SAGE ALERTING SYSTEMS

June 16, 2016

Via Electronic Filing

Marlene Dortch, Esq.
Secretary
Federal Communications Commission
445 12th Street S.W.
Washington, D.C. 20554

RE: Notice of Ex Parte Communication, PS Docket No. 16-32, Public Safety and Homeland Security Bureau Seeks Comments on Ways to Facilitate Earthquake-Related Emergency Alerts

Dear Ms. Dortch:

On June 13, 2016, Harold Price of Sage Alerting Systems participated in a conference call with Steven Carpenter, Yoon Chang, Gregory Cooke, Behzad Ghaffari, David Munson, Rasoul Safavian, and James Wiley of the Public Safety and Homeland Security Bureau to provide them additional information on the feasibility of disseminating earthquake warning alerts through IPAWS to the public within 3 seconds.

I have provided the list of question and answers that Sage discussed during the call below.

It addition to the answers to the Bureau's questions, Sage provided a general summary. The IPAWS-OPEN system, including origination, dissemination to EAS devices, and the generation of an EAS message by those devices, was not designed to provide three-second end to end delivery. We can discuss incremental improvements of each step of the process, but what is needed is a system built for rapid dissemination of earthquake messages. Such a system can be based on existing standards such as publish/subscribe, or websockets. A system optimized for rapid delivery can be built using some of the components of the existing IPAWS ecosystem, and in particular could use the existing hardware deployed at EAS participants.

The EAS protocol itself, with its 15 seconds of headers and audience attention tone, is not suitable for an event that is only 30 seconds in the future. The EAS expertise acquired at FEMA, the EAS equipment manufacturers, the EAS participants, and at the Public Safety and Homeland Security Bureau can be used to design, develop, and deploy a suitable earthquake warning system that need not start from a blank sheet of paper, and would reduce costs by taking advantage of hardware already deployed at EAS participants. But EAS as currently deployed is not suitable.

Various delays in audio/video delivery via broadcast and cable will also delay presentation of the alert to audience members. Many of these are downstream of EAS devices and are beyond the control of the IPAWS system.

Questions on IPAWS and EAS Timing

Questions on polling the IPAWS ATOM feed:	**Sage Response**
How long does it generally take to retrieve CAP alerts using Http, Https, SSH, and IPAWS?	As much as 4 seconds, to go through all phases of the interaction with IPAWS and signature verification.
Are there any specifications as to the polling interval for the IPAWS ATOM feed?	Not to my knowledge. IPAWS informally suggested 30 seconds in 2012.
Can EAS equipment manufacturers simply choose an interval they like for whatever reason and set their EAS device models to poll accordingly?	For the ENDEC, it is a user settable parameter in seconds, with a default of 30 seconds.
From the perspective of IPAWS, are there any problems with having all EAS Participants pol at a very short interval, say 100 milliseconds? For example, if every EAS Participant polled at the same rate and sought to retrieve the EEW CAP alert at basically the same moment in time, would that surge in data requests cause problems for the IPAWS servers?	You would need to ask IPAWS about that. We are not aware of any performance issues under the current use model, but there certainly could be if you increase the load by a factor of 300. Some metrics: Assume the server sends about 800 bytes of data in response to each poll– this includes the SSL overhead and the IPAWS payload data. Assume one poll per 20,000 users each 30 seconds. (800 x 20,000)/30 = 533,333 bytes per second = 4.07 Mbps load at the server. Polling at 1 second is 122 Mbps Poll at 0.1 second is 1,220 Mbps This is just for one direction, of course, about the same amount of data goes in the other direction for a poll. As the load on servers is also transaction based, we've gone from 666 transactions per second at 30 seconds per poll to 20,000 at one second, and 200,000 at 100ms. That's an aggregate of all servers, load balancing reduces the load per server. When it comes time to actually fetch the alert, assuming 7k, and if you wanted to deliver it to everyone within three seconds, it depends on how many users actually need the alert. The only data I have is that an RMT to Washington state a few days ago was picked up by 663 devices. 7k x 663 is 38 MBits, you'd need a bit rate of more than 1/3 of that, about 13 Mbit to serve all the users in 3 seconds.

Is it feasible or desirable for all EAS Participants to poll in synchronization or do they have to be staggered to some extent to prevent overloading of the IPAWS ATOM feed?	Spreading out is better, though there is no mechanism to make that happen other than the basic randomization that occurs across the internet.
From the perspective of EAS device manufacturers and EAS Participants, what is the minimum polling period that EAS devices could meet, and could this polling be synchronized on some external clock? What impact, if any, would a near-constant polling interval have on the EAS Participant's IP connection? For example, if the EAS Participant had only one IP service it used for all of its Internet needs, would the constant polling represent any material issue?	Polling at a constant rate would also affect the CPU time available on the EAS device to perform other tasks, and the rate at which non-IPAS servers could be polled. From the EAS participants point of view, at the 30 second polling rate they use 100 Mbytes/month against their total transfer limit, at 1 second it is 1.9 GByte per month and at 100ms, it is 19.3 GByte/month. Multiply by two if the ISP also charges for upload. With some sort of Pub/Sub or HTML 5 Websockets, there is no polling at all. Only a keep-alive ping/pong is used, but it uses the existing connection and is only a few bytes, as opposed to 800. There are many technical details to be worked out, but the basic technique is now well known. It was not widely available when IPAWS-OPEN was first released.
As a purely technical matter, would a "push" or "publish/subscribe" mechanism provide for faster relay of the CAP alert from IPAWS to EAS Participants deployed EAS devices than using the IPAWS ATOM feed? Could IPAWS and deployed EAS devices support such mechanisms?	It would be faster, and would use less CPU and bandwidth resources. EAS devices that are already SSL capable should be able to support such a scheme. On the server side, some things are harder, such as load balancing, with many other problems to solve. But this (using a pegged-up connection) is the only way to remove polling from the equation. Note also that a publish/subscribe or a websockets design will still permit EAS participants to connect outbound to FEMA and would not require a hole in the firewall to allow FEMA to initiate a connection, or require FEMA to keep an up to date list of IP addresses to connect to.

Questions on EAS Participants' Internet connection to IPAWS:	**Sage Response**
Assuming that the CAP alert is ready for retrieval from the IPAWS ATOM feed, is the actual time to retrieve that alert impacted by the Internet service used by the EAS Participant? For example, EAS Participants are basically required to have a broadband connection to process IPAWS CAP alerts, but there is no definition as to a minimum level of throughput that would qualify as broadband connectivity for these purposes. It is likely that some EAS Participants have very poor connectivity, with fairly low throughput. CAP alerts are short, but would any aspect of the IP connection have a material impact on how long it takes to retrieve the alert from IPAWS?	All aspects of the IP connection matter. EAS Participants access IPAWS over the internet using commodity connections, not a carefully controlled in-house LAN. Dropped packets, and the resulting delay due to TCP retry backoff, can be more painful than a slow link. Ignoring embedded audio, we're talking about the delivery time of a small number of packets. This isn't something that can be concealed by use of a jitter buffer or forward error correction. CAP alerts are not short if they contain embedded audio, as one-way satellite messages would be. A CAP message with audio could be as much as 960k bytes for an alert with two minutes of audio. Also, a big message that is being downloaded could delay receipt of a smaller message. Typical text-only alerts are about 7k bytes. An alert with polygons could be larger, as would alerts with multiple languages. Yes, some EAS Participants have low bandwidth connections, and some use Wi-Fi in a shared environment. If audio is included, transfer time is an issue. Otherwise, for a single small alert, throughput is less of an issue, but shared connections, in particular Wi-Fi, could be an issue.
Questions on how long it takes EAS devices to process CAP alerts from input to output:	**Sage Response**
Once the CAP alert is fully acquired in the EAS device, how long does it take for the device to process the alert and output the data stream?	For the Sage ENDEC, less than two seconds. That is the time required to start the sequence. If standard EAS message formats are used, the header portion lasts about seven seconds, the attention signal lasts 8. The start of actionable audio to the audience is thereby delayed by 15 seconds – not counting downstream delays. The text information via a video crawl would start sooner. Each additional location code in the EAS message would add 276 ms to the header portion. (15.36 ms per character x 6 characters per location * 3 headers).
Is this processing time impacted by whether there is an audio file and/or enhanced text included?	No for text, yes for audio file – the ENDEC currently downloads the entire audio file before starting the alert – to avoid audio pauses due to data delivery pauses. There is no FCC-approved method of including enhanced text or audio files to legacy EAS messages.

Is this processing time impacted by whether the device has to encode the alert for the benefit of downstream monitoring stations (as opposed to simply decoding the alert for local consumption)?	No – with EAS, the alert is the same in both cases. Unless you are asking about the time to send the EAS header bursts – that adds a minimum of seven seconds plus attention tone.
Do your EAS device models allow EAS Participants to program the timing of when state and local EAS alerts are broadcast (e.g., setting the device so that it waits 15 minutes before broadcasting a state / local alert)? If so, (i) what, to your knowledge, is the typical setting used by EAS Participants that have deployed these models, and (b) can such feature be turned off for selected EAS event codes?	Yes, we allow users to program the timing, in minutes. Delay is on a filter by filter basis, the user has complete control over delay times.
Could requiring a given state / local event code to be aired "immediately" – i.e., just like an EAN, but secondary in priority to an EAN – be programmed into your currently deployed and marketed EAS device models?	Yes, this is a feature that has been a part of all Sage ENDECs since 1996 – including the ability to start a legacy alert while it is still being received.
Questions on how long it takes for the EAS data stream (including audio, if present) to be broadcast:	**Sage Response**
Once the EAS device has converted the CAP alert to SAME and output a data stream (including audio, if present), how long does it take for that data to actually be broadcast as a visual crawl and audio (if present)? Is any of that timing impacted by station equipment upstream from the EAS device? Is this answer different for broadcast versus cable and satellite? For cable operators, would this time period change depending upon whether the cable operator overlayed EAS onto each channel at the headend or ran the EAS on a control channel to which subscribers were force tuned?	Devices after the EAS device (downstream) do control how long it takes for a member of the audience to receive the alert. Encoding/compression delay, profanity filters, HF radio encoder, etc., will add delays. HD radio typically adds at least seven seconds, sometimes more, to both analog and digital listeners. Some stations place the profanity filter downstream of the ENDEC, which can add 15 or more seconds.

Respectfully submitted,

/s/

Harold Price
President
Sage Alerting Systems, Inc.

cc (via email) Greg Cooke
David Munson
James Wiley
Rasoul Safavian
Behzad Ghaffari
Steven Carpenter
Yoon Chang