ELECTRIC POWER
RESEARCH INSTITUTE

# Improving Electric Power System Situational Awareness

## Leveraging "Information Bursts" and an Internet Message Manager

*November 2015*

**Data Analytics for Utilities**

- *Moving more and more data is untenable. Moving just the actionable information is the goal.*
- *Managing and obtaining insights from millions of inputs per second can be accomplished with Internet-scale technologies.*
- *As of August 2015, over 500 million tweets are sent per day. Applications are readily available that enable analysis and visualization of this massive information volume.*

**Asset Optimization**

**Load and DER Awareness**

**Outage Management**

✓ **Practice and Technology**

**System Optimization**





# Overview

*Big Data Is Getting Bigger. Let's Do More with Less.*

Electric service providers are installing and evaluating many new sensor technologies, conferring a greater end-to-end awareness and visualization of their power systems. Experts agree that in the near future, the majority of the electric power grid will have vast, integrated communication networks and intelligent monitoring and sensing capabilities. The tentacles of the new grid will touch virtually all points within the electricity business and infrastructure, from transmission assets to field crews—even certain types of customer equipment. How will utilities handle the constellation of data points that convey information vital to understanding and managing their modern grids? The authors of this concept paper propose a new paradigm using "information bursts" to visualize and to manage massive amounts of system status messages.

The idea is to leverage distributed computing to substantially preprocess sensor data into compacted and actionable intelligence packets, or "information bursts," before that sensor data ever moves from its collection point. By packaging data into information bursts—similar to the way texts, tweets, and photos are transmitted from a smart phone—a standardized message could be transmitted from any point on the power system and handled with methods and tools similar to those used by Internet-scale companies such as Google and Twitter, which manage and visually represent tens of millions of simultaneous texts and tweets per minute. Use cases presented in this paper illustrate the tangible benefits of harnessing these messages. Assuming that the optimal technologies are identified, electric service providers would ultimately employ the systems in secure private networks and finally achieve optimal "right time" visibility for operation and management activities.

# A Sensor- and Data-Enabled Power System

*Vet the Bytes before Sending.*

As more and more of the electric power grid becomes rich with sensors and communication systems, electricity providers are faced with challenges related to managing and leveraging an unprecedented flood tide of data. The most conspicuous of the new data sources is the smart meter. However, nearly all modern utility assets—including voltage regulators, transformers, switching devices, and even the modest wooden pole—have or will have options, or add-ons, for monitoring electricity and various operational status parameters.

Consider the trend of increasing sensor deployment and the abundance of external data feeds such as video, weather, social media, and communication-enabled customer devices. The challenge is clear. How will electric service providers manage and integrate all of that data? Existing approaches to transfer, warehouse, and analyze such a torrent of data may prove inadequate to the challenge. Categorizing the challenge to overcome deficiencies in data management as a bandwidth and storage problem curtails innovation when innovation is needed most.

This paper proposes that *it is not necessary to transmit more and more data in order to achieve the desired levels of system visibility and situational awareness. Instead of more data, the goal should be to transmit only the necessary data. The method proposed here is one of moving just the right messages—information bursts—to the appropriate points where that information can be leveraged most efficiently.* Internet-scale technologies employed by "big data" companies like Google and Twitter—which are already being deployed for smart metering systems by some utilities—may be the next indispensable technology for the electric utility industry.

The supplemental message layer proposed in this paper would not impact any of the existing data-collection and analytics processes. The idea is that the new layer would provide an option for "on demand" situational awareness activities, while the existing and legacy systems like SCADA and AMI would continue to provide their intended functionality and redundancy.

# What Is an Information Burst?

*Thrifty Preprocessing Reduces Data Burden for Electric Utilities*

An "information burst," as the name suggests, is a message containing actionable information that automatically transmits (bursts) from a remote location without the need for a user to poll the device for the message. Information bursts can range in type from simple status messages to more detailed incident reports. For example, a device may be reporting just a heart-



beat (that it is aware and able to communicate), or it may be transmitting a parameter such as voltage or a temperature on an hourly basis. The sensor or device may also burst an unplanned message about a change in status or an "out of limit" condition.

To understand how the concept of the information burst can be used and applied, consider the mapping application on a smart phone. Without even being aware of it, every driver with a smart phone is continuously transmitting information bursts back to the cell towers that their phones are linked with. Now with just a few simple distance and triangulation algorithms, the mapping application is able to create traffic-congestion maps by knowing the speed and position of each phone sensor (the drivers). With this traffic layer, the roadways on the map can be color-coded to show red areas where traffic is stopped or backed up and green areas where there is no congestion. Additional algorithms can even recommend to drivers which routes to take for optimal time efficiency based on the congestion analysis.

To understand how information bursts would work for the electric power system and all of its sensors and hardware, consider how information is transmitted from customer smart meters. A smart meter at a customer location captures all of the customer's usage data but sends only an aggregate of the usage over a given time period and will additionally send an outage notification or a "power restored" notification in the event of a power interruption. A sensor employing an information-burst system will do likewise, sending only information contained in a very compact message packet with date and time stamps, and there is never a need to poll or to query the sensor for a message, which is pushed as soon as the message is ready. The information bursts from each power system sensor and each asset can be different, so a transformer might just send one hourly burst while another asset may send a burst once every few seconds.

## How Data Is Preprocessed into Information Bursts

### *An Information Burst Is a Small Slice of the Data Pie*

As represented by the blue area in the graphic to the right, the current process flow from data capture to useable information requires transmission, storage, and processing of full datasets—every bit of sensor data must be transmitted before any analytics take place in order to create actionable insights. In the near future, any utility still applying this method will bump against an ever-increasing volume of data and a static bandwidth—trying to push more data through the same pipe.

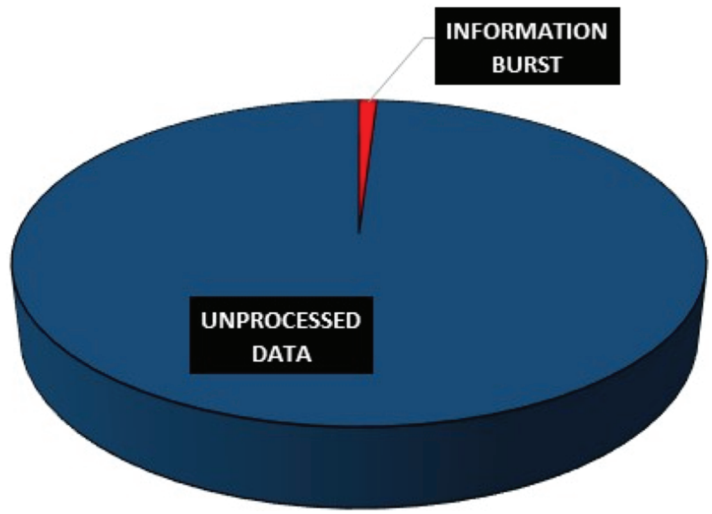


*An information burst consists of only actionable information processed from raw data from a sensor.*



On the other hand, preprocessing data into the much smaller information bursts significantly minimizes the volume of data that must be transmitted. In essence, information bursts are an important compression and bandwidth-conservation technique that minimizes the amount of data that has to move across a network. As depicted in the red slice of the figure, several megabytes of raw sensor data can be distilled down to just a few kilobytes of actionable information before that information ever moves from its collection point.

With minor modifications, much of the current utility hardware and sensors could be used to accomplish the proposed method of distributed data processing. The challenge for electric utilities is dealing with the paradigm shift from central to distributed processing—in short, creating an "information burst infrastructure" that centers around programming, software customization, and standardization.

Internet-scale data managers such as Twitter already employ this new paradigm of high-speed distributed processing. In fact, as of August 2015, there are on average over 500 million tweets sent per day. Applications are readily available that enable analysis and visualization of this massive information volume. A similar type of data manager designed for the utility industry must be



capable of managing and processing millions of information bursts in a matter of seconds and be able to visually display the outcomes in sub-minute timeframes.

Consider a moderately sized city and surrounding community like Chattanooga, Tennessee. The number of sensor-enabled devices would be estimated at double the number of customers, or about 1 million communicating sensors, smart meters and assets. At a rate of one information burst per minute, in just eight hours, this moderately sized utility would have transmitted the same number of tweets that Twitter manages in a full day. The key difference between a public Twitter system and a private utility version is the hardware devices themselves—smart phones versus sensors, distributed generators, and other utility assets.

## Proposal to Demonstrate Information Bursts and an Internet-Scale Message Manager

### *Distributed Intelligence Will Reduce Data Flow to a Manageble Level*

With a combination of distributed intelligence at each sensor node to preprocess data into smaller packets of information containing only the important insights and a private version of a message-management system like Twitter, it would be possible to achieve close to real-time situational awareness of an entire power system. Internet-scale message mangers are not used today for electric power systems, but all of the pieces are available to demonstrate such a system. EPRI is working with the University of North Carolina at Charlotte to develop a demonstration plan and a prototype of the message manager and the message burst data structure. Because this work is sponsored by EPRI's Distribution Modernization Demonstration, the initial focus will be on use cases that address objectives related to distribution systems. The project will investigate data-management techniques similar to those currently used to analyze tweets from Twitter accounts and data from other social media. The data analytics and interactive visualization methods used by such Internet-scale companies will be applied to the specific use cases.

The demonstration will address the following areas:

1. Conceptualization document. This is the white paper you are reading.
2. Computing requirements (at the device or sensor) necessary to accomplish distributed data analysis.
3. Development of the device and asset-specific information bursts (the message sets).
4. The process by which the remote devices transmit their information bursts.
5. Development of use cases (the most relevant and valuable use cases that crystalize the concept of the information burst).
6. Value quantification. How is the information used to manage and optimize the grid?
7. Information management. Where and how is the aggregated information used or collected and visualized?
8. Field demonstrations.

It is beyond the scope of this paper to completely detail all of the demonstration activities in the previous list, but there is some value in discussing the topic of how the information burst would be converted into actionable intelligence by way of an aggregation and visualization interface. The idea behind the visualization tool would be to provide every user with a customizable dashboard to enable a more personalized interaction with the interface.

Because the idea is to simulate an Internet-scale message manager, the system would require a data-management interface and suitable database structure to manage all data sets of interest from streamed information bursts to enrichment data sets such as images or weather information. There is a wide range of visual analysis options that can be leveraged to support streaming analytics, historical trending, predictive analytics, and machine learning tools.

The figure to right provides an example of the kinds of visual tools that could be set up to drag and drop onto a given user's customizable dashboard. This would be analogous to the way that a "homepage" can be customized for each user.



# Use Case for Information Bursts: Rooftop Solar Installations

## *Packaging Data from Solar Installations Reduces Data Flow from Megabytes to Bytes*

The increasing density of rooftop solar panels in today's electrical grid serves as a useful example of how a layer of situational intelligence might handle message information. Consider a single distribution circuit with hundreds of individual rooftop solar installations. Ideally, we would want to be able to understand the exact status of each solar installation at any given time. More specifically, the kind of situational awareness that would be useful might include whether any given unit is offline, whether it is generating at full capacity, or whether the unit is underperforming due to shadows or dirt buildup, and so on. In terms of situational awareness, we would additionally want to know how each solar unit helps to support the power system's steady-state and dynamic load and generation. Imagine the analytical burden of downloading megabytes of data from each solar installation and adding this extra data stream to the current system analytics.

Now consider a scenario leveraging information bursts. Our hypothetical circuit—now replete with sensors—is capable of communicating data about voltage and power flow from end to end. And at each solar installation, instead of megabytes of data, each solar node incorporates a properly designed and configurable set of information bursts that can be transmitted as needed. The information bursts from all of the solar installations would be aggregated at a master interface where a forecasting application could achieve a much more accurate "day ahead" solar generation forecast. Similarly, a load flow application could attain useful insights into forward and reverse power flows throughout the day.

Information bursts could also be applied at the circuit level. For example, if clouds pass over some of the solar installations, reducing their individual generation, the voltage regulator upstream could be immediately aware and change taps accordingly. Essentially, all devices that were subscribed to the messages sent by the solar nodes could participate in this society of intelligent devices by receiving the information burst, responding accordingly, and perhaps contributing to the community conversation.

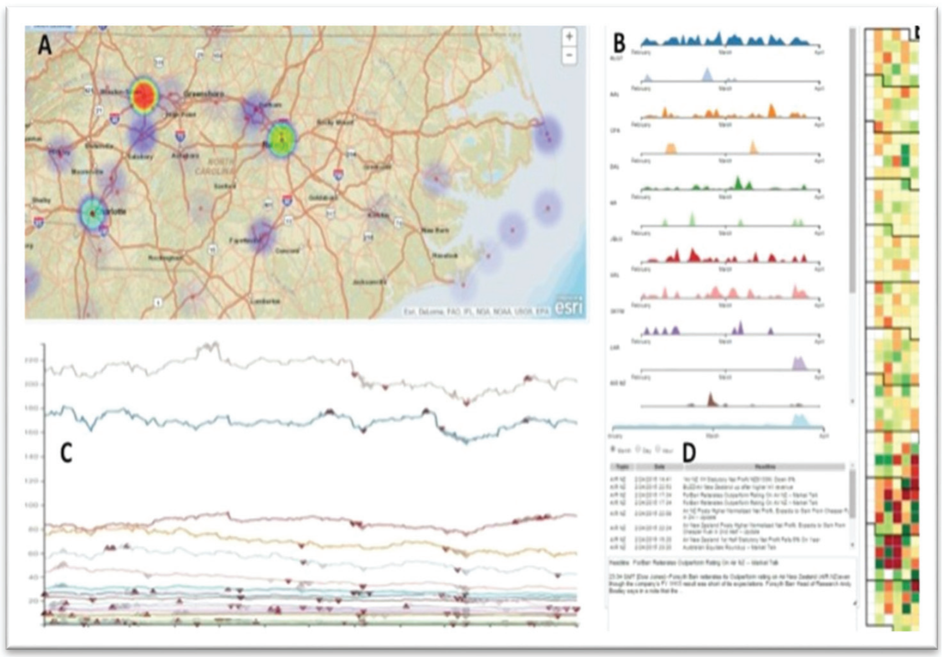


*Information Bursts—Transmitted Only When Data Is Needed—Can Be Aggregated and Viewed in a Master Interface to Achieve More Accurate Day-Ahead Forecasts of Solar Generation.*

# Data Tsunami or a More Visible Power System?

## *A Flood of Data May Conceal Vital Information.*

Improving smart grid situational intelligence as it relates to grid automation, flexibility, interactivity, and predictability will certainly be beneficial for society. The challenge is that achieving such visibility and insight requires a significant amount of *information* from every part of the system. Many data-analytics experts agree that current methods used in the power industry for data management are untenable. How will electric utilities manage vast sensor data with the constraints of a finite bandwidth and a competitive economy? The initial reaction is to figure out how to collect, manage, and store this data, but the



industry needs a completely different mindset on sensors and data collection. In other words, what is the best way to move just the bytes of valuable information instead of terabytes of unprocessed data?

The new paradigm of distributed data processing ensures that only the key information is transmitted, eliminating many present-day data bottlenecks. A few examples of the many applications of information bursting are:

- Power monitors and smart meters could send notifications that could easily assist in identifying the exact location of a power system fault.
- Assets could "tweet" or send text messages containing their health status or maintenance needs.
- Video cameras on vehicles (or cameras at stationary grid locations) could send message coordinates where a vegetation fault is imminent.
- Weather stations could send "change in intensity" updates directly to visualization tools to reduce the inherent forecast uncertainties, as opposed to just using data for the next time delayed forecast.

## Conclusions

EPRI fully expects that the cost of sensors and communication hardware and software will significantly decline and that data-processing hurdles will be overcome. When the cost of sensor deployment becomes minimal in the near future, the majority of the electric power grid will have vast, integrated communication networks and intelligent monitoring and sensing capabilities. This paper proposes a methodology and demonstration to assess the efficacy and feasibility of leveraging information bursts to visualize and to manage hundreds of millions of system status messages. The idea is to employ distributed computing to substantially preprocess sensor data into compacted and actionable intelligence packets, or "information bursts," before that sensor data ever moves from its collection point. By packaging data into information bursts—similar to the way texts, tweets, and photos are transmitted from a smart phone—a standardized message could be transmitted from any point on the power system and handled with methods and tools similar to those used by Twitter. A technology demonstration of the concept would focus on showing how any electric service provider might employ the technology within a secure private network and finally achieve optimal "right time" visibility for operation and management activities.

# What Is the EPRI DMD?

Actionable information derived from data and analytics is an important objective for the electric power grid of the future. This smarter grid of the future will have five attributes—enabled by information and communications. Each attribute will play a role in the ability of a utility to deliver energy to customers in the most efficient, cost-effective, and environmentally responsible means possible. These five attributes describe a grid that is more automated, more flexible, more intelligent, more interactive, and more predictive compared to today's power-delivery system.

One of the best ways to understand the role and the value of analytics as they relate to these smarter grid attributes is to accomplish demonstrations of technologies and the associated research concepts that benefit one or more of the attributes. To that end, EPRI's Distribution Modernization Demonstration (DMD) focuses on distribution-specific applications, supporting infrastructure, and demonstrations of data-related innovations around the world. The project has three key focus areas:

1. ***Use Cases***
   Documenting valuable and innovative data-driven use cases that leverage distribution sensors and enrichment data sources.
2. ***Roadmaps***
   Supporting utility members with shared learnings, leading practices, and value-assessment tools as they deploy analytics technologies.
3. ***Faster Insights***
   Providing members with informational materials and demonstrations that enable a better understanding of emerging analytics opportunities.

This document series, called *Data Analytics for Utilities*, is designed to describe research under the DMD that addresses one or more of the five smart grid attributes. This document series will cover DMD activities where data is leveraged to demonstrate enhancements in Asset Optimization, Load and DER Awareness, Outage Management, Practice and Technology, or System Optimization. Welcome to the second edition, *Analytics Opportunities and Leading Practices for Remote Detection of Energy Theft.*

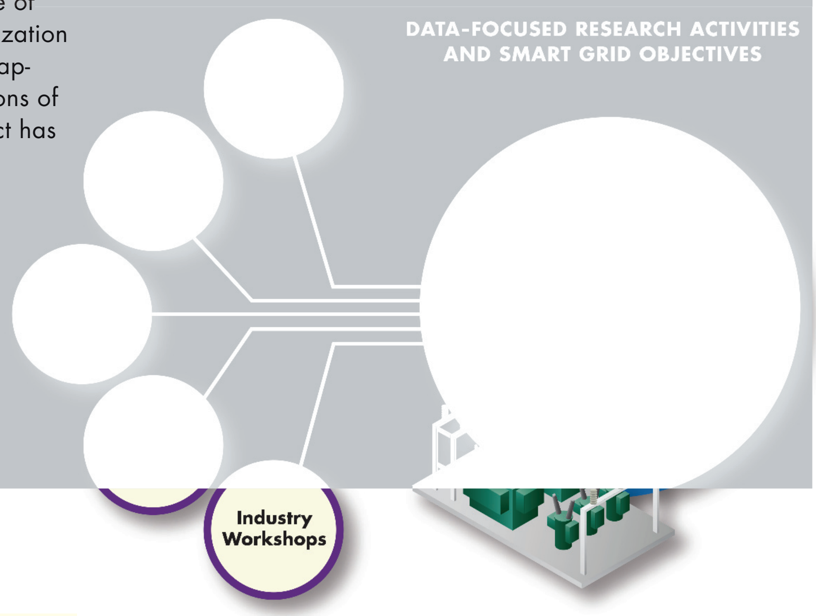


## EPRI Resources

**Doug Dorr**
*EPRI DMD Project Manager, Information & Communication Technology, EPRI, 407.240.5049, ddorr@epri.com*

**For technical questions, contact:**

**Tom Geist**
*Senior Project Manager, Energy Utilization, EPRI, 865.218.8014, tgeist@epri.com*

***Information & Communication Technology***

**The Electric Power Research Institute, Inc.** (EPRI, www.epri.com) conducts research and development relating to the generation, delivery and use of electricity for the benefit of the public. An independent, nonprofit organization, EPRI brings together its scientists and engineers as well as experts from academia and industry to help address challenges in electricity, including reliability, efficiency, health, safety and the environment. EPRI also provides technology, policy and economic analyses to drive long-range research and development planning, and supports research in emerging technologies. EPRI's members represent approximately 90 percent of the electricity generated and delivered in the United States, and international participation extends to more than 30 countries. EPRI's principal offices and laboratories are located in Palo Alto, Calif.; Charlotte, N.C.; Knoxville, Tenn.; and Lenox, Mass.

Together . . . Shaping the Future of Electricity

3002006985 November 2015

**Electric Power Research Institute**
3420 Hillview Avenue, Palo Alto, California 94304-1338 • PO Box 10412, Palo Alto, California 94303-0813 USA
800.313.3774 • 650.855.2121 • askepri@epri.com • www.epri.com