Cybersecurity Aspects of Pub/Sub Communication Architectures (Discussion DRAFT)

November 9, 2018

Operation of the electric system requires precise real-time balancing of power system demand and supply against infrastructure and resource constraints. As grid modernization brings with it the increasing use of digital and computing systems, a growing concern is the possible introduction of vulnerable points throughout the grid network which could be exploited to bring down portions of the electric grid network.

Evolving Smart Grid capabilities, such as advanced metering infrastructure (AMI) and active load management, and increasing reliance on distributed energy resources as substantial contributors to system health, elevates the importance of distribution system automation to grid efficiency and reliability. As greater automation expands reliance on computer networks, data transmissions, and internet connectivity, cybersecurity weaknesses may eventually pose as great a threat to grid reliability as it now does to utilities' and customers' data security.

Security risks within grid communications depend on the ways in which grid components are accessed. One approach is for an information system to uniquely identify and authenticate an organization-defined list of devices before establishing a connection. Devices requiring unique identification and authentication may be defined by type, by specific device, or by a combination of type and device as deemed appropriate by the organization. The smart grid information system must protect the integrity and confidentiality of communicated information. It is feasible to implement these protections at one or more various locations within the communications stack; each placement location carries varying benefits and downsides. Organizations can employ integrity and confidentiality cryptographic mechanisms to prevent unauthorized tampering and disclosure of information during transmission.

Publish and Subscribe Communications

A common communications method is devices publishing and subscribing to message topics, referred to a publish and subscribe (pub/sub) communications. In this communication method, messages are associated with topics, which can have a hierarchical structure (e.g. topic1/topic2/topic/…/topicN). When a device publishes a message, a topic must be specified and the message is published under that topic (e.g. topic1/…/topicX/"Message Text"). Devices subscribe to specific topics they want to receive messages about, so that when a message is published to that topic the devices subscribing to it will receive the published message. Example scenario:

- Device A subscribes to topic devices/Information
- Device B publishes message "I am Device B" to topic devices/Information
- Device A receives message devices/Information/"I am Device B"

Devices can publish and/or subscribe to as many topics that are allowed by implementation specific policies.

Brokered vs. Brokerless Communications

Information flow between smart grid devices can happen in a brokered or brokerless manner. In the brokered model, devices can subscribe to message topics they should receive communications about, and/or publish messages to message topics that devices are subscribed to. In the brokerless model, devices exchange messages directly with other devices; no central service handles how published messages are distributed. Each method of communication must deal with discovery, availability, and management of devices. Device discovery is creating and maintaining a list of which devices can publish or subscribe to messages on specific topics. Device availability deals with how to handle messages intended for devices that do not have continuous connections with each other. Device management is the process of ensuring devices are configured correctly to perform identification, authentication, support secure communications protocols, apply configuration, and remediate vulnerabilities. Brokered and brokerless communications each have different ways of addressing device discovery, availability, and management with varying pros and cons.

Brokered Communications

Each device has a connection established with a common broker, which controls message distribution to and from all of the devices. When a device publishes a message to other devices, the message is first sent to the broker and then the broker transmits the message to all of the devices who are subscribed to receive the message. Within this model, an example of which is shown in **Figure 1**, the broker is a central location that all devices are registered with and all messages flow through.

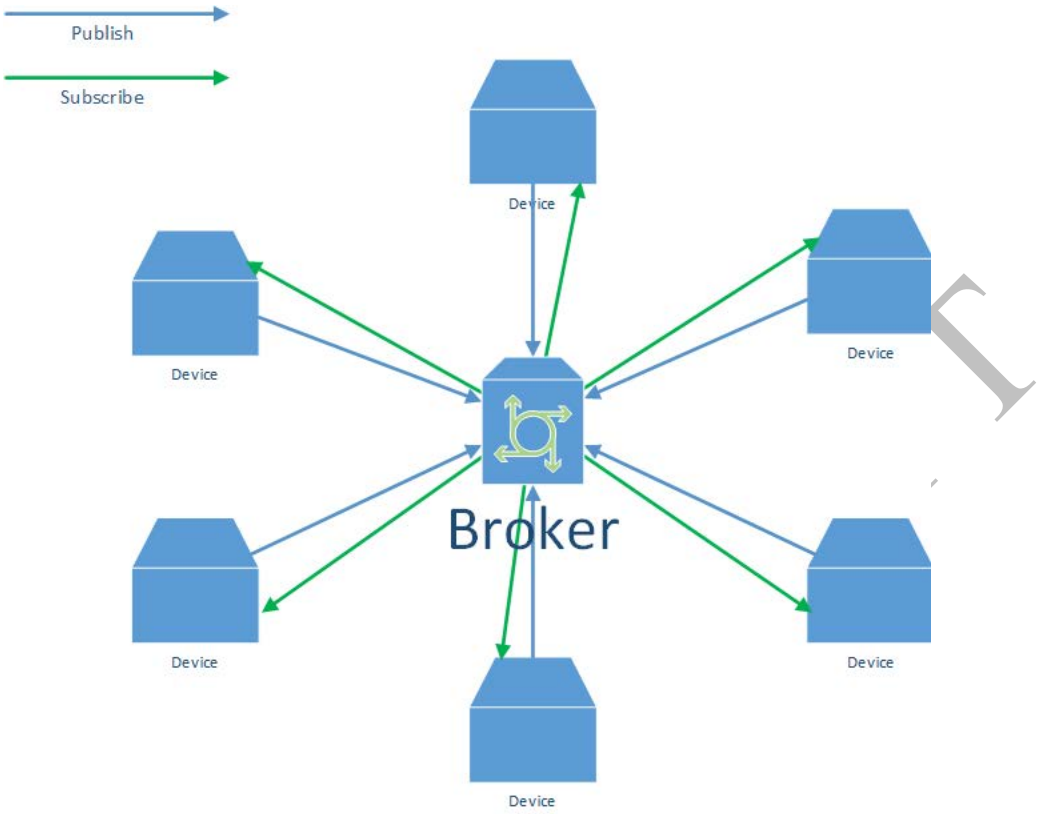

Figure 1 - Brokered Communications Network

The broker maintains device registrations, which stores information about which devices are allowed to publish messages about specific topics, and which devices can subscribe to receive messages about a specific topic. Any particular device could be allowed to publish messages to any relevant topic, as well as subscribe to any relevant topic. The policies that govern which devices can publish and subscribe to topics should be created by the organization operating the network of devices. Device registration can be performed in a variety of ways: MAC address, IP address, Trusted Platform Module (TPM) key attestation, username and password, shared secrets, cryptographic key authentication, etc. Since the broker maintains device registration, the rules for each device to publish and/or subscribe can be stored within a central location for the broker.

The broker also aids device availability by ensuring that all devices receive messages that are in intended for them. The broker can maintain a list of all active and inactive connections that it has with devices. Since the broker also maintains the device registration list, when it publishes a message intended for a set of devices the broker can send the message to devices have an active connection. For the devices that are supposed to receive the message but have an inactive connection, the broker can cache the message for this subset of devices and send it to devices as they reconnect.

Device management can be uniform within a brokered communications environment. Since all devices will be connecting to a common broker, similar device configuration could be required of all devices before the broker grants each device its access. This fact could enable a standardized way for organizations to manage their smart grid devices because all of the devices must meet a common policy to successfully register with the broker.

Performance of devices before and after enabling security controls for establishing connection with the broker. The results in **Table 1** show that performing the TLS handshake is the cause of the increased time of operations.

Table 1 - TLS handshake introduced latencies for establishing broker connections

Authentication / Encryption	None	TLS 1.2
None	.002855	.108071
Username & Password	.003223	.108458
X509 Certificate	N/A	.201232

Performance of devices before and after enabling security controls for publishing a message. The results in **Table 2** show that performing the TLS handshake is the cause of the increased time of operations.

Table 2 - TLS handshake introduced latencies for message publication

Authentication / Encryption	None	TLS 1.2
None	.003242	.109013
Username & Password	.003216	.108034
X509 Certificate	N/A	.186208

Table 3 - Potential Pros & Cons of Brokered Communications

Pros	Cons
Reduces computations on devices	Centralized broker adds extra hops/time for messages to be sent and received
Device management can be standardized	Difficulty in getting new/specialized devices registered
Broker can ensure devices receive all intended messages	Single point of failure if broker goes down
Logging and aggregation of all messages	Device registry with access rules could become very complex

Brokerless Communications

Each device has a mechanism to publish its own messages while simultaneously listening for messages from other devices. Each device is responsible for how it publishes its messages. Brokerless communications could have different types of implementations; e.g. a publishing device could broadcast its published message and rely on subscribers' ability to receive the message, or conversely the publishing device could send individual messages to all of the devices subscribing to its messages. Devices are also responsible for handling the message subscriptions they have in place. The subscribing devices must monitor for relevant messages being published and assure that the published message is coming from a reputable device.

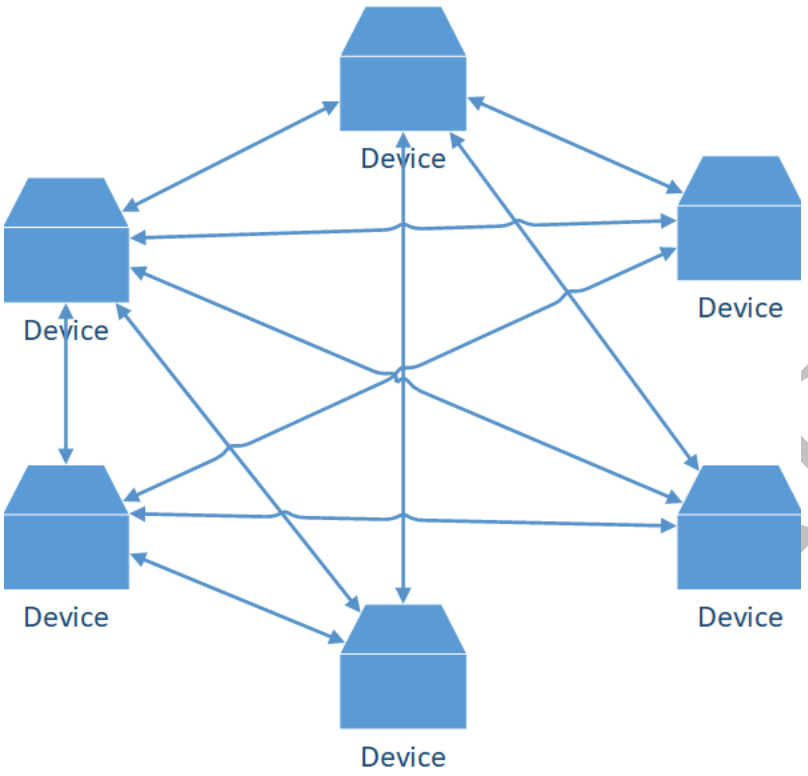

Figure 2- Brokerless Communications Network

Within the brokerless model, an example of which is shown in **Figure 2**, each device is responsible for maintaining its own device registration list. The device registration list contains information on how each device publishes and subscribes to messages. For publishing messages, this list will have rules on what types of messages are sent to which devices. For example, a PMU may publish a general status message to any device that may be listening, while publishing frequency information only to other PMUs. In regard to message subscription, devices must also maintain a list of devices from which they will accept information for specific message types. A certain level of trust must be maintained between devices transmitting and accepting information. Maintaining the appropriate device registration information on each smart grid device adds to computational overhead.

As with device registration, in the brokerless model handling device availability is also managed by the individual devices. Each device knows—according to its device registration list—which devices it should be publishing messages to, and from which devices it should accept subscribed messages. Each device will need to monitor its connections to ensure its published messages are received, as well monitoring to ensure it has not missed any messages it has a subscription for. Careful consideration should be given to this aspect because the order of informational messages can be important for grid operations.

Due to the decentralized nature of brokerless communications, device management may be the responsibility of multiple organizations with the communications network. There is also the

potential for devices to have varying criteria for publishing and subscribing to messages. This fact could make it difficult to establish a standardized way for organizations to manage their smart grid devices because all of the devices may not use a common policy to successfully register with another device.

Table 4 - Potential Pros & Cons of Brokerless Communications

Pros	Cons
Direct device to device communication	Increases computations on devices
Device can be updated dynamically per device	Possibly no standard way to manage devices/message access list
Distributed messaging enables resiliency if any device fails	No guarantee messages are received by all intended devices
Simple message access rules on each device	Difficulty logging and aggregating messages