Security Fabric – Tailored Trustworthy Space
Part 2: End-to-End Management and Security

Charles Speicher

McAfee
3965 Freedom Circle
Santa Clara, CA 95054-1203

Abstract

The Security Fabric framework is a commercial implementation of the "tailored trustworthy space" strategy developed by the White House Networking and Information Technology Research and Development (NITRD) Program and promoted by the Department of Energy for maintaining security of end-to-end intelligent grid environments. For end-to-end security, no one size fits all implementation is possible because of the variety of specific installation needs. The approach must have the flexibility to dynamically adjust to the policies that are appropriate to each individual situation. It must be suitable for the very smallest of situations, but it must also scale uniformly to support the largest of situations which involve millions of managed objects. In that there will be no single victor in the commercial marketplace for a single proprietary design for matters such as key management or other major functional concepts, the Security Fabric provides an interoperable framework that comfortably supports many solutions for individual components using varying standards that can be tested and certified for interoperability.

This *Part 2: End-to-End Management and Security* provides more details on the proposed Security Fabric management system used to provide security and management of the end-to-end system. Beyond just the policy capabilities described in Part 1, the management system is essential in operating a TTS.

1. END-TO-END MANAGEMENT AND SECURITY

The fundamental security and management functions needed for all secure distributed processing are used. Although security management is a subset of the overall management scheme for an end-to-end system, security is somewhat interdependent on all the other management services. It is actually harmful to try to consider security management in isolation from the next larger context of configuration management, fault management, and the other tools that are available for use in the overall management scheme.

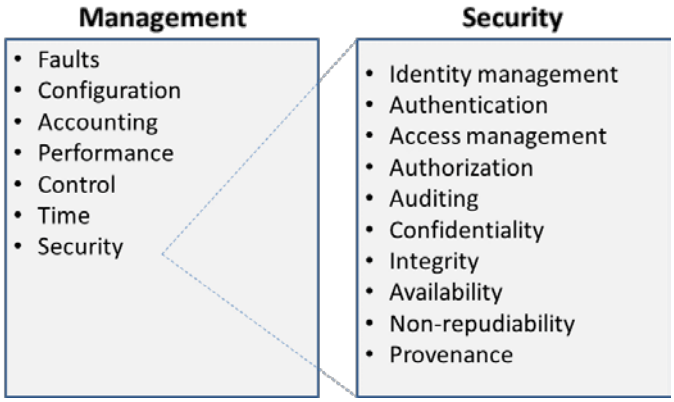

Figure 10. Security and system management are intertwined.

The Security Fabric is a framework for implementation for all these areas of management. None can be left out because of the inter-relatedness of all the parts. If an area were left out, this would be the immediate target of attackers as the easiest point of attack in the fully managed environment of the TTS.

The following sections describe some important details of each of these areas of the framework.

Fault management monitors management event messages signaled from managed devices for problems in the TTS and initiates both automated remedies as well as situational awareness actions.

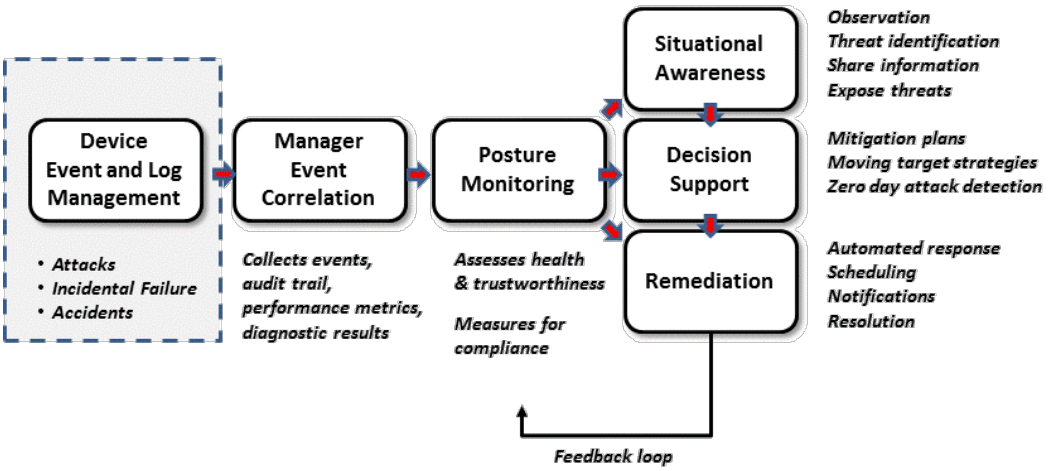

Figure 11. Fault management process flow from the device to the manager.

As depicted in the preceding diagram, event messages from all managed devices in the TTS originate in the device event and log management service on the left and proceed to the right for review and appropriate action. The management device contains a manager event correlation function that collects inbound event messages. Sometimes these event messages are copies of the audit trail from devices that have been momentarily offline. Sometimes these messages are performance metrics that the policy enforcement point logic deemed important enough to forward for central review. Sometimes these messages are the results of test or diagnostic operations currently in progress. Sometimes thousands of event messages arrive in a flood of trivia that really only indicate that there is a single major problem, but it was noticed thousands of times. The role of the correlation logic is to weed out the trivia and refer to human cognizance a simplified picture designed to focus the human mind on what is important and what is actionable, in a recommended priority sequence.

Events are offered to the posture monitoring function to try to determine automatically the health and trustworthiness of the end-to-end system, subsets of the system, and that of individual devices. Here measures for compliance are important profiles that are used to determine if remediation is necessary. Depending on the assessment, one or more of three eventualities can be triggered: situational awareness, decision support, or remediation.

2. SITUATIONAL AWARENESS

If posture management detects that something significant is amiss, but that there are no canned procedures for dealing with the situation, it can refer the event to the *situational awareness monitor* and signal through color, sound, or other attention drawing techniques that human judgment is required with some sense of priority weighting as to its severity in a world of multiple priorities. This monitor is designed to identify threats and share information among multiple specialists in multiple locations as an alert.

3. DECISION SUPPORT

Whether action is warranted, or in what sequence, or whether the alert makes sense in the current situation, it is largely a matter of human judgment – either now or pre-planned for such an event. Sometimes more information is required as to the situation, detailed records, or an assessment of operational readiness to address the problem. The decision support subsystem support needs to have access to logistical and analytical information needed to support decisions by those in command at an operations center.

4. REMEDIATION

Once a command decision is made, there must be automated command sequences for different scenarios that can be broadcast to devices from the command center to initiate certain remedies. There must be notifications to mobile action teams of instructions as to what to do, and how to report back status, success, and failures in resolution. In some cases the public needs to be informed through proper emergency jurisdictional channels. Sometimes test and diagnostic tools need to be employed for resolution, and the results maintained for ongoing quality optimization of the end-to-end system.

Sometimes the problem found is that the end device is genuinely under attack. This aspect of fault management that may require lockdown, isolation, reconfiguration, attack tracking, countermeasures, and return to normal operation following resolution will be discussed in further detail later in this paper. But in general, this is all part of the fault management subsystem.

Configuration Management provides an authoritative view of what devices, services, policies, and control data profiles are authorized to be in the environment plus the current trusted path connections in place.

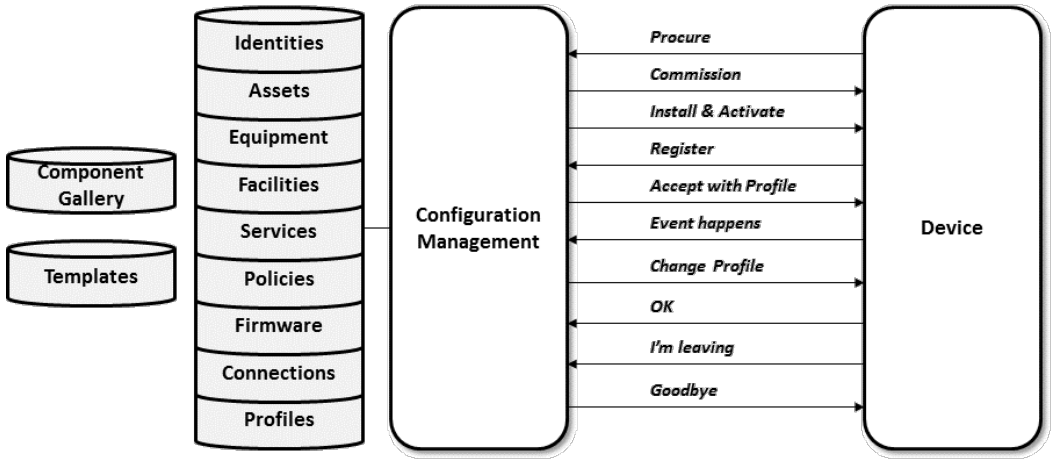

Figure 12. Configuration management interaction.

The life cycle of a device that is to be used in a TTS must be carefully managed and tracked from procurement through retirement. During the later stages of manufacturing or assembly of a device, use of the Security Fabric framework dictates that both a secret identification key and a public key counterpart be loaded into the secure portion of the silicon base of the machine. It is also useful for logistics purposes if the public key is bar coded on the outside of the device. But in no case is the secret key ever allowed outside of the secure silicon – it is not even known by the manufacturer! Immediately after procurement, a new device must be commissioned and tested by the utility for successful loading of initial firmware and data selections indicating the capabilities that the utility uniquely wants to provide when the device goes into inventory.

The public identification and configuration information must be captured by the utility into the utility's configuration management database as a part of the commissioning test. When a device is pulled from inventory for installation, it is important that the public device identification be used by the installer during the installation and activation process. The activation process for a device uses information from the work order or service order that uniquely identifies the GPS location of the service delivery point as well as descriptive parameters that identify the specific service that the individual customer is expecting to receive. This unique profile information is necessary for the service logic and policy logic within the device to use to operate properly.

The use of the Security Fabric component gallery, rule templates, and policy profiles support all tailoring. There are usually not an infinite number of useful selections for a device, so common configurations are easier to provision correctly if done from a standard profile. In custom service situations, new "standard" profiles can be created as minor adjustments to previous profiles if this will help (and it usually does). But what starts as a custom profile can sometimes evolve to be the norm as opposed to an exception. Most popular new services start in exactly this fashion. The core idea is the genetic pattern properties of the component gallery and the rule templates greatly simplify the operation since custom arrangements operate in much the same way as "standard" arrangements.

Accounting management collects metrics needed for billing for power delivery services. There are two styles of computing the actual amounts owed, depending on the creditworthiness of the customer – either before or after service is rendered as shown in the following diagram.

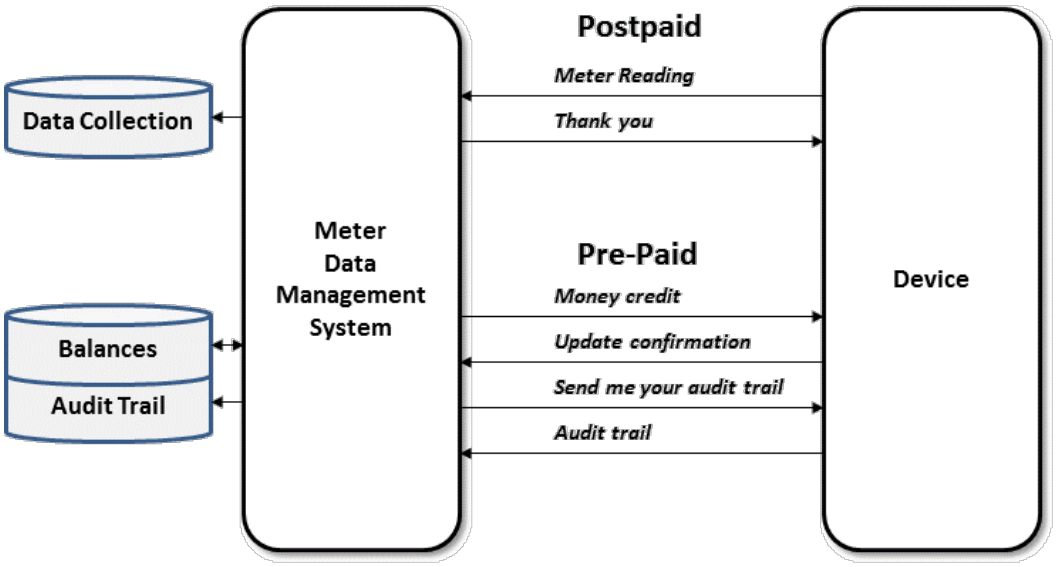

Figure 13. Accounting management supports two kinds of billing.

With postpaid billing, the metrics for what amount of service was actually delivered are sent to the central meter data management system where they are interpreted by rating engines and billed on a periodic invoice, after which the utility collects on the amounts owed. With pre-paid billing, the customer pays an amount in advance for power services and that amount of electronic value is stored into the actual meter immediately. The meter then collects the metrics for what is being delivered, dynamically rates the usage, and then decrements the electronic value in the meter in real-time.

The complexity of real-time pre-paid billing is actually in the process for collection of funds in advance from the customer at various merchant locations such as banks or convenience stores, and then delivering the electronic value securely to the meter. Any time money is moved, there is a tendency for fraud and theft to occur. This always is the case, no matter what the country or what the circumstance. The security of money, or electronic value exchange, is similar to but different from the security of operations. The storage of electronic money inside devices that are located at consumers' homes or businesses requires the same security measures that are used in credit card operations. Attempts to tamper with the elements of electronic money need to be reported and managed as in any other operational aspect of the system where services are subject to theft.

Performance measurement collects metrics associated with how the power service was delivered as well as how the communications systems fared. These metrics are gathered from an end device as shown in the diagram below.

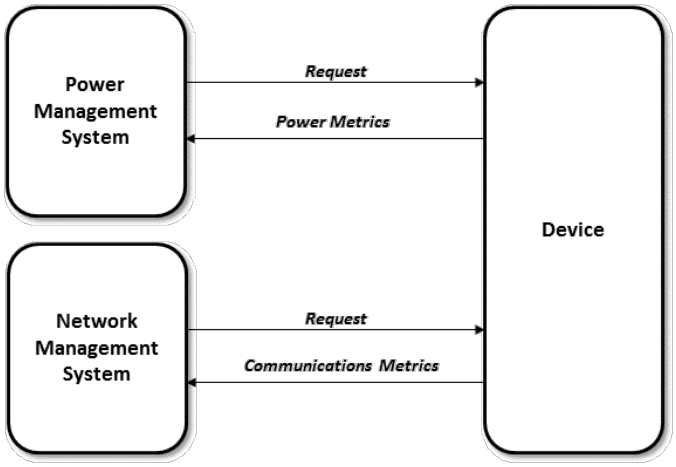

Figure 14. Performance management supports multiple system dimensions.

The specific metrics to be collected vary over time, but the collection and analysis pattern remains the same. The end device needs to be able to adjust over time which metrics it will collect, both from an application point of view and from a management point of view. This means the logic needed for collecting these metrics must be suitably adjusted whenever the metrics are adjusted. This also means that the central systems used to interpret the metrics need synchronized adjustment at the same time. The elements of coordinated change of configuration provided by the Security Fabric framework allow for simultaneous change of policy logic and configuration control variables at multiple points in an end-to-end system.

Typically, different central systems are used to interpret applications metrics from those that interpret management

metrics. Application metrics are used to make real-time decisions about how to regulate the flow of both voltage and current in various areas of the power blanket so that actual power distribution is maximized. However, management metrics are used to monitor the performance of the management system itself to determine whether the control network is configured sufficiently to keep pace with the demands of the situation. Many times the analysis of trends can be made such that additional capacity can be added to the speed of the management network itself prior to any significant degradation that would prevent the management system from being adequately responsive to the evolving situation at hand.

Control management provides a way to deliver management commands to a device, a group of devices, or broadcast to all devices participating in a TTS as shown in the diagram below.

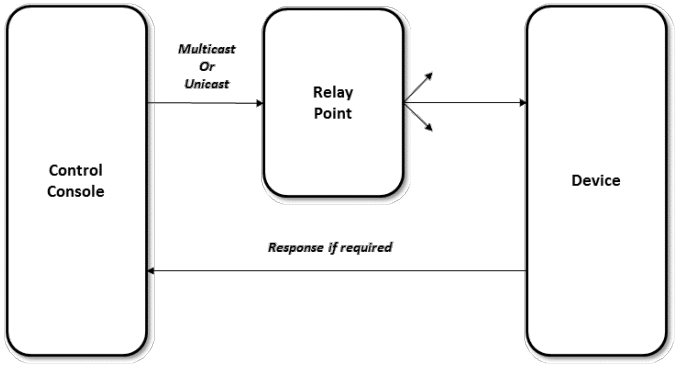

Figure 15. Control is a command pattern.

The command structure needed for the Control Console to talk to millions of devices securely at the same time requires a hierarchical arrangement of relay points – just like the stock exchange. If specific control of an individual device occurs, then a closed circuit response may be prudent. (Forwarding policies at the relay point can monitor for accidental responses to multicasts that could have devastating denial of service effects on a large network.)

Security of control focuses not only at the center and the end point, but also at the relay point.

A secure network connected to a secure network does not automatically mean there is a secure network end-to-end. The elements of the relay point need to be secured and monitored just as carefully as the central console. Yet these relay points are typically located in the field in unmanned locations where physical security is rarely sufficient. Therefore the logical security must be sufficiently strong to circumvent any tampering with the relay point itself.

Secure time service is as essential for security as it is for precision operational measurement. Many management and application controls depend on knowing with precision what time it is, sometimes to thousandths or millionths of a second. The accuracy of audit logs in various different interacting devices must be synchronized so that diagnostic activities can reconstruct what actually happened at any particular point in time to determine how anomalies occurred.

Most network-based systems depend on the Network Time Protocol (NTP) to do this continuous time synchronization since it is fairly difficult to set the time correctly at each endpoint at initialization, or to keep the time properly set after a power outage. However, NTP in its usual deployment is inherently insecure and easily spoofed. Therefore, the Security Fabric framework also includes a secure NTP as an expanded function of the distributed Manager in a communications hierarchy. This is shown in the following diagram.

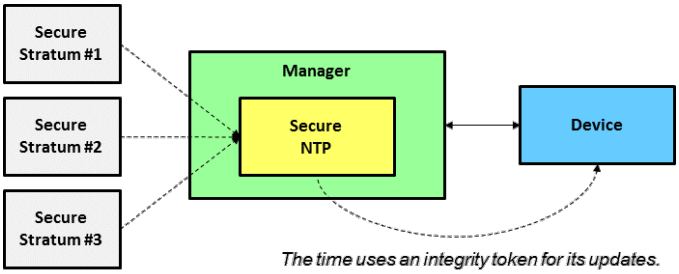

Figure 16. Time management is essential for security.

The National Institute of Standards and Technology (NIST) provides a "stratum 0" atomic clock at the U.S. Naval Observatory that can be referenced as the official time here in the United States. Clocks that reference the stratum 0 clock are called stratum 1 clocks. Clocks that reference stratum 1 clocks are called stratum 2 clocks, and so on. The problem of course is that the speed of light is finite, and any distance at all or latency in the relay of what time it is caused by electronic circuitry causes each stratum to be off a little bit. Knowledge of the distances and latencies involved can help offset this difference, but each opinion is never quite as accurate as the stratum 0 clock itself.

In some cases, a distributed manager has access to a GPS satellite precision timing clock. This is a good thing except for the fact that GPS time has been compromised from time to time. The Security Fabric framework takes this into consideration, and given fiber optic connections, it can securely deliver time synchronization within nanoseconds of stratum 0 time wherever time accuracy is essential to operations or management.

5. SUMMARY OF PART 2

End-to-end management of a distributed system such as a TTS requires all the classic controls, including the following:

- Fault management
- Configuration management
- Accounting management
- Performance management
- Control management
- Time management
- Security management.

This second part in the series on the Security Fabric has introduced these capabilities as they are implemented in the framework.

Fault management handles automated posture monitoring and remediation as well as human-oriented situational awareness functions.

Configuration management supports structural control as well as software and control data management.

Accounting management supports both directly and indirectly the billing of services as measured and metered within a TTS.

Performance management supports the collection of metrics for the management of power as well as the systems and network elements of the management system itself.

Control management supports the transmission and response of operational directives within the end-to-end management system.

Secure time service synchronizes the clocks in all the devices in the TTS, even given attempts by others to compromise the system.

Part 3 of this series will further develop the details of the security management system – which depends on the features of the end-to-end management system elements described in this Part 2.