

The background of the top half of the page is a blue, glowing tunnel-like structure composed of a grid of lines that recede into the distance, creating a sense of depth and movement. The lines are illuminated with a bright blue light, giving the impression of a data tunnel or a futuristic network.

Building Successful IoT Solutions

Background

Have you ever tried to deploy an Internet of Things (IoT) project but gotten stuck trying to manage the data? Chances are if you tried to use a traditional relational database to capture and analyze the data you probably found it worked fine in the lab with smaller amounts of data, but when you moved the project to scale (or deployed it in the real world) those tools just didn't suffice. Or, perhaps your analytic requirements, both in querying and visualizing the data, weren't achievable as the number of sensors grew.

The demands of real-time sensor data provide a new set of challenges for IT professionals to implement. The proliferation of time-series data has put strains on traditional tools – and putting new requirements on systems for achieving success.

Problem Summary

The term "IoT" has been around for years and refers to capabilities that are made possible by a collection of different technologies. While electronic sensing technologies have been around for decades, the IoT refers to the following capabilities:

- Sensors, often small and embedded in other devices, that collect data about the environment
- Network connectivity to get real-time access to the sensor data
- An aggregation technology for creating historical collections of sensor data
- Security systems to identify and assert the fidelity of the data being acquired
- Software for providing information about the sensor data streams in real time
- Software for learning from sensor histories to reveal patterns in the data that are predictive

For a free consultation, visit www.envigilantsystems.com

These capabilities are all domain-independent. Any IoT implementation will have these capabilities in different measures, depending on the domain and the problem being solved. The range of systems that are billed as “IoT Systems” is quite varied. They can be “smart home devices” to monitor and control your home environments (temperature or security, for example); medical instruments in a hospital; agricultural sensors monitoring and controlling growing conditions; location and operational sensors on public transit vehicles; or detailed sensors measuring air flow, gas flow, temperature, current, etc. on a manufacturing floor. All of these constitute IoT Systems being used for very different purposes.

Database Requirements

Let’s look at more detail in the requirements and important attributes of the database systems that are deployed in typical IoT settings. [Note: the database technology deployed when discussed in terms of IoT data is typically referred to as a “historian” – it’s just a database that is used for dealing with sensor streams of data.]

First, and most importantly, the database has to support (and in some ways be optimized for) time series data. This implies being able to store measurement data indexed by a timestamp and must be able to accommodate different streams concurrently. Remember that a measurement may be a complex item itself, not necessarily a scalar value. For example, a GPS unit provides measurements for latitude, longitude, altitude, and accuracy (one or more dilution of precision values).

Second, the database must be able to keep up with all of the sensor streams it is receiving. This typically means that the write performance of the database needs to have the highest priority in terms of operation speed. As data comes from the sensors, of which there can be multiple streams each with multiple values, the database must be able to add them to the data store fast enough to keep up with the aggregate flow.

After those two requirements, there are a number of other factors that need to be considered when building an IoT platform.



High Availability

Do you need to ensure that all data is captured by the database, regardless of hardware issues? This is typically expressed in the number of “9s” of availability that you need – 5 – 9’s meaning you want the system to be available 99.999% of the time. Remember that providing a highly-available system implies redundancy in hardware, networking and software. All three must be considered – and this normally comes with increased cost and power requirements.



Visualization

Do you need to visualize the sensor data streams in real time? Typically the answer to this is “yes” (but not always). If you do need this, then that implies that your database needs to be able to provide data fast enough to the visualization engine to make it current, or your architecture needs to support another access method to data streaming in from the sensors.



Real-time Analysis

Do you need to be able to “score” data as it comes in from the sensors? This can be as simple as thresholding information (aka “descriptive analytics”) or as sophisticated as models built from machine learning algorithms. Depending on the type of analytic, this can require looking at the most recent values from the sensors, or possibly taking data from a time window for further processing.



Offline Analysis

With the advances in Artificial Intelligence such as machine learning, deep learning, or other analysis paradigms, data scientists typically need a larger data set than the historian will manage and, even better, labels on the data streams of “error” conditions. These larger, labeled data sets form the basis of analysis that can produce the real-time analytics that mines true value from your sensor data. Your database and architecture may need to support the aggregation of longer-term data for your data scientists to successfully do their job.

Envigilant Sensor Platform

The Envigilant Sensor Platform (ESP) was designed and built as a new kind of IoT Platform – one that can leverage open technologies and be deployed in a truly distributed fashion. ESP is built in a way that incorporates a variety of different database technologies that can be utilized based on the precise requirements for each individual deployment.

Options include:

- A time-series RDBMS for more resource-constrained deployments that need to collect data fast and repose them until they can be aggregated
- A wide-column NoSQL database that can be installed in a clustered environment to unify data over a longer period of time and provide time-series visualizations and analytics
- A Lucene-based (search-engine) store for aggregating data for other search type operations

The ESP allows for a federated architecture, which provides:

- A distributed architecture where remote nodes have compute and store capabilities locally, with no uplink networking required
- A central server that can be used to aggregate slices of data from the nodes (or entire data sets) at the users' discretion
- Flexible node software to provide anything from small-footprint data collector to a high-availability edge compute service
- The ability to deploy these edge compute nodes that can capture data, run analytic functions, and notify users of possible imminent problems all without access to a central server
- Data visualization and analytics processing that to allow for the development of machine-learning and deep-learning based approaches to data investigation
- Support for deploying ML/DL scoring on edge nodes for maximum autonomy

The ESP was organized specifically to allow for scaling from large, multi-tenant data management and analytic capability as well as compact node configurations to allow the problems' requirements to dictate where resources are placed and how problems are solved.

Without leveraging these newer, more modern data storage and access technologies, you may find it quite challenging to get your lab IoT project into full-scale production. The good news is that the same underlying technology improvements that make these ubiquitous sensors possible have come hand-in-hand with new hardware and software techniques for managing these data. By broadening your platform and keeping an eye on what your ultimate needs from the data are, you can successfully architect a system that will meet all of your needs.

The team at Envigilant Systems (<http://www.envigilantsystems.com>) has built data integration solutions for many challenging real-world applications. Their proven experience and award-winning solutions can help you design, build and deliver successful IoT projects, large or small.



Secure IoT Solutions & Services

3041 Orchard Parkway, San Jose, CA 95134, USA

1-800-424-4282

info@envigilantsystems.com

ENVSB009 RevA

About Us

Envigilant Systems is a business unit within Allied Telesis that has engineering capabilities to build security-focused IoT solutions for manufacturing, transportation, and smart building applications. Backed by 30+ years' experience building secure, reliable networks, Envigilant Systems is ideally placed to design, implement and provide services for tailored solutions to meet demanding requirements. With a global sales and support network, Envigilant Systems is solving challenging IT problems worldwide.