1200 G Street, NW
Suite 500
Washington, DC 20005

P: +1 202-628-6380
W: www.atis.org

June 30, 2016

Via Email
3GPP Liaisons Coordinator
3GPPLiaison@etsi.org

Thomas Tovinger, 3GPP SA5 Chairman
thomas.tovinger@ericsson.com

Jean-Michel Cornily, 3GPP SA5 Vice Chairman
jeanmichel.cornily@orange.com

Christian Toche, 3GPP SA5 Vice Chairman
christian.toche@huawei.com

RE: Establish a Metric to Determine a Drop in Registered Users in an IP-Based Network

Dear Messrs. Tovinger, Cornily, and Toche:

The Alliance for Telecommunications Industry Solutions (ATIS) Network Reliability Steering Committee (NRSC) kindly requests 3GPP SA5 to standardize a new performance-based metric that may be used to monitor for service impacting events in the evolved packet core (EPC).

ATIS' NRSC has explored existing industry standards for performance SIP-based metrics that yield information comparable to what is available in time-division multiplexing (TDM) networks; however, NRSC has discovered that no metric exists that would yield data showing what percentage of registered users have lost connectivity or are no longer requesting a reregistration. The need for this metric is two-fold—not only is it critical for all service providers to have insight into their networks when an event has occurred that impacts network availability, service and/or customers, but North American service providers are also mandated under the *Disruption of Communications Title: 47; Part: 4 Rules*[1] to report certain outage events (with specific details) to the Federal Communications Commission (FCC).

NRSC proposes a metric that will monitor a drop in registered users on the network minus successful deregistration multiplied by 100 to yield the percentage of registered users dropped in the network:

[1] *Disruptions to Communications*, 47 CFR Part 4 (§4.9 - *Outage Reporting Requirements—Thresholds Criteria*), 77 FR 25097, Apr. 27, 2012.

$$\textbf{Drop in Registered Users \%} = \frac{(\text{\# of IRAs}^{2} - \text{Total \# of REGISTER Requests } \textit{Polled registered users every 15 minutes}) - [\text{Deregistration (Expires: 0) 200 OK}]}{(\text{\# of Registered Users from HSS/HLR}^{3}) - [\text{Deregistration (Expires: 0) 200 OK}]} \times 100$$

A majority of network impairments come from the last mile connectivity and it would be useful to monitor for a drop in registered users at the access point using the Proxy-Call Session Control Function (P-CSCF) or Interrogating-CSCF (I-CSCF). As well, in order to identify localized events, a core metric is needed to capture widespread or network related outages. Provided in an appendix is one use case with data collection needs for both access and core.

Service providers would use this calculation to set network thresholds alerting the service provider that their network may be experiencing a service impacting outage.

The creation of this metric is of high priority for ATIS NRSC and we would appreciate this request being considered as a maintenance item during the 3GPP July 2016 meeting. Please contact us if you have any questions.

Sincerely,

Andy Gormley, NRSC Co-Chair (andy.gormley@t-mobile.com)
Stacy Hartman, NRSC Co-Chair (stacy.hartman@centurylink.com)

cc: Mirko Cano Soveri, 3GPP SA5 Secretary (mirko.cano@etsi.org)
Mark Peay, NRSC IP Reliability Subcommittee Co-Chair (mark.peay@cox.com)
Chris Oberg, NRSC IP Reliability Subcommittee Co-Chair (chris.oberg@verizonwireless.com)
Tom Goode, ATIS General Counsel (tgoode@atis.org)
Steve Barclay, ATIS Director - Global Standards Development (sbarclay@atis.org)
Jackie Voss, ATIS Manager - Global Standards Development (jvoss@atis.org)
Rebecca Goodman, ATIS Coordinator - Global Standards Development (rgoodman@atis.org)

[2] IETF RFC 6067, *Basic Telephony SIP End-to-End Performance Metrics*, Section 4.2, Ineffective Registration Attempts (IRAs).
[3] Home Subscriber Server/Home Location Register.

Appendix A

Outage Alerting and Reporting Use Case

Description & Business Value

As the Public Switching Telephone Network (PSTN) and wireless networks transitions to all-IP communications, many aspects of the way networks are managed must also change, presenting new challenges. One such challenge is how regulatory-relevant outages are measured and reported in IP networks and how faults that cause reportable outages can be identified for reporting and restoration activity.

On February 15, 2012, the FCC approved a *Report & Order*[4] expanding the mandated outage reporting rules to include Voice over IP (VoIP) telephony services, effective December 17, 2012. The new VoIP requirements are similar to current voice regulations for circuit-switched telephony. Specifically, the following outage conditions require reporting:

- Telephony user outage condition: ≥30 minutes duration and 900,000 user-minutes.
- E911 rules redefined (≥30 minutes duration and ≥900,000 user-minutes and (a) PSAP loss of communication or (b) E911 loss of call processing (tandem, router, switch, etc.) or (c) loss of Automatic Location Identification (ANI)/Automatic Number Identification (ALI).
- Mass blocked call events that result in ≥90,000 real-time blocked calls and last ≥30 minutes.
- Further complicating factors are that both network operators (i.e., wireline, wireless, cable, satellite,) and over-the-top providers are required to report VoIP outages, and the FCC requires localization information for outages to the city state level.

In the PSTN, discrete alarms for voice switches, voice lines, and trunks of deterministic voice call capacity makes the determination of the number of lines impacted by any given fault (e.g., a switch, port, card failure, or TDM trunk failure) fairly straight forward in identifying the location and scope of an outage. By contrast, the meshed IP networks afford resiliency in the form of reroute and convergence to self-heal in many instances. Purely monitoring critical equipment alarms is no longer a valid approach in determining an outage condition. These networks are also converged service networks where voice traffic typically represents a small proportion of the aggregate traffic through any given link or switch, and the throughput per voice connection varies significantly and continuously over time. In addition, IP congestion control mechanisms, depending on how the network is engineered, may not fully restore impacted voice connections. So the question arises: how to achieve the level of visibility and control needed to identify, alert, and accurately measure network outages in IP networks?

The goal of this use case is to improve outage-causing fault visibility.

Variations

There are multiple levels of scope and sophistication of this use case that can be contemplated:

- *Automated outage threshold detection* – Tracking locations of faults, calculating the total number of users impacted, and tracking the duration of those faults in order to identify when thresholds have been crossed for an outage requiring reporting, as well as aggregating the total user voice minutes impacted by an outage.

[4] *Extension of the Commission's Rules Regarding Outage Reporting to Interconnected Voice over Internet Protocol Service Providers and Broadband Internet Service Providers*, 77 Fed Reg. 25088, Apr. 27, 2012.

- *Automated early outage warning* – An early warning version of the base variation, which would alert network operations staff of potential network faults impacting customers and reducing network availability. These early warning indicator would allow service providers to better priorities limited break fix resources.

- *Automated network repair/mitigation* – Taking the previous variation to a significantly higher level, it may be possible under certain types of network faults to employ automated instantiation of new virtual network functions and/or rearrangement of network connectivity in order to eliminate or reduce the outage impact of a network fault. For example, if a large scale power or connectivity failure made local Media Gateway (MGW) or Session Border Controller (SBC) elements inaccessible, new virtual MGW (vMGW) or virtual SBC (vSBC) capacity could be brought online quickly at another data center location and virtual private networks (VPN) could be reconfigured under automated software-defined network (SDN) control to route VoIP traffic to that replacement virtual network function.

Input Data

The types of inputs that need to be collected and analyzed to determine the magnitude of user-minute outage impacts due to network faults falls into the following categories are:

- *Fault alarms & metadata* – Collection of primary (point of fault) Network Element (NE) and link alarm information (both general and E911 related), localizing the elements and interfaces impacted.

- *Outage localization* – If not available directly from alarm information correlation, localization could be approximated by comparing current vs. historical voice traffic patterns to estimate the geographic reach of an outage.

- *SIP registration statistics* – Total number of expected VoIP Registrations. Registration failures at various points in the network and graceful deregistration for wireless networks identify outage conditions based on users affected.

- *Network VoIP statistics* – Tracking of the total numbers of voice lines on a network (i.e., assigned active and inactive numbers)at various points would aid in determining outage impact for events where access to such data is unavailable. The FCC's outage reporting systems[5] require that total VoIP Subscribers Served vs Total Subscribers Down be reported by county served.

Output Action & Target

Once a network failure impacting voice traffic is detected and the magnitude and extent of the voice traffic impact is assessed by the analytics systems, a number of alerting and reporting actions need to take place. Examples are outlined below:

- *Voice-impacting fault alerts* – Operations alert highlighting a critical network fault(s) with customer impacts and a threshold alarm based on outage duration and impacted VoIP lines, so that service restoration can be prioritized and outage impact reviewed against reporting and escalation procedures.

- *Reportable outage threshold crossing alert* – Operations alert indicating that network fault(s) have resulted in a reportable E911 outage, so that service restoration can be expedited and collection of outage reporting information can be triggered.

[5] FCC Network Information Reporting System (NORS) and Disaster Information Reporting System (DIRS) < http://transition.fcc.gov/pshs/services/cip >.

- *Outage summary report* – An automated report generated to summarize the information needed for an outage report, either after services have been restored or ongoing "in-progress" reporting of cumulative outage statistics.

Data Correlation Needs

Assessing the extent and localization of outages will require extensive correlation of traffic data and network information culled from a variety of sources across the network, and may vary significantly depending on the level of information already correlated in operational support systems (OSS) employed in any given network. The following are examples of the types of data correlation needed to drive this use case:

- *Timestamping* – Accurate timestamps for all network alarms (on and off), IP traffic monitoring samples, and session traffic statistics.

- *VoIP address localization* – Mapping of VoIP traffic (or call attempt) IP addresses to data plane control point (SBCs) and/or user endpoint locations.

- *SIP registration statistics* – The number of registered VoIP uses should be compared against a known data set of billable users. The registration statistic will need to capture the number of failed registration attempts and gracefully deregistration requests for wireless networks. Registration data needs to retain location information to identify specific location information to aid in potentially identifying to a city state level a drop in registered users. Denoting a drop in registered user equipment (UE) at both the core and network's edge, various points in the network can then be used to identify outage conditions based on users affected.

- *Alarm correlation* – Reduction of multiple related alarms to primary (causal) failures, so that each network failure is only counted once for outage reporting purposes. Also correlation of node and link alarms to before and after traffic monitoring data (i.e., correlation to the failure group size – card, node, link, etc.).

Standards Gaps:

- Telephony User Outage Condition: ≥30 minutes duration and 900,000 user-minutes.

Gap: The IETF RFC 6076 IRA metric monitors failed registration attempts. However, in most failure scenarios the UE cannot reach the network as a result of last mile access issues, so the request never reaches the registrar.

Metric #1: A standard that shall monitor a drop in registered users at the P-CSCF or I-CSCF minus successful deregistration multiplied by 100 to yield the percentage of registered users dropped at the network's ingress.

$$\text{Drop in Registered Users \%} = \frac{(\text{\# of IRAs} - \text{Total \# of REGISTER Requests } \textit{Polled registered users every 15 minutes}) - [\text{Deregistration (Expires: 0) 200 OK}]}{(\text{\# of Registered Users from HSS/HLR}) - [\text{Deregistration (Expires: 0) 200 OK}]} \times 100$$

This metric would monitor the network's ingress for last mile outages and those only affecting localized areas.

Metric #2: A standard that shall monitor a drop in registered users at the network's core minus successful deregistration multiplied by 100 to yield the percentage of registered users dropped on the network.

$$\text{Drop in Registered Users \%} = \frac{(\text{\# of IRAs} - \text{Total \# of REGISTER Requests}) - [\text{Deregistration (Expires: 0) 200 OK}]}{(\text{\# of Registered Users from HSS/HLR}) - [\text{Deregistration (Expires: 0) 200 OK}]} \times 100$$

This metric would monitor for more systemic issues across the network.

A correlation between the drop in registered users at the network's core and ingress may need to be further understood to avoid duplicate alarming on the same network event.