influxdata®

AN INFLUXDATA CASE STUDY

ADLINK and InfluxDB Deliver Operational Efficiency for Defense Industry with Edge IoT

Chris Montague
Senior Solutions Architect, IoT Solutions and Technology, ADLINK

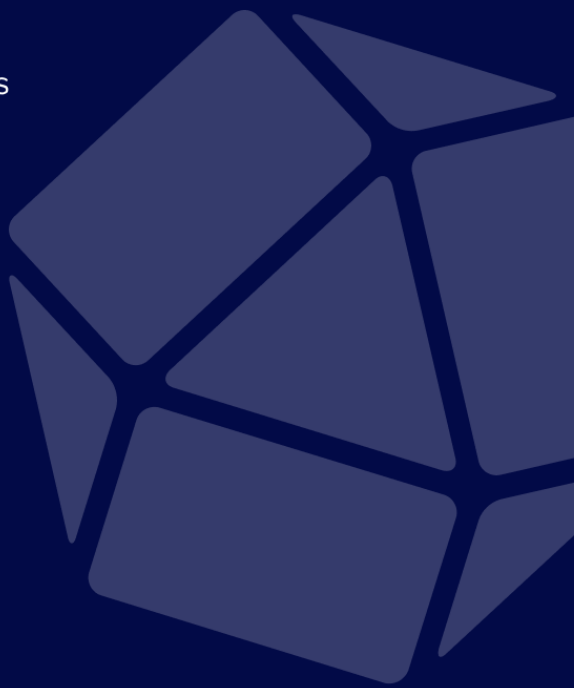

ADLINK TECHNOLOGY INC.

JUNE 2019

Company in brief

Headquartered in Taiwan, ADLINK has operations in the United States, UK, Singapore, Beijing, Shanghai, Shenzhen, Japan, Korea and Germany. ADLINK products are currently available in over 40 countries across five continents, with worldwide distribution networks and more than 1,800 employees. ADLINK is proud to be associated with many major technology leaders and Fortune 500 companies. With design and technology centers in the U.S., the Pacific Rim regions and Germany, ADLINK is a technology-leading platform provider in the embedded computing industry.

Case overview

ADLINK needed to meet the time series data collection and visualization needs of its client: A major military and aerospace contractor operating globally. The client wanted visibility into why their test machines were breaking and real-time diagnosis and predictive maintenance that would reduce loss of revenue caused by unexpected downtime.

ADLINK uses InfluxDB as a part of their ADLINK Edge™ IoT smart gateway solution to store IoT data so that they can visualize it for monitoring and analysis and thereby extract actionable operational efficiency insights. Their client is using their ADLINK Edge™ to connect existing software, equipment and systems securely, seamlessly and cost effectively. The company is now able to extract real-time operational data from its materials testing chamber and stream real-time data to drive efficiency by enabling predictive maintenance.

The benefits that ADLINK's defense contractor client gained include reduced downtime from planned and unplanned maintenance shutdowns, improved machine performance, greater accuracy and reduced cost. ADLINK Edge and InfluxDB fit together in a highly effective Edge IoT deployment which delivers operational efficiency for the end user.

> *"The issue really is loss of revenue. And each callout (maintenance window) was costing about $100,000 per year for the company because you have to be a specialist to fix these machines."*
>
> ***Chris Montague**, Senior Solutions Architect*

The business challenge

As a global company and one of the leaders in Edge computing, ADLINK provides a robust platform for real-time data connectivity, solutions and applications. It is sought by enterprises seeking to extract value from their data.

ADLINK: A Leader in Software-Defined IoT

One of ADLINK's customers in the defense industry has large Thermotron test chambers where they test equipment to ensure that it's fit for use in the field. These machines are a simple device in terms of what they do because they're typically either baking or freezing something. But the mechanism associated with the refrigeration units, the heating and cooling involved, and the circuitry lead them to often fail.

Materials Testing for Defense Industry:
Connecting and Streaming Data to Enable Predictive Maintenance

Since the client had different machines to schedule tests with, the schedule would be easily upset when one of their machines is out of commission. And since the machines typically took a few hours to power up, It's only at that point that they would indicate a malfunction. At that stage, the client would have to move any of the priority materials being tested to another machine and then authorize overtime. With testing conducted around the clock, it puts more stress on the other machines when a single machine is out of commission. Any machine outage, and its associated repair costs, caused serious loss of revenue. Repairing these machines required specialists, and each callout cost the company around $100,000 per year.

The business challenges that ADLINK sought to solve for their client were:

- Extract real-time data to enable predictive maintenance
- Reduce cost of planned and unplanned downtime
- Monitor machine performance for accurate diagnostics
- Enable training of engineers to maintain machines

The technical challenge

ADLINK set out to empower the client through its technology, to go from reactive maintenance to proactive maintenance. When a machine breaks, the first action the client would traditionally take is to have an engineer come punch in some codes in the machine's panel, and it would render back a code and state that a certain part is needed. Parts often take a week to arrive, which means that machine is out of commission for a week, and they couldn't necessarily always get hold of the engineer immediately. That wait time impacts all the devices that need testing. The client wanted to be able to connect to a machine and its data, its Programmable Logic Controller (PLC), to work out why machines

were breaking down. Thermotron themselves were very open to fixing the machines but also wanted to know why these machines break when they do.

ADLINK realized that the only way to answer that question is by gathering lots of data. There were external aspects affecting the machine's performance, that the machine could not sense. For example, water is pushed into the machine chamber at a certain pressure. Water pressure and water temperature needed to be monitored, which also required sensors.

ADLINK had to solve, for their client, the following Edge IoT challenges:

- Connecting the test chamber and external sensors then extracting and streaming real-time and historic data to a secure internal database
- Visualizing data easily, on a clear user interface, to enable effective decision-making on maintenance scheduling and requirements

The solution

> *"One of the most important things, when you're shipping lots of data anywhere, is somewhere to store that data. We do have a way of storing some data in transit. You do tend to need a database for that, and that's where InfluxDB has been used in most of our applications."*

Why InfluxDB?

ADLINK chose InfluxDB because it met their storage and performance requirements. ADLINK customers, especially in Europe, are quite sensitive about data privacy and don't trust the data going to an enterprise cloud service. The only way to overcome that concern is to store that data in a database that can be installed locally and that can handle multiple data streams in a highly performant way. That's exactly what InfluxDB does.

When they do product demos or leave demo kits for customers to play with in a digital experiment or a proof of concept, they need that data to be stored securely and efficiently on their IoT Edge device, ADLINK Edge™. This enables them to use the data in any way they see fit. So, it made sense to use InfluxDB for their solution.

ADLINK's solution uses InfluxDB to ingest monitoring data, which is time series data. And since data visualization is the first step to deciding what data to retain and start filtering that data based on your needs, ADLINK's solution connects InfluxDB with Node-Red and with Grafana for real-time and near-real-time visualization respectively.

To understand ADLINK's solution for their client, it is necessary to introduce ADLINK Data River and ADLINK Edge™.

The value of the ADLINK Data River

The ADLINK Data River is the technology they built to transport data in actual real time, from whatever device is connected to the Edge. Once the data makes it to ADLINK's Data River, that transformed data can be pushed to wherever it needs to go.

Introducing the ADLINK Data River:
From Data to Information to Value in Real-Time

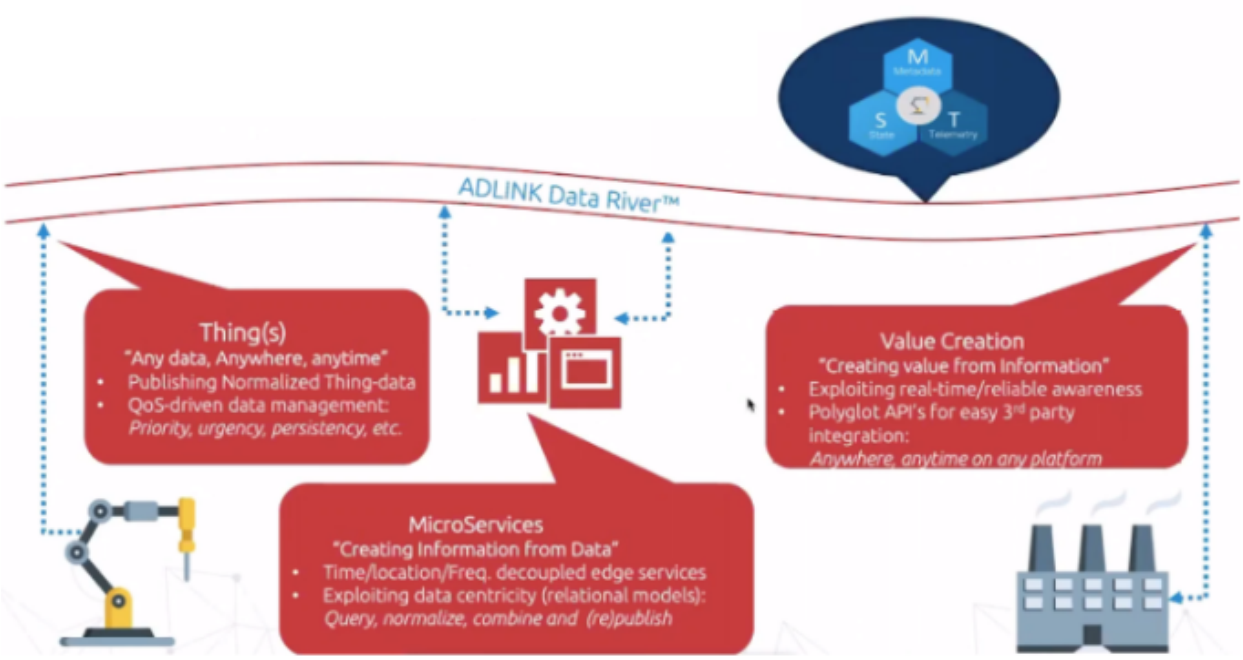

ADLINK Data River delivers value because it is:

- **Easy to integrate** - Data-centric architecture that hides topology details, enabling true plug-and-play
- **High-performing** - Latencies as low as 30 µsec, throughput of millions of msgs/sec
- **Scalable** - Scaling across thousands or even millions of participants

- **Secure** - End-to-end secure data connectivity with authentication, encryption and access control capabilities
- **QoS-enabled** - Full control over data distribution: timeliness, prioritization, reliability and resource usage
- **Open** - Based on the mature, proven and open Data Distribution Service (DDS) standard with future-proof APIs plus wire-protocol
- **Fault-tolerant** - Peer-to-peer communication without message brokers or servers
- **Widely applicable** - Available for embedded, mobile, web, enterprise and cloud applications

Because of ADLINK Data River's low-latency throughput (since it's a publish-subscribe model), the data is only published once, and any number of nodes can subscribe to see that data securely.

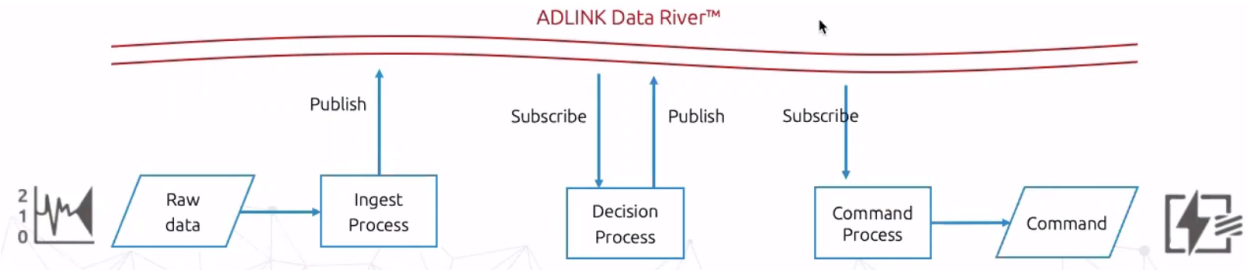

ADLINK can ingest real data in any form (such as Modbus data, MQTT or Canvas). As soon as data is normalized and pushed to the Data River, it can then be pushed to any node in any format needed.

ADLINK Edge™ makes IoT simple

Connecting to the Data River is ADLINK Edge™: a unique mix of edge hardware, data connectivity software, and services to get data moving in the correct format to whoever needs it and where it needs to be, precisely when it needs to be there.

ADLINK Edge™ can be defined as follows (each of the below items in the formula can include a variety of options and solutions):

Edge hardware + Data connectivity software + Services = ADLINK Edge™

ADLINK Edge™ releases the power of operational data to optimize operations and transform business by enabling AI at the edge.

ADLINK EDGE™ Capabilities:
Comprehensive Software Architecture to Rapidly Develop Edge IoT Solutions

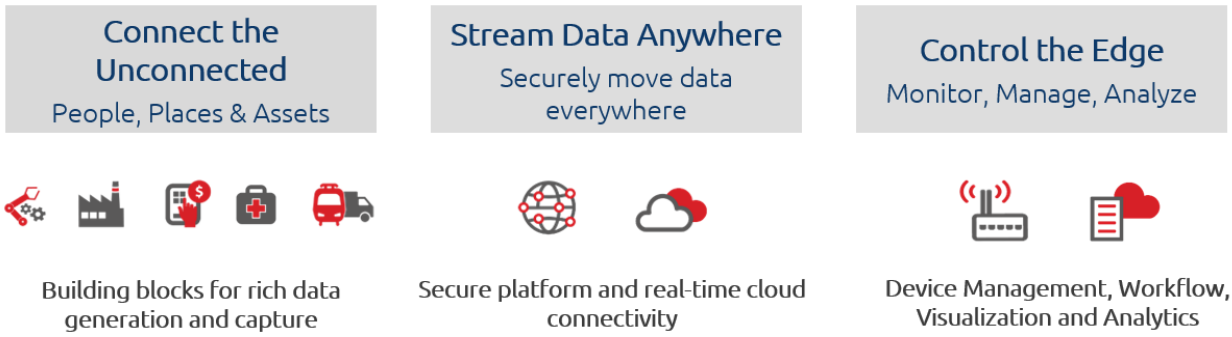

With no programming necessary, ADLINK Edge quickly connects previously unconnected operational equipment, sensors and devices.

How ADLINK Edge™ Works

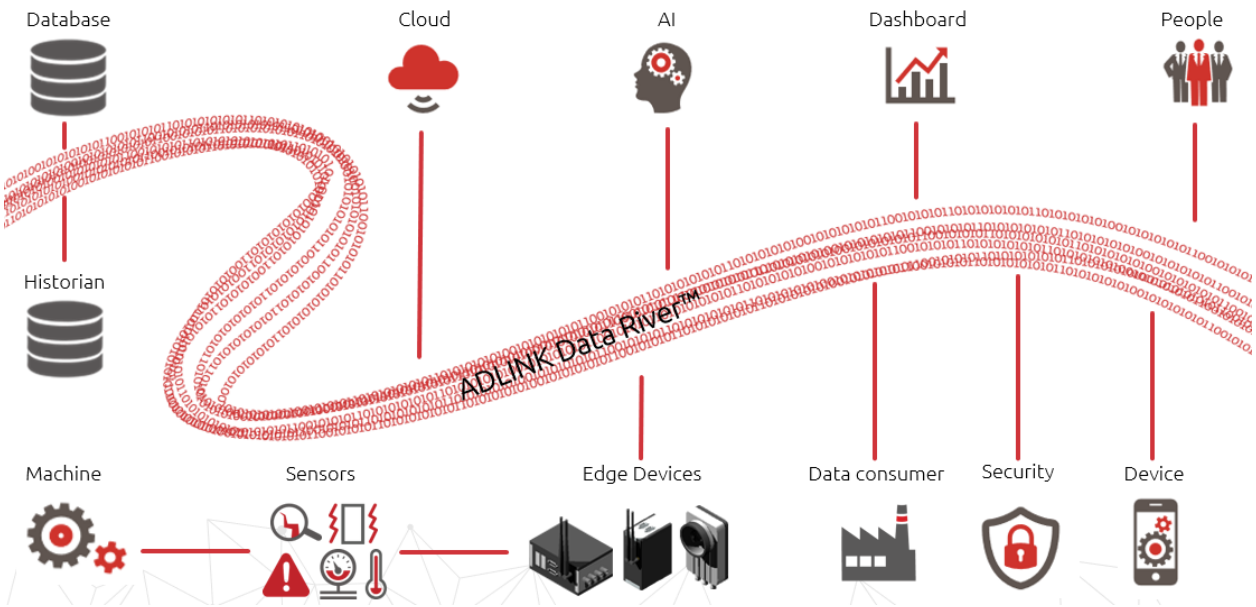

Once ADLINK EDGE is connected to the ADLINK Data River, data can flow freely northbound or southbound, to any cloud analytic platform database and even east to west between devices to control operations at the edge. The ADLINK platform is a platform with ever-expanding services and is capable of reading data from various sources and pushing it to IoT Edge. They transform and interpret that information then push it to multiple streams to display live and historic data through dashboards and enterprise-type services as required.

Overview of ADLINK's solution for their client

Using their Data River technology and IoT Edge device, ADLINK set out to solve their client's challenge. They added sensors to the machine chambers and to a pipe to monitor the water going into the machines from a pumping station was over half a mile away. This enabled them to monitor what's going into the machine and monitor the PLC controller within the machine itself.

Within a few days, they were streaming data from the machine, and at the same time, from the external sensors. When connecting to a machine or site, whether in the field or factory, the need arises to visualize data in order to make it meaningful. ADLINK provided dashboards so they could see a certain temperature of the chamber, when it's ramping up or down, and a lot of other stats which were identified as being important. The engineers available on that shift could see the key metrics and key tags that were pulled out from that PLC visualized on a screen.

ADLINK added the water pressure and water temperature to the dashboard as well to provide comprehensive visualization. Now, if the machine breaks, and somebody wasn't looking at the dashboard, the first thing they need to check is what time frame it broke down in. The only way to do that is to look retrospectively, even in the case of near-real-time data. This is why ADLINK needed to collect data in a time series database, and InfluxDB was perfect for that.

ADLINK set up a secondary Grafana dashboard, but this was in near-real time (running with a delay of about 10 seconds). They used Grafana so they could visualize and then batch that data to see what was going on for the key metrics at that certain point in time.

They can also extract data and put it into management reporting. Upon analyzing the data, they started to understand why a given machine is failing. By capturing data in real-time, they could see machine performance live and were able to reduce maintenance time.

How ADLINK visualizes customer data to solve customer problems

ADLINK's small, powerful gateway (at the center of the image below) contains an Atom processor, and some memory SSD storage which is used to both store and run their Edge software solution.

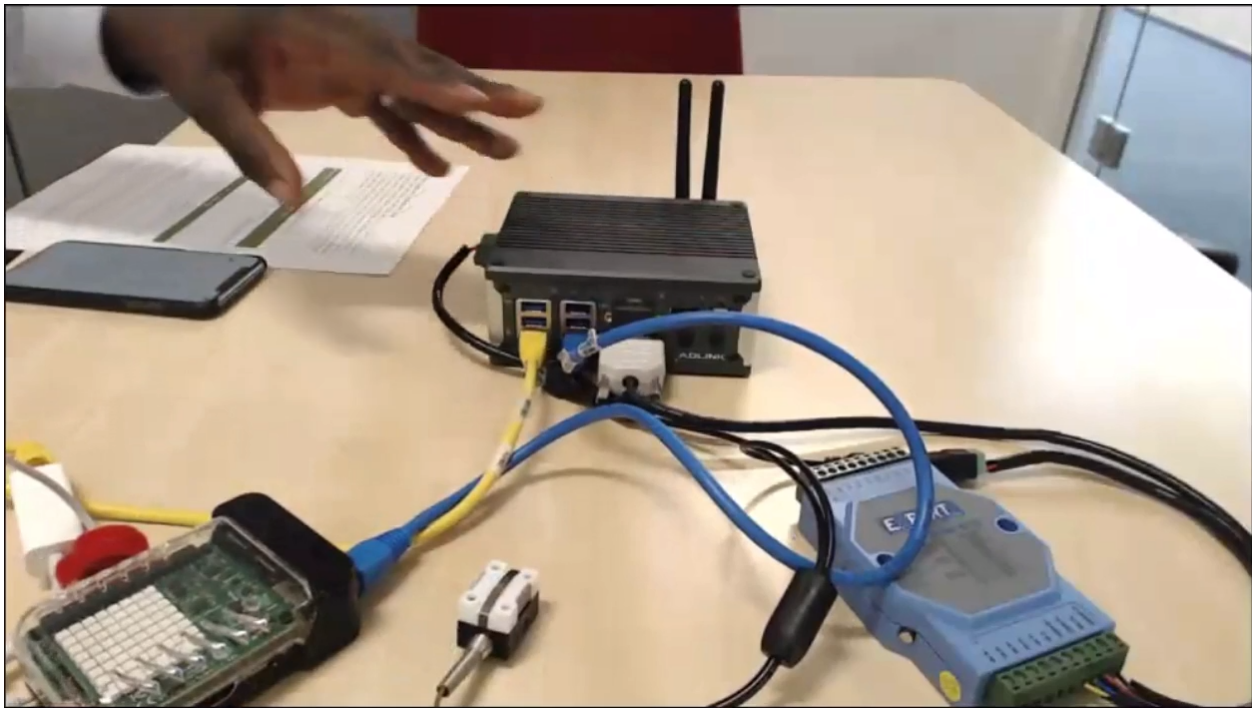

- The gateway is attached to the temperature sensor through the Data Acquisition Module (DAQ).
- The DAQ connects on RS485 to the ADLINK gateway using the Modbus protocol, still the protocol of choice for most machines out there today.
- Also attached to the gateway through the Ethernet port is a Raspberry Pi that runs the OPC-UA protocol.
- The data is generated using Sense HAT (an add-on board for Raspberry Pi that enables it to sense the world around it) to produce pressure, humidity and temperature as a value that can be interpreted and used.

A variety of services can be installed and configured on the gateway. All the services shown below run locally on the gateway and are used for configuration or dashboarding purposes.

Sample Screenshot of ADLINK Edge Software UI

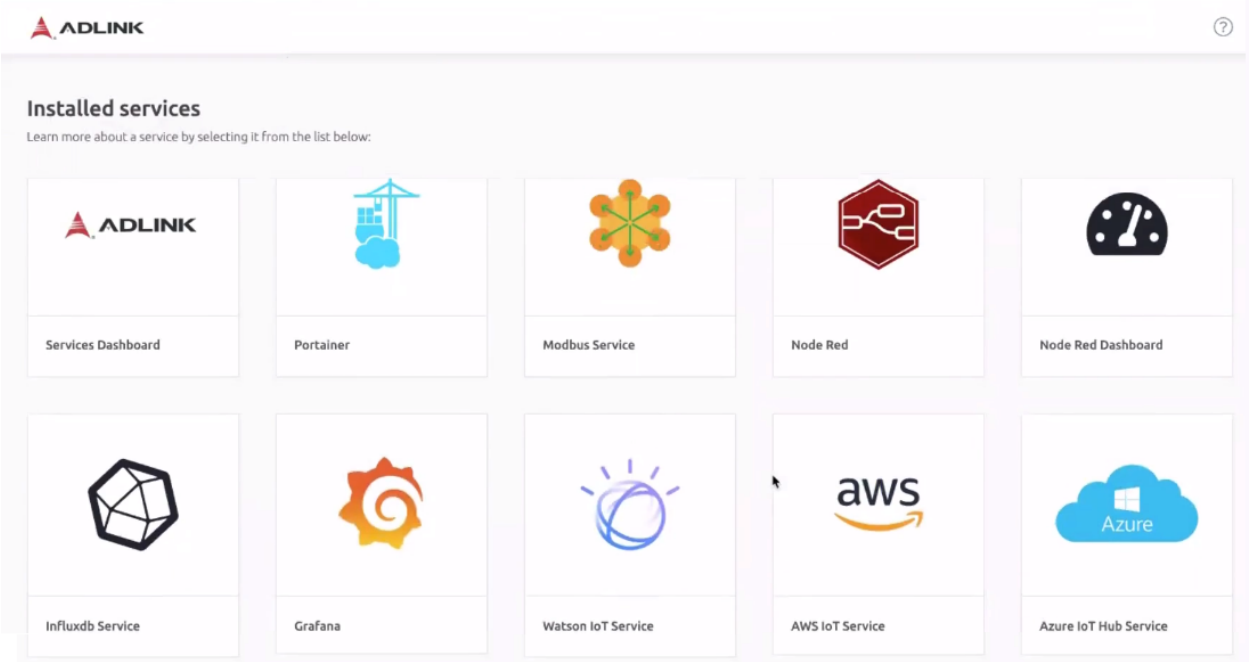

Node-RED, which is built into ADLINK's solution, is used to transform data and generate a dashboard with its built-in capabilities.

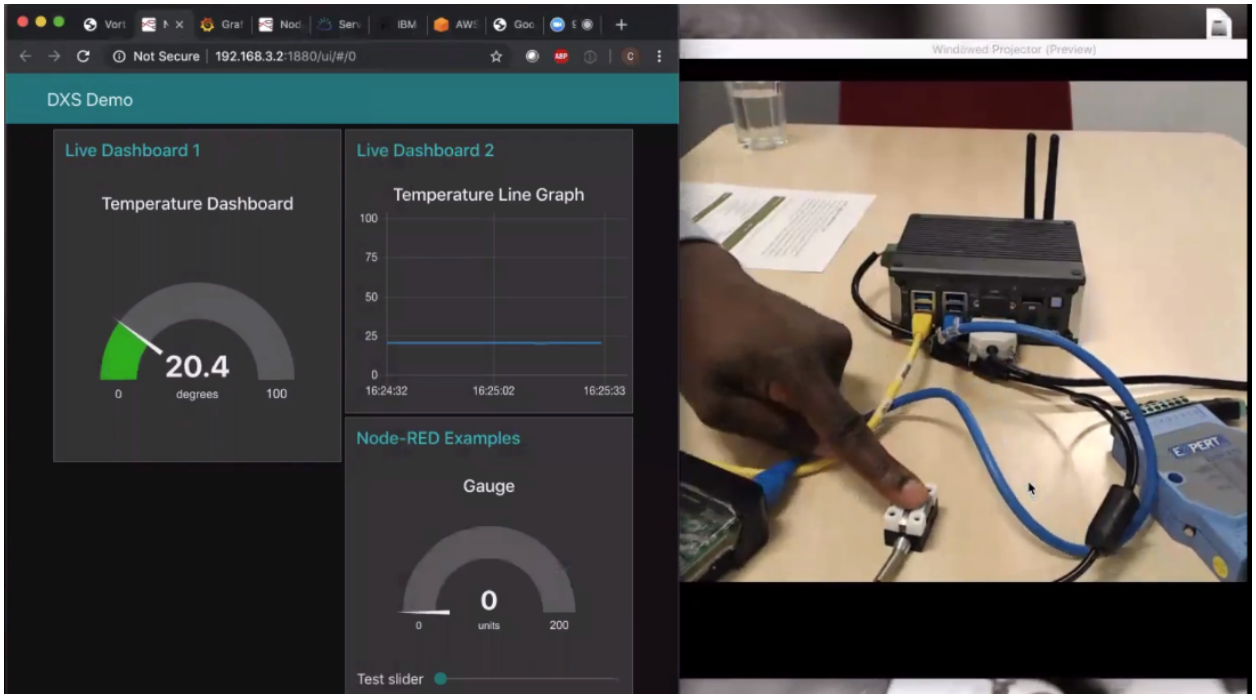

ADLINK can stream data in real-time. They can read that same source of information and then send it simultaneously to a database to be published in near-real-time in a Grafana dashboard, as shown

below. In this instance, Grafana is connecting into InfluxDB to display that near-real-time and historical data.

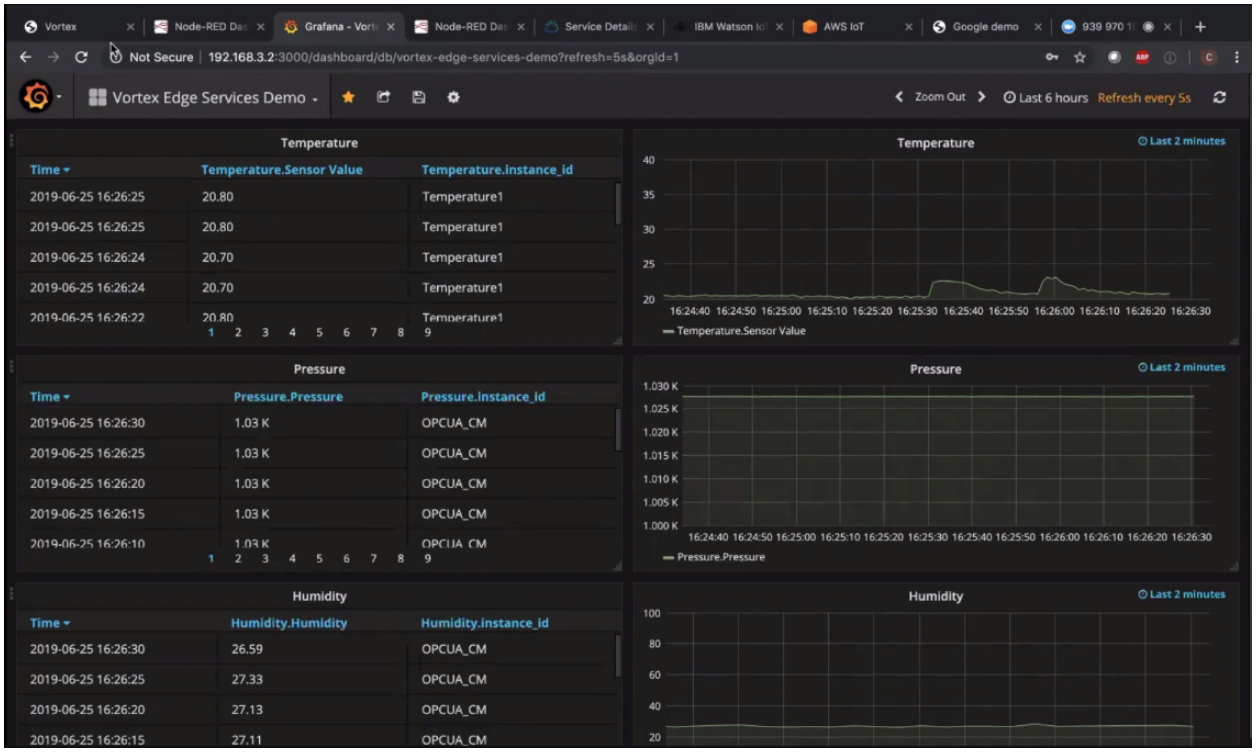

They have a combination of data sources with different protocols displayed at the same time in near-real-time on a dashboard, being streamed through InfluxDB.

Streaming live data and controlling the edge remotely

They can stream live data from multiple sources, using their Data River technology. The same published event on their Data River is being subscribed to by Node-RED service and InfluxDB Service, and a Grafana service through InfluxDB. They read the data once and display it multiple times or push out data into multiple streams at the same time. It's an extremely efficient way of doing things.
If companies with branches in different locations (such as ADLINK's oil and gas customers) want to see that data, the data can be pushed to a remote station.

ADLINK stream raw data as it is, such as in the example below showing humidity data streamed to a live enterprise cloud service (IBM Watson IoT Platform).

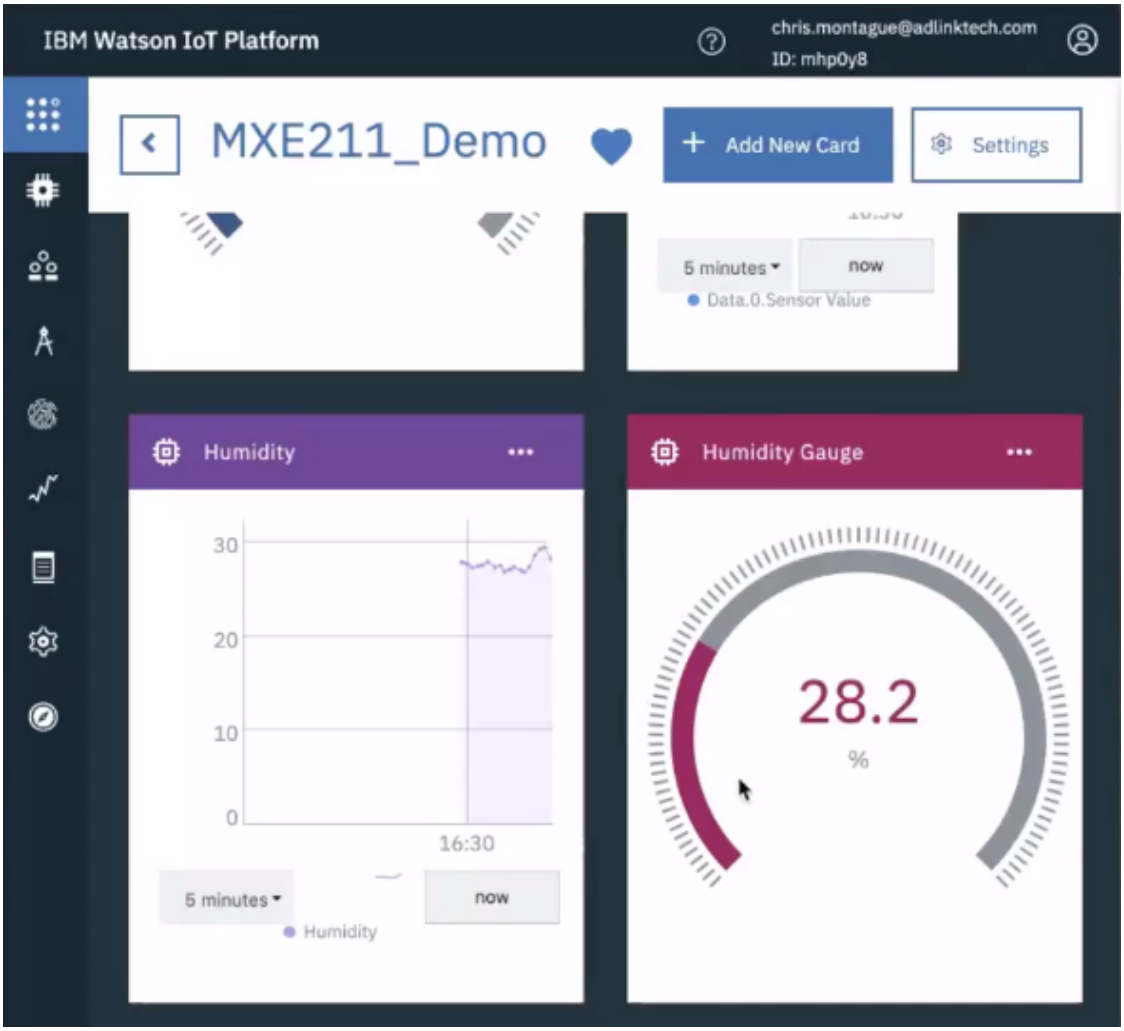

They can stream data in real-time to multiple nodes, internally and externally, reading data once. They can also consume different types of data from their Data River simultaneously.

If it's not possible to stream locally or if connectivity is limited internally (as in a factory), the Wi-Fi card on the back of their gateway can connect to an IP address. This provides a remote version of their dashboard running in Node-RED. The local and remote dashboards are in sync (for example, displaying the same change in temperature as shown below).

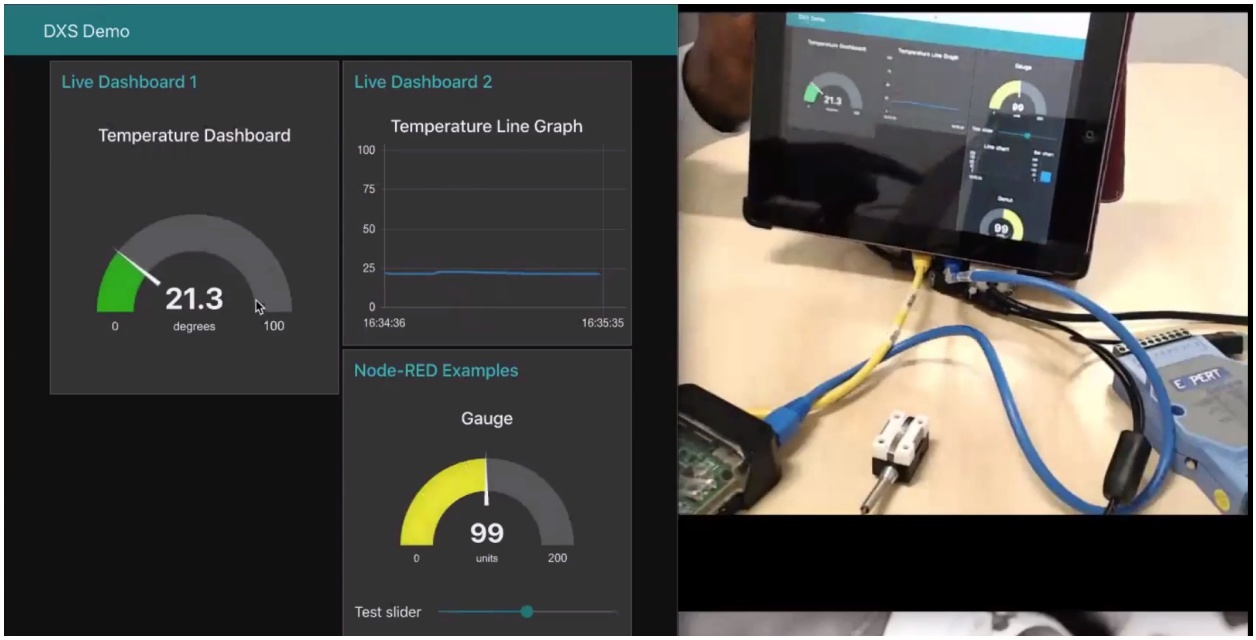

ADLINK can stream data in real-time from the machine, and if they want to push a command remotely back to an industrial machine (for example in a hazardous industrial environment), the machine can be switched off without going into a room, using a slider on the remote desktop.

ADLINK's solution shows how ADLINK can connect the unconnected using all types of connectors (serial and digital), stream the data anywhere, and control security at the Edge.

Technical architecture

> *"To enable visualizing the data, you also need to enable storing data. You need to be able to extract that data to do something with it, whether to push it to an enterprise service, or use it for some form of reporting, or both."*

InfluxDB Time Series Database Central to ADLINK's Solution

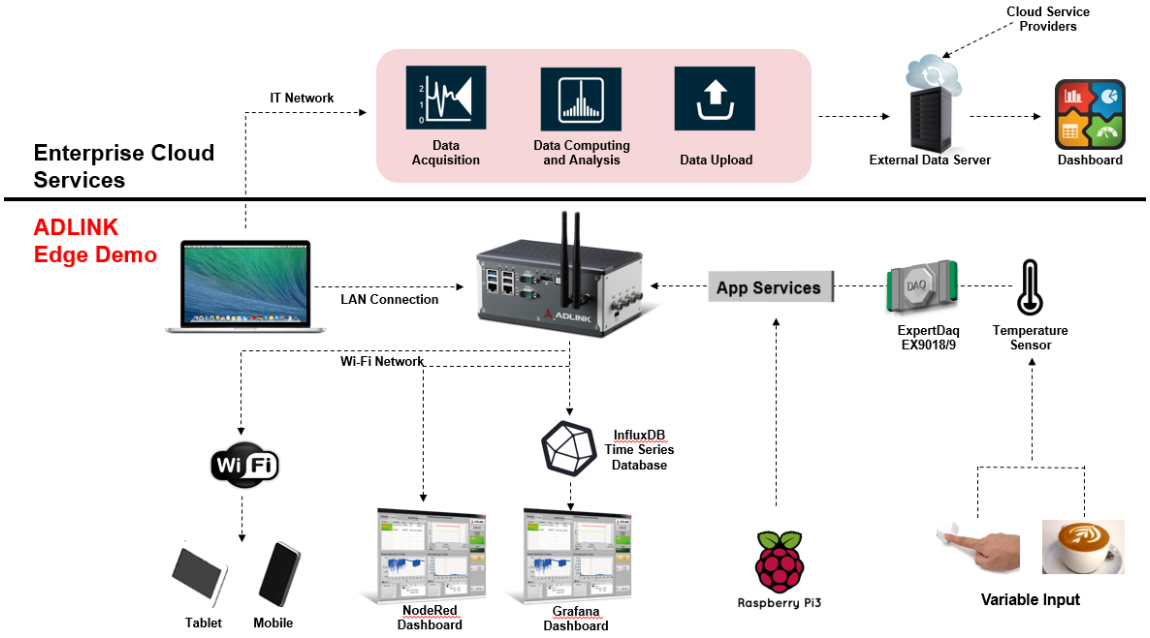

The above conceptual architecture diagram shows the dataflow in ADLINK's solution:

- A temperature sensor detects variable input.
- The temperature data is pushed through a digital unlock converter (ExpertDAQ) and picked up by one of ADLINK's app services which runs on their gateway.
- The data for humidity and pressure is pushed using OPC-UA for App Services from their Raspberry Pi and is then displayed onto a Node-RED dashboard.
- Sensor data is then displayed onto a Grafana dashboard through InfluxDB.

ADLINK's actual solution for their defense contractor client has some of the sensors bolted into a pipe given the distance between the testing machines and pumping station.

The below diagram further shows how information is pushed into the ADLINK Data River and the role performed by InfluxDB and ADLINK Edge in that process.

Enterprise and On-Prem Services in ADLINK's Solution

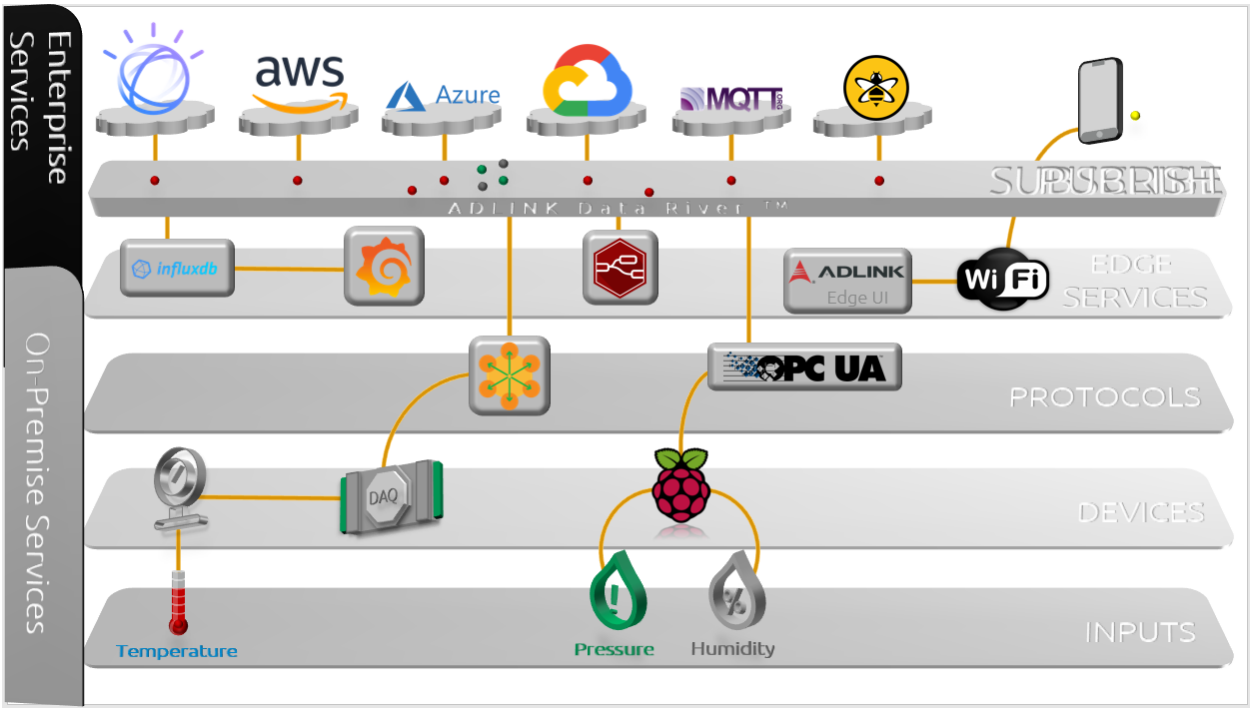

- Temperature data is pushed through the DAQ (Data Acquisition Module) to the Modbus protocol and then onto the ADLINK Data River. Once that data is published on the ADLINK Data River, they can then push it back through InfluxDB to Grafana dashboard onto Node-RED directly.
- Raspberry Pi pushes water pressure and humidity data onto the Data River. That data is consumed by a few InfluxDB nodes they've configured to be subscribed to, and is then pushed to Grafana and IBM Watson.
- The Edge UI allows connecting through Wi-Fi and pushing data back to the Edge as well.

How ADLINK's solution uses InfluxDB

Under the hood, they use a containerized platform. Each of their apps is a Docker container. The screenshot below shows InfluxDB running.

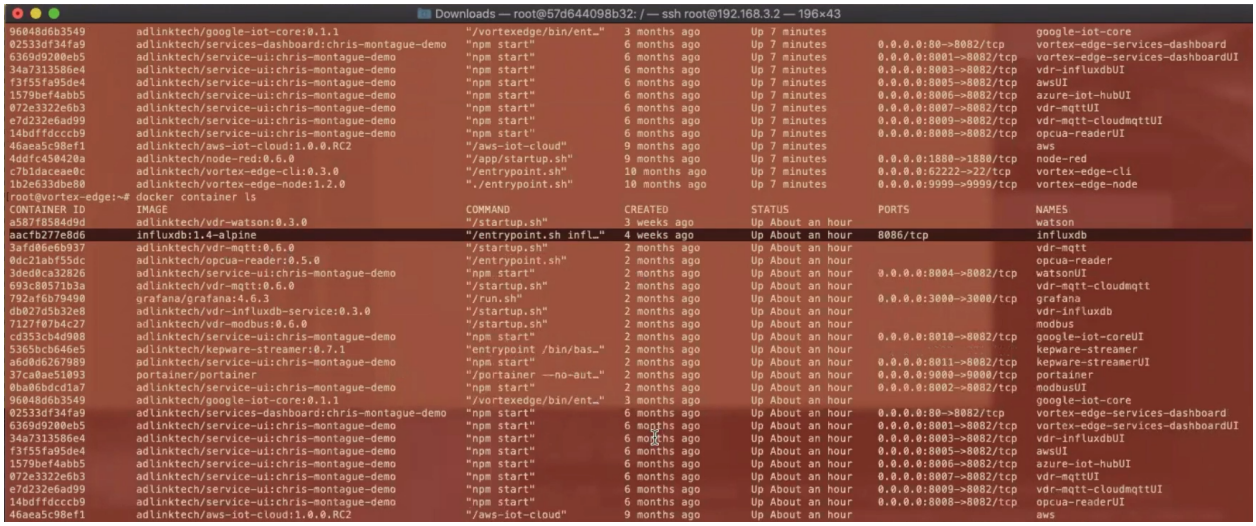

They have written their own connector, which they call “InfluxDB Service”, to connect to InfluxDB and stream the data onto their platform.

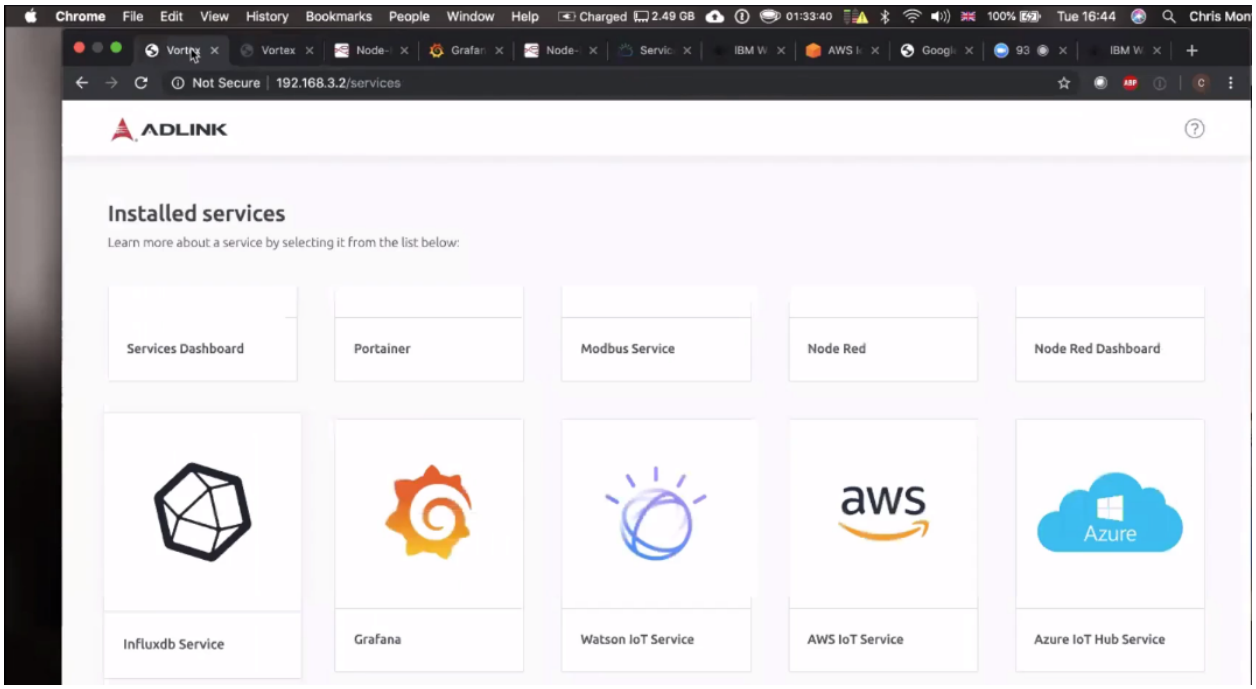

How ADLINK gets data from certain services

The InfluxDB Service is accessible through the ADLINK Edge UI. This service is a local host because they have a local network set up for all their Docker containers. They can subscribe to a number of different topics depending on the type of data. They set it up that way because ADLINK’s underlying technology is called Data Distribution Service (DDS). It's a military-standard way of shipping data from A

to B and back again. It is based on a publish-subscribe model. There are a number of topics (for temperature and other variables).

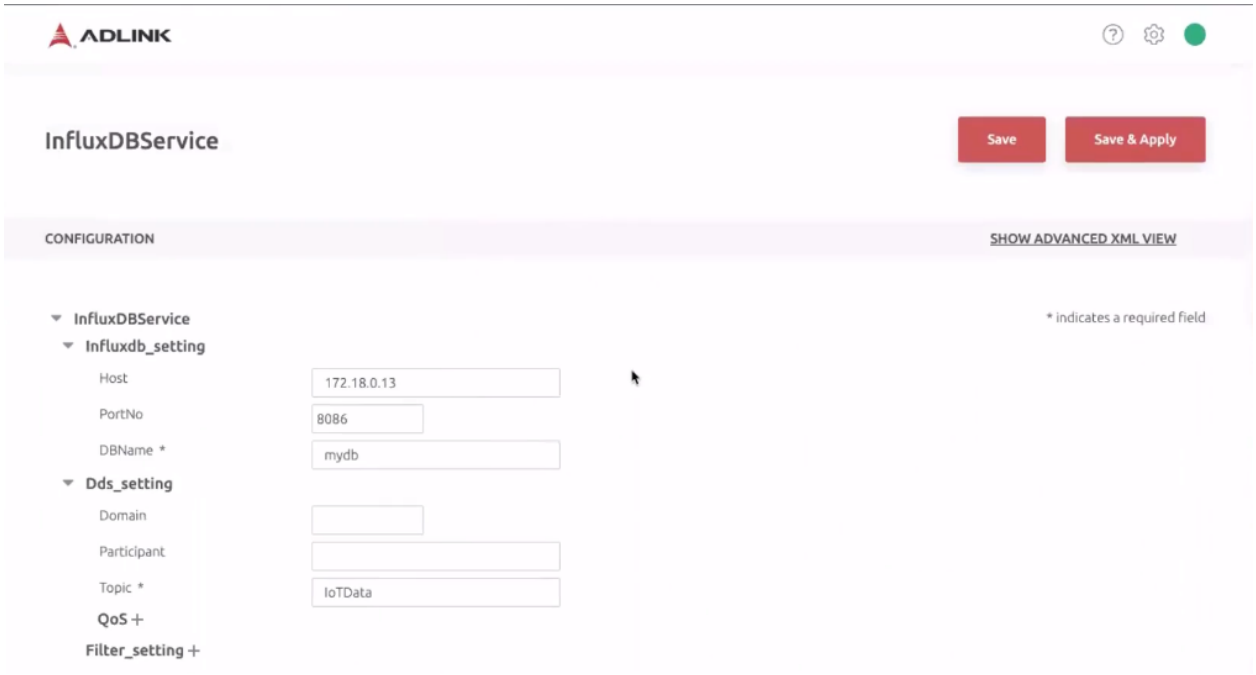

On the InfluxDB Service configuration screen, the relevant values are entered in their respective fields, as shown above.

Below is the simple XML configuration behind the “InfluxDB Service” written by ADLINK.

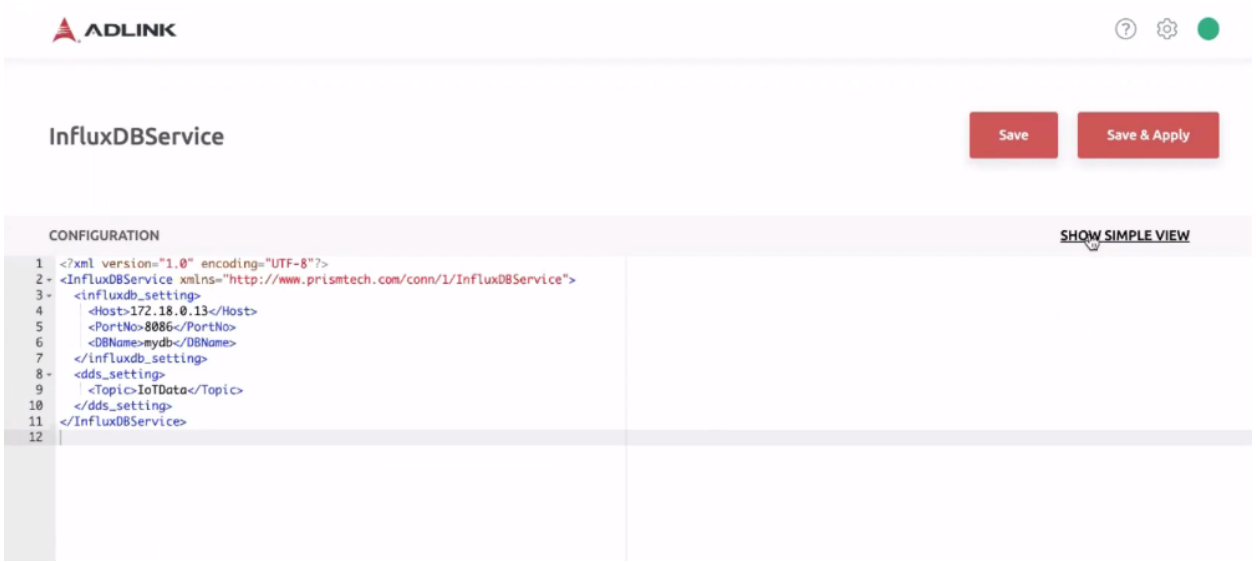

ADLINK get data from certain services through the Modbus Service, which consists of a serial port connection using Modbus RTU. There are fields where the values for Data_bits, Stop_bits, and Serial_mode can be entered. This is where InfluxDB pulls the data and the tags it understands — and where ADLINK enter the values they programmed into their InfluxDB Service. In the Modbus Service screenshot shown below, the values for “Identifier” and “Entity” are “temperature1” and “temperature”.

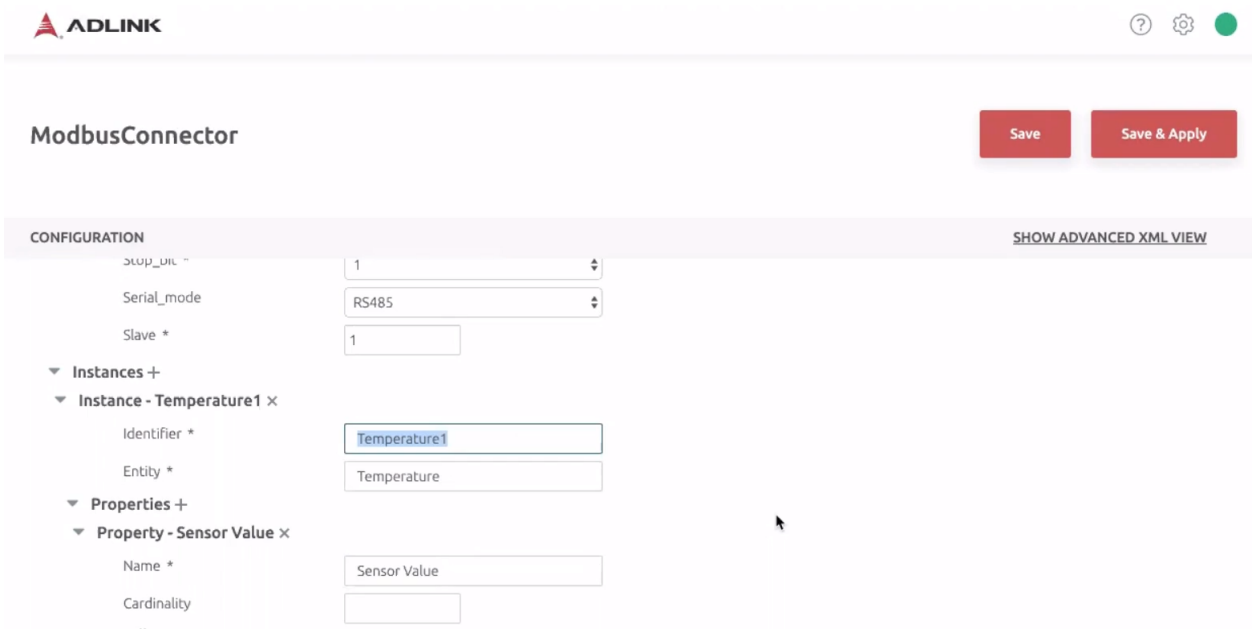

Through the same interface, they can specify the desired poll rate in nanoseconds, and can poll at a very high frequency to get a temperature value over a serial protocol.

By dropping the InfluxDB Service into the InfluxDB container, and using the SHOW SERIES command, they can show the data which they've configured (“temperature” and “temperature1” for instance).

For OPC-UA, they similarly configure the OPC-UA connector from their UI.

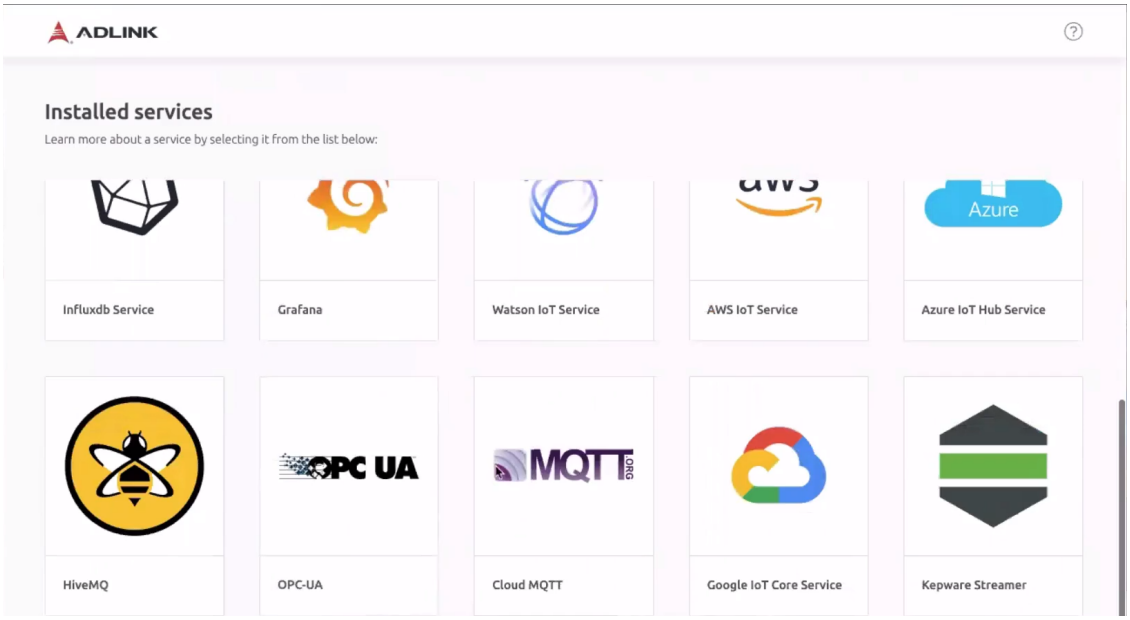

There, the values for pressure and humidity are set, as shown below.

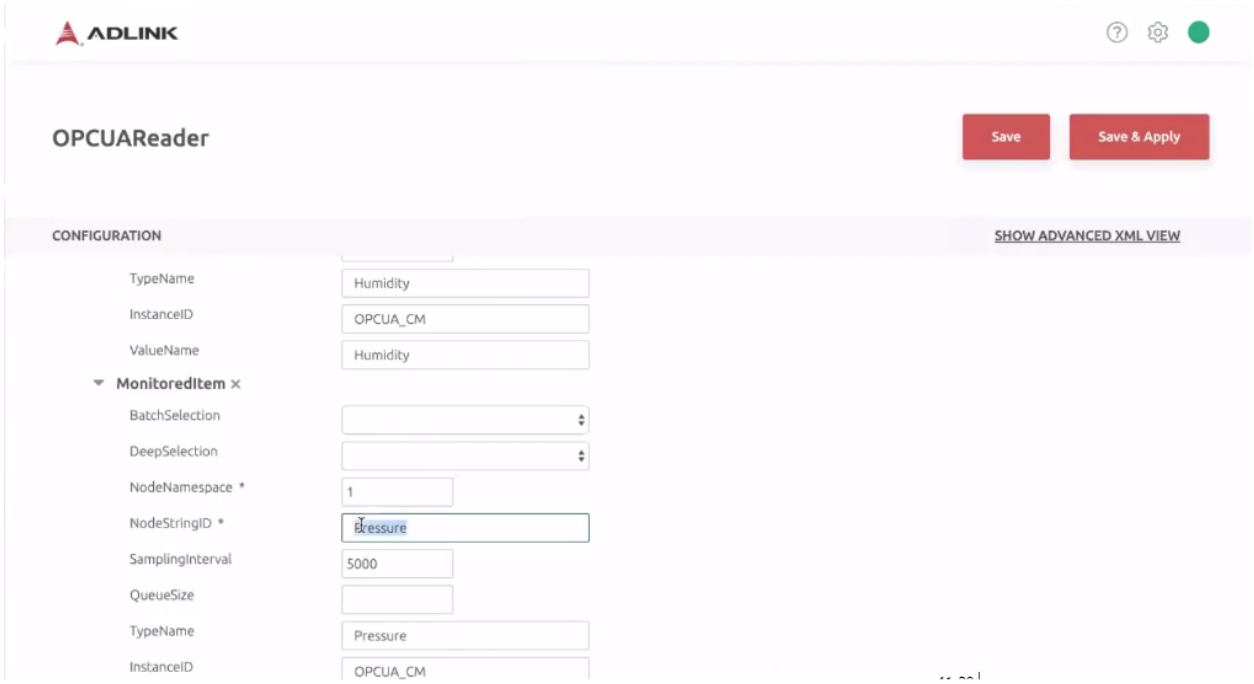

The “InstanceID” that is set above matches what shows below. So do the other values entered above. That's how ADLINK get data into InfluxDB.

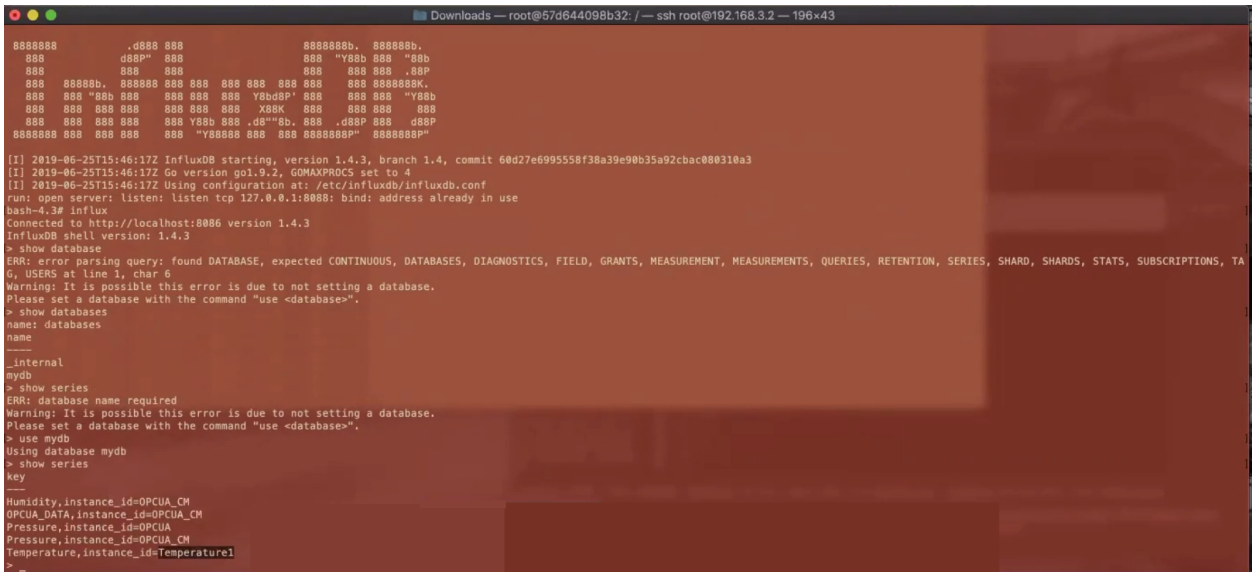

They can stream data into InfluxDB in real-time from multiple desired sources, regardless of where it's coming from, what type of data it is, what rate it's coming at, or if they need to sample it up.

Results

> *"The only way to diagnose any issue is you need the data collected in a way that you can monitor and analyze it. And the best way to do that is in a time series fashion."*

Companies around the world are looking to ADLINK for leading Edge IoT solutions and guidance, and for finding real ROI at the Edge.

For example, ADLINK is putting its Data River technology and InfluxDB to solve challenges for another client: a large drinks manufacturer, working within the UAE, that supplies soft drinks for the entire UAE region. When it gets hot in the factories there (which don't have interior cooling), the plastic wrap that goes around the bottle stretches and causes the machine to jam. When the machine jams, production on it stops for 40 minutes. When it breaks — for example, for that minimum of 40 minutes, and then takes an hour for that machine to come back online — all of a sudden, they've lost a run of 50,000 bottles. In a low-margin business, that loss is huge.

To increase their efficiency, the client got around it initially by putting their big wrap in cold storage. This meant that, instead of breaking at 12 o'clock, it broke at 1 o'clock. They bought some time but didn't

understand the exact issue: what's the maximum temperature you can run these machines at before the plastic wrap starts to stretch and machines start to jam?

ADLINK has installed sensors on the soft drinks machine. Unfortunately, the machine has to break so they can work out what thresholds it can handle. But they're monitoring that and collecting its data to be able to control that machine, and even slow it down or pause it, depending on the speed that the plastic is being pushed through it and the resulting temperature generated by the machine.

The other problem ADLINK has successfully solved for this client is replacing manual data collection with automated data ingestion. The client had employees looking at counters of how many bottles are being produced. The numbers they enter in their spreadsheets are always a rough estimate because there are logging issues as well. ADLINK solved that issue for them just by collecting data using InfluxDB, thereby providing the added benefit of solving their logging issues and also freeing employees to focus on other tasks.

Powering its solution with InfluxDB, ADLINK is fulfilling its mission to “effect positive change in industry by connecting people, places and assets with artificial intelligence through the delivery of leading Edge, robust and reliable hardware and software solutions that directly address mission-critical business and technology challenges.”

What’s next for ADLINK?

ADLINK plans to use the built-in dashboard capability of InfluxDB 2.0 to show live data, which would enable them to tune machines in real time. This is because:

- Though Grafana allows for complex dashboards, when you have a lot of time series data, it starts to slow down considerably when handling large volumes of time series data. It is also not live (has a five-second delay in showing data). One solution they attempted was to clear data regularly, but while it’s updating, Grafana also takes up processing time.
- While they use Node-RED to get live visualizations and while you can connect to a database with Node-RED, it's not as easy for a customer to do it because it requires some customization. With InfluxDB 2.0, you get visualization in real time out of the box.

They also plan to develop forecasts with the time series data that they're collecting (such as with the Thermotron example above), to let clients know when a part might need to be replaced and further enable predictive maintenance.

About InfluxData

InfluxData is the creator of InfluxDB, the leading time series platform. We empower developers and organizations, such as Cisco, IBM, Lego, Siemens, and Tesla, to build transformative IoT, analytics and monitoring applications. Our technology is purpose-built to handle the massive volumes of time-stamped data produced by sensors, applications and computer infrastructure. Easy to start and scale, InfluxDB gives developers time to focus on the features and functionalities that give their apps a competitive edge. InfluxData is headquartered in San Francisco, with a workforce distributed throughout the U.S. and across Europe. For more information, visit influxdata.com and follow us @InfluxDB.