WHITE PAPER

Envigilant™ Systems

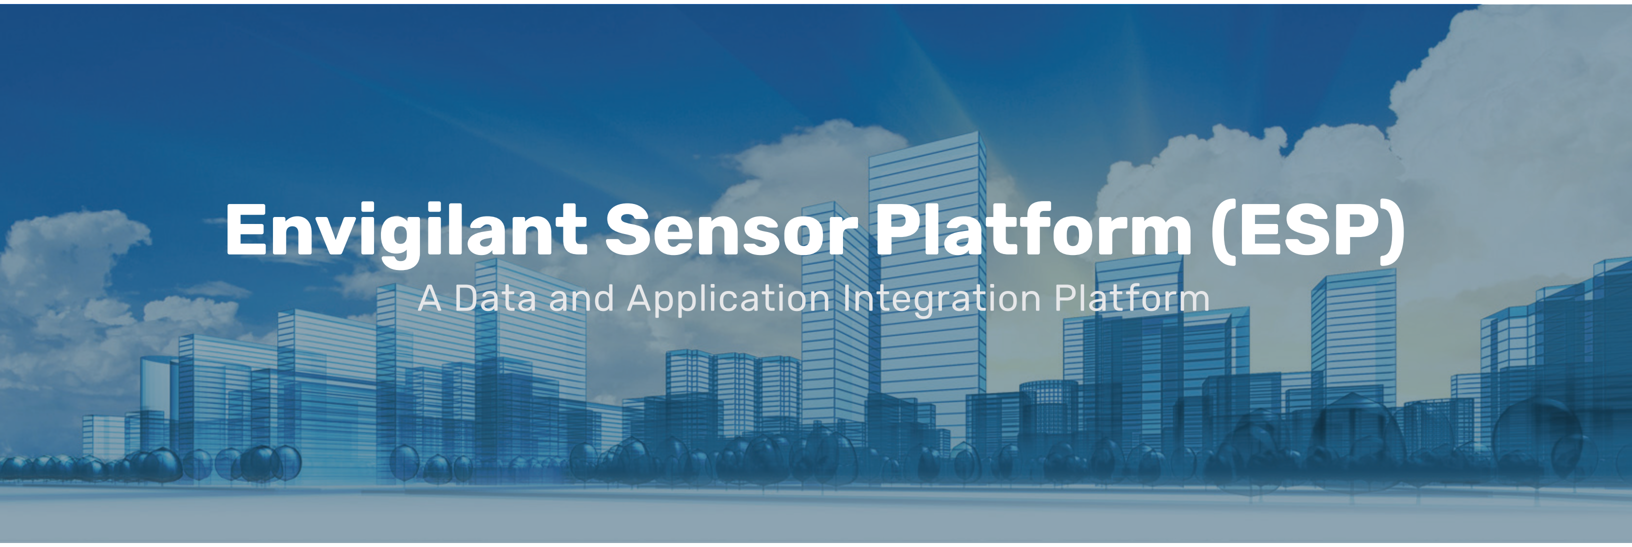

Envigilant Sensor Platform (ESP)

A Data and Application Integration Platform

Contents

For a free consultation, visit www.envigilantsystems.com

1 Executive Summary

When technology advances quickly, it is up to the user community to help find innovative ways to deploy these new capabilities to improve the lives of their members. For local and municipal governments, this type of work typically falls into the Smart City and Community (SCC) category. The promise of new sensing technologies, advanced machine learning algorithms, and more capable edge computing devices is to implement systems that improve the quality of life and security in urban areas. Being able to deploy these technologies as they develop requires a robust, flexible platform for integrating their data and understanding what they are telling us.

This white paper describes the Envigilant Sensor Platform (ESP) from Allied Telesis. It has been built around standards for data interchange and with scalability in mind.

Not all cities are the same. Every city has its very specifics given by its history, its culture and its resources. So, even the most appealing and innovative approach cannot have the same impact in two different cities around the world that are actively looking to became smarter. To really make a difference, to get closer to having a Smart City, the approach and the selected technologies should match the specific circumstances.

The simple adoption of technology does not necessarily make the city "smart." The technology infrastructure only creates the environment for the smart development of the community in a sustainable way. For this, decisions made on the real and actual information are crucial. Either we speak about emergency response or "if-then" scenarios for mid and long-term projects.

In order to support these decisions, the city must be ready to collect, analyze and capitalize on vast streams of data from different data sources, all arriving at varying volumes and high velocity. The ability to effectively utilize this constant flow of data is dependent on implementing an effective software platform that can capture, manage, and prepare data for analysis and visualization.

Once implemented, ESP-based solutions connect any data source the city needs, manages both the source itself and the collected data, stores and secures the data, and allows the user to create and build specific applications needed for the city to make life easier, safer and happier for the citizens. The ESP platform must support on the infrastructure layers, from data collection to data processing and storage. It is designed to support a wide range of application programs, irrespective of how small or big they are, focusing only on solving business problems.

Built on open standards to connect anything and manage everything, the solution has crucial features that help the cities to transform their vision into reality, by providing:

- A uniform mechanism for data collection and data distribution leading to better ROI and lower TCO;
- A common data warehouse providing independence from application providers;
- Open standards-based interfaces providing long term sustainability of projects;
- Compatibility with analytical suites covering multiple areas of interest for the City;
- A reduced Time-to-Market due to uniform data access interfaces for all the applications; and
- An easy expansion of the number and types of data sources covering future city projects and adapting itself to multiple types of data.

Analyzing the data turns information into intelligence that helps people in charge to make better decisions. The greatest benefits come when data is processed across multiple sources and silos, and a common data collection platform is key to making this a reality.

2 Reference Architecture

Following the principle of using open standards to future proof the solution, ESP was designed following the industry reference model issued by ITU [Y.2060/Y.4000] where the main layers are as follows:

1. **Device Layer**: contains all the sensing devices together with any gateway functions needed.
2. **Network Layer**: includes network connectivity and transport functions together with resource control and management functions
3. **Service Support & Application Support Layer**: includes common capabilities which can be used by different IoT applications (e.g. data processing and data storage) combined with any other specific capabilities support functions to different IoT applications.
4. **Application Layer**: contains all IoT applications needed

Following this model, ESP supports construction of solutions composed from the following types of components:

1. **Devices and Sensors**: Includes multiple categories of sensing devices (e.g. video-cameras, traffic sensors, environmental sensors for gases, particles, noise, temperature, humidity, temperature, etc.). All sensor types may interact directly with the network layer through dedicated standard interfaces (Ethernet, WiFi, Bluetooth, Zigbee), being able to send the collected information to a central processing location. If needed, the system may provide (included with the sensing equipment or separately), edge computing and storage capabilities able to store and process data near to the collection point for fast resolution or analytics preparation.
2. **Communication Network**: Includes all network elements (switches, VPN gateways, etc.) that connect the primary data sources (devices and

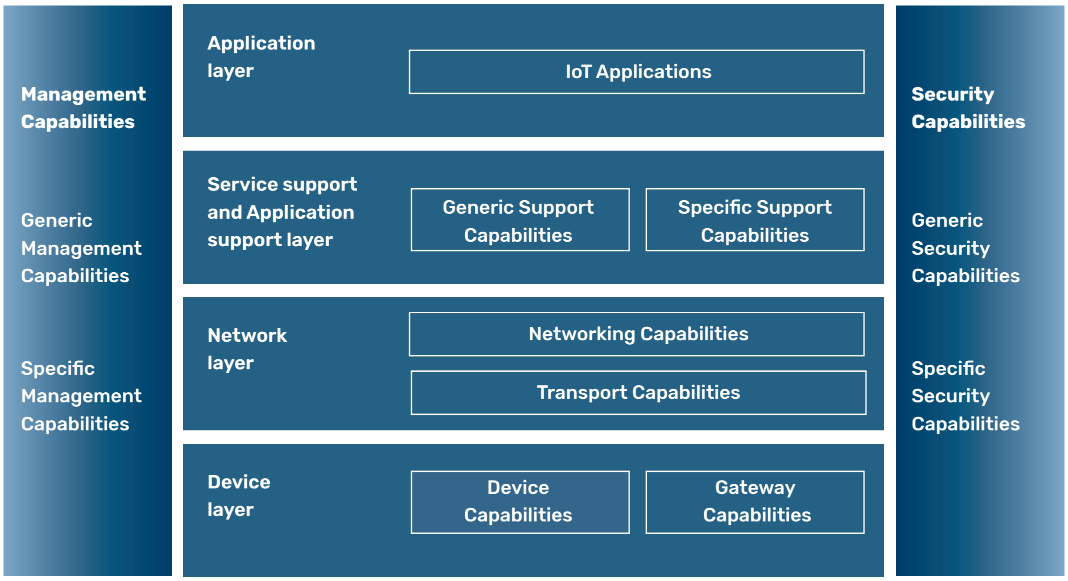
Management Capabilities
Generic Management Capabilities
Specific Management Capabilities
Application layer
IoT Applications
Service support and Application support layer
Generic Support Capabilities
Specific Support Capabilities
Network layer
Networking Capabilities
Transport Capabilities
Device layer
Device Capabilities
Gateway Capabilities
Security Capabilities
Generic Security Capabilities
Specific Security Capabilities

Figure 1: ITU Reference Architecture

sensors), together with the management and control software that allows continuous operation and monitoring.

3. **Services and Applications**: Includes both hardware and software components. The software components include IoT platform functionality as well as components for managing both the data streams and devices themselves:

    a. IoT Platform: provides all the necessary functions for collecting and storing all operational data captured by existing devices, including the result of pre- and post-processing of data. It also provides the interaction with other service sub-systems like Video Management Systems (VMS), License Plate Recognition (LPR), etc. In addition, the platform allows data import/export operations and provides the necessary elements for stream processing. The IoT platform provides the northbound API interface developed to interact with any business applications and dashboards needed by the city.

    b. Data stream processing framework: includes all the capabilities to process the data streams extracting operational information from the input data.

    c. Device management: includes a flexible and scalable solution to monitor the entire infrastructure including the sensors, networking devices and computing and storage elements, offering an accurate and near real-time status of the systems

4. **Business Applications**: the system allows for direct or indirect integration with the following types of generic applications:

    a. Dashboard: real-time and historical visualization of operational data (e.g. values read from sensors, result of data processing, etc) grouped and sorted according to operational requirements specified the end-user

    b. Data Processing: any type of stream or batch data processing by adding software modules created by custom request and adjusting the existing resources (computing and storage) to match the required performance.

Depending on the operational role (from the business/operational view-point) the components of the bottom three layers remain largely similar for a significant number of business applications, all of them taking advantage of the activities performed on "Edge" and "Platform" levels. This is the part that dictates and enforces the overall data collection policies (e.g. what data sources are allowed for input), security policies (e.g. access restrictions, additional cyber security rules) or data processing policies (e.g. related with the processed parameters and how the result is delivered) to achieve a specific outcome.

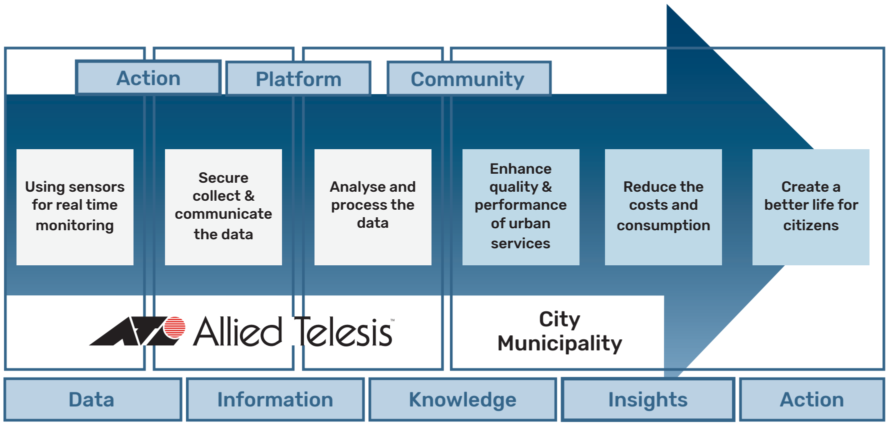
Action
Platform
Community
Using sensors for real time monitoring
Secure collect & communicate the data
Analyse and process the data
Enhance quality & performance of urban services
Reduce the costs and consumption
Create a better life for citizens
Allied Telesis
City Municipality
Data
Information
Knowledge
Insights
Action

Figure 2: Business Architecture Model

This capability allows different behaviours of the system to be implemented depending on the current needs of the entity that controls it. Based on this approach and depending on the community interests, a set of pre-defined "operating modes" can be defined.

3 Physical Architecture

Allied Telesis and our partners have used the architecture described above to implement a physical information system and technical architecture for a variety of different Smart City and Communities (SCC) projects, each of which customized the blocks in order to respond to the concerns of the city stakeholders.

The Envigilant Sensor Platform (ESP) was constructed to facilitate options for deployments that utilize edge compute devices when available, by allowing the processing to be shifted onto ruggedized devices that have a growing amount of resources. In the past, the capabilities of edge compute nodes only allowed for implementing gateways, perhaps with some protocol translation. Now that edge computing has matured, we offer the ability to bring visualization and real-time analytics right at the point of sensor data acquisition.

3.1 Sensors & Edge Services

The sensors are the primary collection point for information from the field. The system includes sensors capturing any time-series data including environmental data, traffic monitor data and other types useful for SCC projects. Additional intelligent sensors can be added in the future, by using the existing software programmable interface to connect them.

The instruments (e.g. cameras and environmental sensors) are either connected directly to industrial switches using Fast Ethernet (FE) and Gigabit (GE) copper ports, or, when the available interfaces are incompatible or the computing power needed for that role is not enough, via an intermediate equipment (IoT gateway) that provides the additional resources.

In order to accommodate special requirements in this area, special capabilities are implemented under the generic name of "Edge Computing". Edge Computing brings computation closer to the data source to accomplish a wide range of functions starting from data aggregation, meta-data augmentation and protocol conversion (e.g. converting a MODBUS message into an MQTT payload), which allows for fast decision and action. Edge computing significantly decreases the volume of data that must be transferred

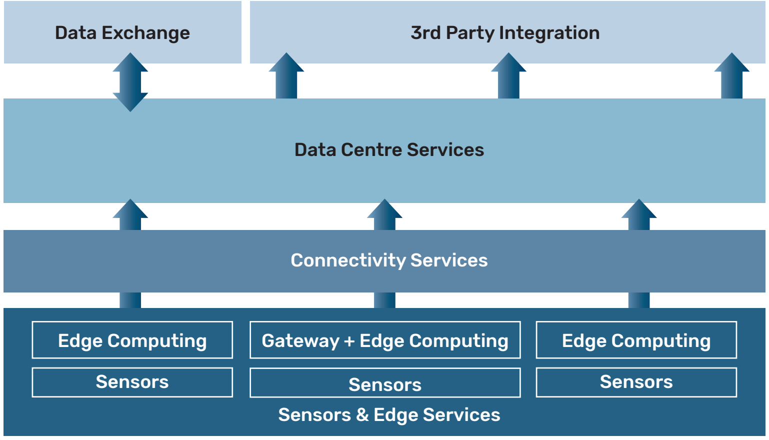

Figure 3: Solution's Block Architecture

between data source and data center, thereby reducing latency and improving quality of service. It also acts as a data buffer allowing for temporary retention of the collected data in the event of a loss of connectivity to the data center. In Smart City applications edge servers are the ideal choice for handling video processing, analytics and other time sensitive workloads allowing for a faster response to the detected events.

3.2 Connectivity Services

The network role is to provide secure interconnection between all of the components in the solution, including the transportation of sensor data into the Edge system, as well as between Edge, Central, and third-party applications. The general scheme of the network follows the Allied Telesis solutions for Autonomous Infrastructure, bringing more intelligence to the edge to allow to the all parts of the system to function in a coordinated manner. Given the distribution of these components, from sensors at the point of data acquisition, to edge servers where those data are brought into the first level of storage, all the way back to the data center and to other cloud applications as needed. The resilience, security, and durability of the network is of primary concern.

Splitting the networking domain into multiple distinct functional zones, based on the primary functionality provided by each zone, is a good method to allow further expansion of coverage both in terms of covered areas and number of sensors. Usually, three distinct infrastructure zones are implemented:

1. **Access Zone**: connects together in a resilient manner the sensors scattered throughout the city, independent on their type
2. **Distribution and Aggregation Zone**: connects together all the access zones, collecting the traffic originating there and transporting it for processing to Data Center
3. **Central Zone/Data Center**: connects all the computing and storage used for data processing, including management and security applications as well as remote VPN connection aggregators.

This classical structure brings the flexibility required for the future development by establishing a baseline definition for the functions and policies of each zone.

The Access Zone of the network is built up with suitable equipment for connecting the sensors (e.g. network traffic surveillance cameras, environmental monitors, etc.) mainly in outdoor environments.

Upstream, the switches are connected using GE FO ports, with optical transceivers fitting the distance requirements for each segment. The uplink ports should support redundancy protocols (ITU G.8032 or similar) in order to provide a highly reliable topology (e.g. ring) that has a minimum re-convergence time of tens of milliseconds. Local intelligence, running on the switch level, may be used to automate regular tasks (e.g. performance check) or incident processing response tasks (e.g. unresponsive sensors). As a security measure, the fibre optic links are constantly monitored for any variation of the signal level that could indicate interception attempts. This capability to constantly monitor the reception level on all FO ports should be integrated in each device part of this area, for early detection of security incidents (e.g. FO tapping) or modifications of the FO infrastructure (e.g. excessive cable bending).

An important capability in the Access Zone is to provide the right amount of power to the sensor instruments installed. The need can range from as low as 7.5W up to 90W. PoE capabilities embedded in the access switches serve this need by negotiating and delivering the requested power.

The presence of edge computing capabilities in this layer (e.g. triggers, scripts, containers, APIs) make it possible to process data on the fly with significant benefits in terms of operational efficiency.

The Distribution & Aggregation Zone is built with higher bandwidth switches (10G) using suitable SFP+ optical transceivers. These will ensure fibre optic connectivity with access rings as well as connectivity between core network members, like the Data Center and operational sites. Usually, this is a Layer 3 type of network to give more control over the traffic. If possible, the topology of this network will be also a ring-type, potentially running same kind of redundancy protocols beneath the use L3 routing protocols for faster re-convergence. The virtual chassis (stacking on short and long distances) is a feature that can help provide a better organisation of this zone.

The edge computing model tends to blur the functional space defined by the static Access/Aggregation/DC model, as some capabilities specific to Data Centers (like processing and analysis of data) are now performed near the data acquisition point. This enables ESP to support a redesign the above structure of the connectivity service level into a much more flexible structure with only two levels, focused on data transformation rather than data communication:

1. Data collection and preparation: the zone where the data is acquired, prepared for processing and then transported to the business applications that consume it;
2. Data analysis and processing: the zone where the data is processed according to business rules dictate by various applications needed by the city operations.

This new structure is made possible by the existence of the API interfaces on all levels from collection, preparation, transport, and analysis of the data. This leads to a model where the fundamental operations of collection, preparation, transport, delivery and processing of data are interlinked, moving the focus on data flows rather than just network links.

For both zones, an important consideration is related to the capacity of the infrastructure to be influenced and commanded directly by the business applications. The Software Defined Network model gives control over the network ports to the applications allowing a unprecedent level of traffic flow flexibility, adjusting the communications parameters (e.g. quality of service, access, etc) to the immediate need of the application. A direct link for the applications to manage the behaviour of the network ports (e.g. via OpenFlow or other southbound protocol) are mandatory for ensuring the level of control required by such complex systems as city infrastructures. This centralised control feature is available to be implemented where the data processing applications will require it (e.g. global security policies, analytics suite, etc.).

Central configuration control provided by agents inside the network operating system will allow zero-touch replacement in case of defect or even zero-touch deployment for new equipment in conjunction with a central master function for individual areas or for the whole system.

3.3 Data Center Services

The ESP IoT Platform allows for the division of processing between edge components (running Envigilant Edge) and the data centre (running Envigilant Central). Any number of Edge components may be installed and related back to the same Envigilant Central instance. The Edge processors allow for real-time analytics to be run where the data is acquired.

3.3.1 Envigilant Edge

As stated above, the traditional IoT platforms have an "edge" component typically provided inside a IoT gateway. If analytics were to be used for

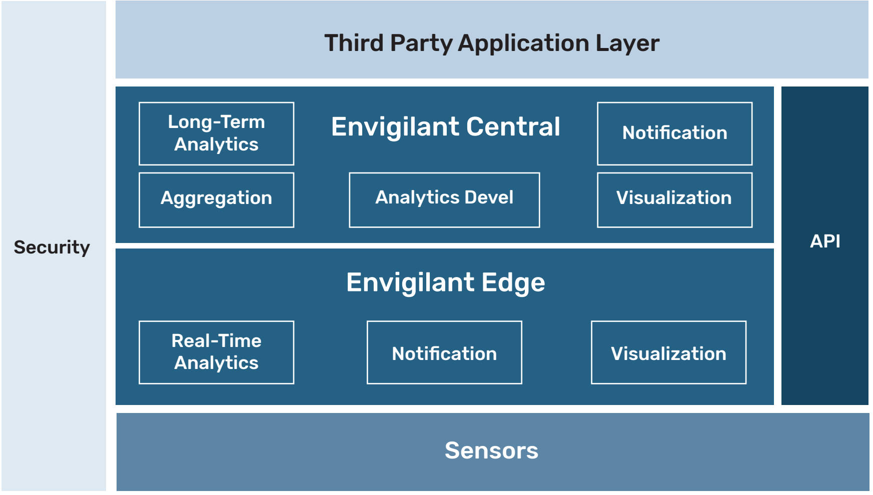

Figure 4: IoT Platform Block Structure

scoring data, that would have to be done in the cloud-based compute resource. This component isn't strictly required to be physically in the data centre – the architecture allows for it to be placed anywhere in the solution. If your sensors are connected into the Data Centre (or physically located there) then the ESP Edge instance can live in the DC itself.

As edge computing devices have become more capable, the days for passing all the data to the cloud are gone. The modern IoT platform has to be able to store, manage, and reason about data at the edge in real time. The Envigilant Edge server is built to provide the following services:

- **Storage** – ESP Edge provides a timeseries database, optimized for data write, so that data acquired from the sensors can be guaranteed captured and stored. This provides a durable storage for the data, and ESP Edge can also be clustered to provide a high-availability option.
- **Analytics on-the-fly** – descriptive analytics that look at data in transit (such as simple thresholding) is easily done by any Edge server. ESP Edge, with its storage ability allows for more sophisticated, including scoring based on machine-learning (ML) algorithms, to deployed at the edge. This means that the power developed by your data scientists and data analytics team can be deployed where it's needed.
- **Visualization** – for those edge deployments that are near operators, such as drivers on vehicles or machine operators on manufacturing floors, visualization of the data can be performed by the ESP Edge server. Coupled with touch-screen controls, this unlocks the data being captured by the sensors as well as those synthesized by your analytics to inform the people who need to make in-the-moment decisions about the local process.
- **Notification** – ESP Edge also provides automated notifications, including rule-based transfers of data to external systems (including ESP Central) as well as emails and other human-based media.

The ESP Edge server is also responsible for device management. Since the sensors connect directly to the Edge server, this is the place where adaptive systems for monitoring the health of devices, suspending them when upstream systems become overloaded, or detecting cyber threats from compromised devices should be implemented. The Edge server itself can gather and store data on device health and use analytics to react to situations as they arise.

3.3.2 Envigilant Central

There are a multitude of applications of the sensor data that are desirable for an IoT platform to facilitate that involve longer term analysis of the data. The role of ESP Central is to provide an aggregation point for data from multiple ESP Edge servers in order to enable the organization to make full use of the data that is being gathered.

ESP Central was built on the foundation of openness. All of the data models, data storage engines, and analytics tools are open and documented for your use and extension. ESP is guided by the philosophy that these types of tools must be able to be configured to individual customers' environments.

- **Aggregation** – ESP Central has a NoSQL data store to facilitate the aggregation of data not only from multiple ESP Edge instances, but also from other operational and back-office IT systems that have information that needs to be fused with the sensor data. ESP Central leverages the power of the data centre to create a data lake for all relevant data.
- **Visualization** – key to understanding data is visualizing it in a way that makes sense to the people trying to solve problems. A good visualization environment is the starting point for any serious data analytics problem. Using Grafana and Kibana, users can explore their data as well as develop meaning dashboards for KPI monitoring.
- **Analytics Development** – the flexibility of the ESP Central data lake has the ability to evolve as the data scientists learn more about the data. With integrated tools such as Kibana, Jupyter, Zeppelin and a host of open-source ML packages, ESP Central was constructed to facilitate the development of sophisticated data scoring models. This includes the development of Long-Term Analytics models for looking at sensor and operational data over weeks, months, or longer to provide insights from the data.
- **Notification** – developing routines for identifying patterns of data that are meaningful is the first step, but it is a requirement that the system can notify people when certain conditions occur. ESP Central has all of the notification capabilities that build on ESP Edge for communication with appropriate personnel.

3.3.3 Security

Keeping IoT systems running securely, especially when widely distributed, requires a variety of tools and solutions. The implementation of the Envigilant components provide a number of different options for securing data at rest as well as in motion, and can be selected at deployment time to match the computing environment restrictions as well as the requirements.

Allied Telesis networking equipment can implement segmented and virtual private networks to keep data in motion from being visible to any resources other than those required by the solution. For network segments that may traverse the public internet this component of the security profile can add an additional layer of trust in the system.

Inbound messages use the MQTT protocol which is protected by TLS encryption. The MQTT protocol is a resilient, inbound-only mechanism which protects the system from data breaches. The TLS security protocol is widely used and supported, and can provide a layer of assurance that the messages are originating from known, trusted devices.

All Envigilant servers use certificates to authenticate themselves and to assure that devices that present themselves on the network can be assured that they are authorized to send or receive content. In addition, data at rest on the Envigilant servers (both Edge and Central) can have their backing data store encrypted, and access to all data can be protected through password challenges.

Every encryption layer consumes both processing and memory resources, so the various security components chosen for any project will be matched to the solution needs and available resources.

3.3.4 APIs

In keeping with the open fundamentals of ESP, RESTful APIs are provided for all of the data storage mechanisms and communication components, so that any technical staff can dig into the data and integrate ESP into operational systems. ESP's goal is to unlock access to data so that deployments facilitate real ROI from the data gathered, without requiring the engagement of specific consultants.

3.3.5 Data Import/Export

The IoT platform has been designed to be interconnected with other external data-bases in order to have access to existing information, potentially in different formats. The goal is to provide the data analytics applications with all the necessary data to produce meaningful results. Symmetrically, the operational (sensor) data, can be extracted and exported to external data bases, to be used elsewhere or for regulatory purposes.

A variety of methods are provided for sharing data between systems. ESP was designed to offer different open methods for connecting to data streams coming from sensors. The breadth of these methods allows for different types of applications and different types of users to manipulate the data.

- A web socket interface is available from ESP Edge instances to get a real-time stream of data to traditional web hosted applications
- OpenAPIs are available on both the Edge and Central systems to programmatically query and summarize data
- NiFi, an open-source data flow engine, can connect to ESP Edge instances and provides a visual environment for capturing and manipulating data coming from sensors.
- APIs on ESP Edge, including support for Spark streaming jobs as well as Java and Python based analytics, are available for custom manipulation of data as it is acquired.

Envigilant Systems

Secure IoT Solutions & Services

For a free consultation, visit www.envigilantsystems.com

Envigilant Systems | San Jose, CA | 1-800-424-4282