes that do not have the subscriber data.

Alcatel-Lucent has tested the 8650 SDM to determine how it will manage re-attach traffic under extreme excessive load. The load tools used generated a re-attach load of over eight times engineered attach load to partially simulate the scenario described in the main body of this paper. The 8650 SDM managed to shed load as required and operate well within the limits associated with surviving an attach storm caused by major failures in a LTE network (Figure 4).

**Figure 4. Alcatel-Lucent 8650 SDM operates well within the limits associated with surviving an attach storm**

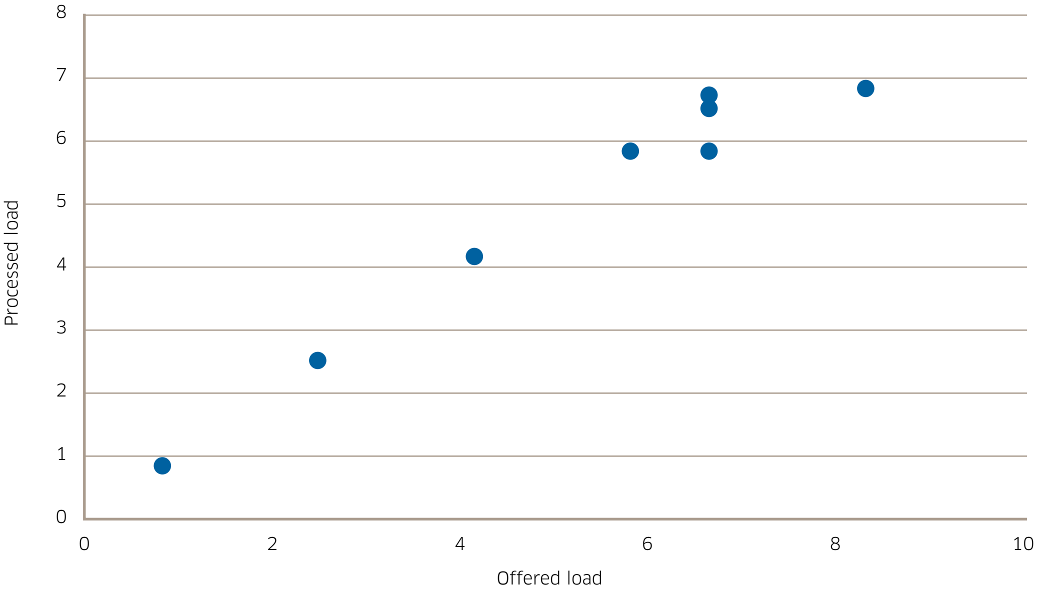

## CONCLUSION

LTE networks do not have the luxury of re-purposing general purpose servers to throw extra capacity at a network failure in an effort to return to a steady state from an overload created by an attach storm. To eliminate potential service outages and reduce the impact of service disruptions on end user services, network elements must be able to shed excessive data traffic load gracefully. This lowers the impact on the network and enables it to return to a steady state condition automatically without causing a long term outage.

An effective load shedding capability in elements like the HSS allows the network to mitigate against and recover from events such as PGW failures and their resulting attach storms. The ability to discard messages selectively at the application layer, based upon factors such as CPU and application queues, in addition to discarding of messages, based upon Diameter stack queues, enables effective shedding of load, minimizes the effect of attach storms on the network, and ensures end users continue to receive the reliable data services they expect and demand.

## ACRONYMS

| | |
|---|---|
| AAA | Authentication, Authorization and Accounting |
| ARPU | average revenue per user |
| AuC | Authentication Center |
| EPC | Evolved Packet Core |
| GSM | Global System for Mobile Communications |
| HLR | Home Location Register |
| HSS | Home Subscriber Server |
| LTE | Long Term Evolution |
| MME | Mobility Management Entity |
| NOR | Notification Request |
| PDP | Policy Decision Point |
| PGW | PDN Gateway |
| SDE | service delivery environment |
| SDP | service delivery platform |
| SGW | Serving Gateway |
| UE | user equipment |
| UMTS | Universal Mobile Telecommunications System |

Alcatel·Lucent