White Paper

What and why - ISA95 is dead, what next?

“... the model has not followed modern technology… **The leading candidate would be intelligent data pipelines.”**

Daniel Riddoch
Date of release 24/01/2025

What do we mean?

ISA95 is a popular derivative of the Purdue model. These models aim to describe the manufacturing process and divide responsibilities between production, operations, IT, and business management. The stratification of responsibilities has several advantages.

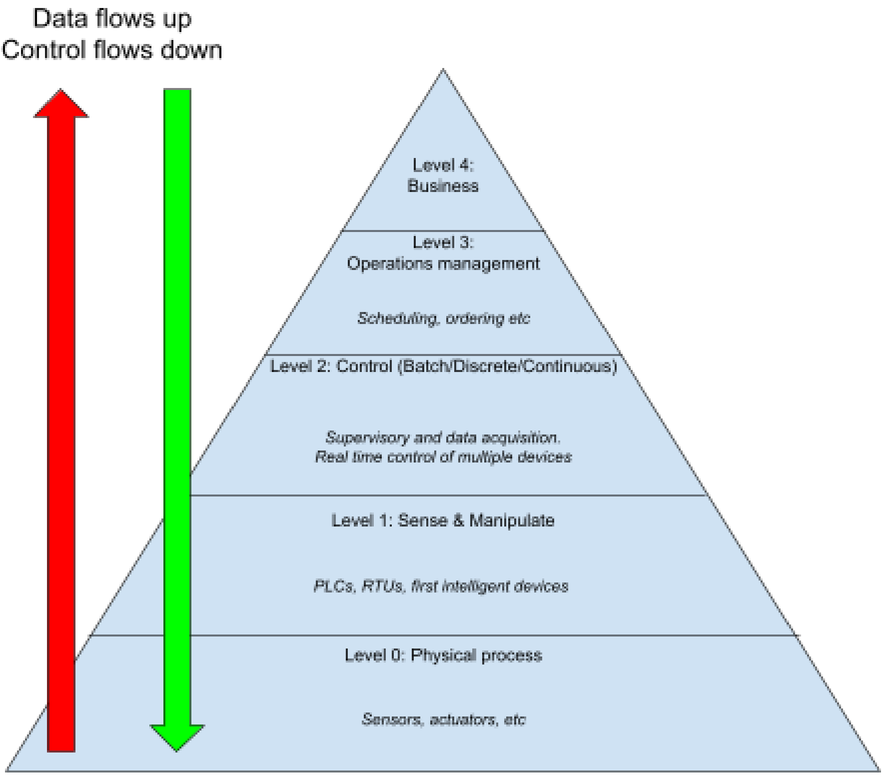

Figure 1: A sketch showing the basic principles of the ISA95 Purdue model

Firstly, it allows specialists in each area. Having specialists working in operations for example, allows the processes to be optimised and for any developments to be monitored by dedicated staff, keeping the processes state of the art.

Secondly, the division of responsibility creates clear lines of accountability, meaning that people within the structure are clear about their role and are thus empowered to take ownership of their part of the process.

Finally, this model is conventional. Over many years, this has come to be an industry standard, and as such, there are lots of resources available, and lots of accumulated knowledge. Furthermore, new hires are likely to be familiar with the system and able to adapt quickly.

How does the problem arise?

Being the industry standard does not mean that this method is infallible, so let us re-examine those advantages from a more critical standpoint. Firstly, by specialising in each area of responsibility there is a risk of becoming blinkered. This means that workers in each section become less aware of the issues faced in other areas. Consequently, a lack of teamwork can lead to inefficient solutions, missed opportunities for innovation, and stifled growth.

Secondly, although clear lines of accountability are created, it will not always be practical to measure the results in such an isolated manner. For example, should a reduction in waste products in a certain process be attributed to extra care taken in the production stage, or a change in operational thresholds for certain actions to occur, or even a change in raw material supplier implemented by management.

Similarly, by encouraging departments to take ownership of their area can encourage competitiveness between them. As such, decisions are perhaps not always made in the best interests of the company. For example, if given a target to reduce wastage, operations and production may extend times between machine cleaning and/or maintenance which could lead to a drop in product quality or expected machine life.

Whilst it is true that the model is conventional, this also means that it has been force fitted to many applications. For example, is this solution the best fit for a digital manufacturing business that utilises many smart machines, where the lines between IT, OT and production become blurred?

Finally, it is also the case that the model has not followed modern technology. The hierarchical approach can restrict and slow down actions and can restrict insights about the process to higher levels within the system.

DIANOMIC

How can we solve the problem?

So, we now can re-examine what we need from such an infrastructure and control model in the modern context. We require a system that preserves the ability for production, operations and IT specialists to innovate and improve. We also hope to retain understanding of the roles each individual has to play.

However, we can take advantage of the new model to challenge a few of the issues that arise. Several of the issues raised come from compartmentalising the problem. As such, a new model should be developed to allow anyone to understand the whole problem and appreciate challenges beyond their specific remit. This should enable clarity around the impact of changes, not just on specific targets for departments but on a more global outlook.

Finally, a new model should also be future proof. By now it is clear that a digitised industry, along with AI and ML revolutions will dramatically change the manufacturing business. As such, any new model should be compatible with these, and embrace the opportunity to innovate that these revolutions present.

So what is such a model? The leading candidate would be intelligent data pipelines. These are an infrastructure model which allows for all devices to be connected, and facilitates IT/OT development and integration. An intelligent data pipeline is a microservices architecture which allows for plug-ins which can perform many different functions.

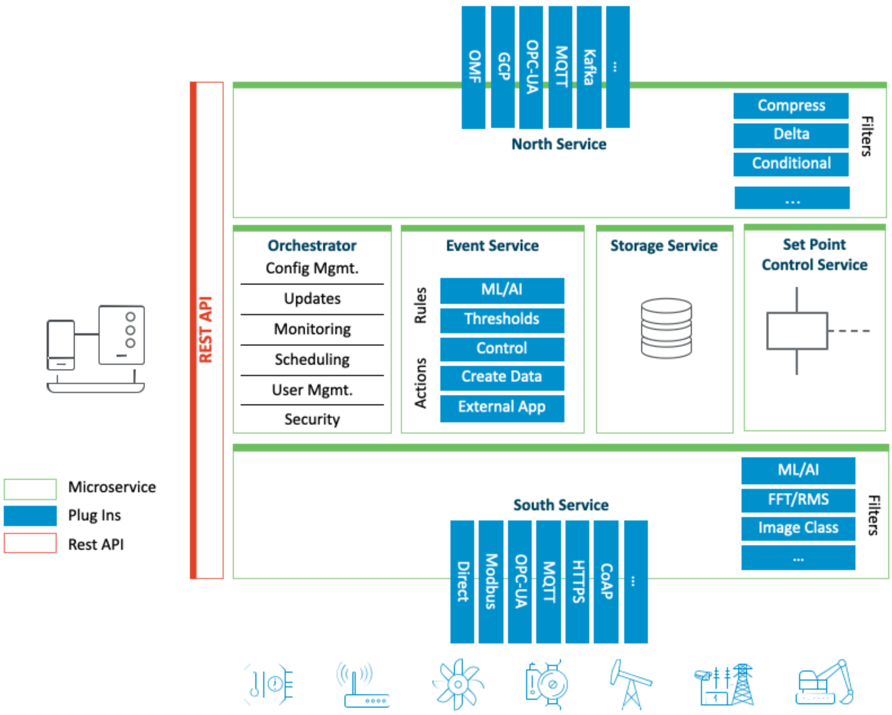

Figure 2: A sketch showing the plug-in architecture of FogLAMP, an intelligent data pipeline system. The architecture allows for flexibility and scalability

The plug-in architecture means that the infrastructure is adaptable to any system configuration and that adaptability is also demonstrated by the fact that pipelines are able to deploy ML/AI models and gain their insights, even on the industrial edge. This gives the platform for innovation in individual sectors and because the system is unified it is straightforward to isolate and display the results of any given alteration.

What other benefits arise?

A data pipeline infrastructure is, because of the microservice architecture, an all encompassing solution. This means that the pipeline begins right at the source of the data(sensors, PLCs, etc) and runs right through to the end of the data flow(display, storage, etc). Consequently security protocols can be addressed within the same device, meaning that attestation data does not need to be sent on, simplifying the data flow and creating curated data more easily.

An end-to-end system also has the advantage of continuity, meaning that a single system can be used to make any and all changes to the IT/OT system and inter-system compatibility will not be an issue. Furthermore, such a system would simplify maintenance tasks and licensing, as this would only require a single point of contact.

Finally, data pipelines are more easily scalable than traditional infrastructure. The pipeline configurations are easily replicated, and multiple pipelines can be managed by a single management architecture, allowing for wider deployment and simpler oversight.

Example

The following example comes from a real use case, where Réseau de Transport d'Électricité(RTE) teamed up with Dianomic and the Linux Foundation to create their next generation control system and data infrastructure using LFEnergy's FledgePOWER, an open source derivative of the FogLAMP suite.

RTE are the French grid operators, responsible for the consistent and reliable delivery of power to over 60 million people and the seventh largest economy in the world, across over half a million square kilometres of land, and overland and undersea connections to neighbouring nations and power systems. This daunting task means that RTE has a huge and diverse infrastructure network in many different and often harsh environments.

With such a system, a Purdue model would require all maintenance, condition monitoring, demand monitoring and network management data to be concentrated in a single central location. This requires reliable and stable network connections, something which is hard to guarantee across a large country. Furthermore, the long term connectivity requirements, and the large and distributed infrastructure network requires a lot of brownfield connectivity, or a large and regular Capex investments. This has been accelerated by the development and roll out of IEC61850, a unifying protocol which has been developed by the International Electrotechnical Commission to standardise and model the condition, function and behaviour of substation infrastructure globally.

Figure 3: A sketch showing the arrangement of devices within an IEC 61850 substation(Credit: SGRwin)

Another issue that RTE could face is data volume. Across thousands of substations, if all the raw data that RTE monitors had to be transferred and then stored for monitoring and compliance, the data quantity would be very expensive and the data transfer would likely cause significant network load and incur significant cost. This latency could result in sub-optimal operation, or even non-compliance.

RTE have partnered with Dianomic and are using Fledge, hosted by the Linux Foundation to develop FledgePOWER. FledgePOWER is an intelligent data pipeline, which allows RTE to connect all the devices within a substation. This means that the substation behaves and appears to be IEC61850 compliant to the RTE network. Furthermore, it means that reliable connectivity requirements can be relaxed, as the buffering service allows for connection to be periodically lost and regained without data loss.

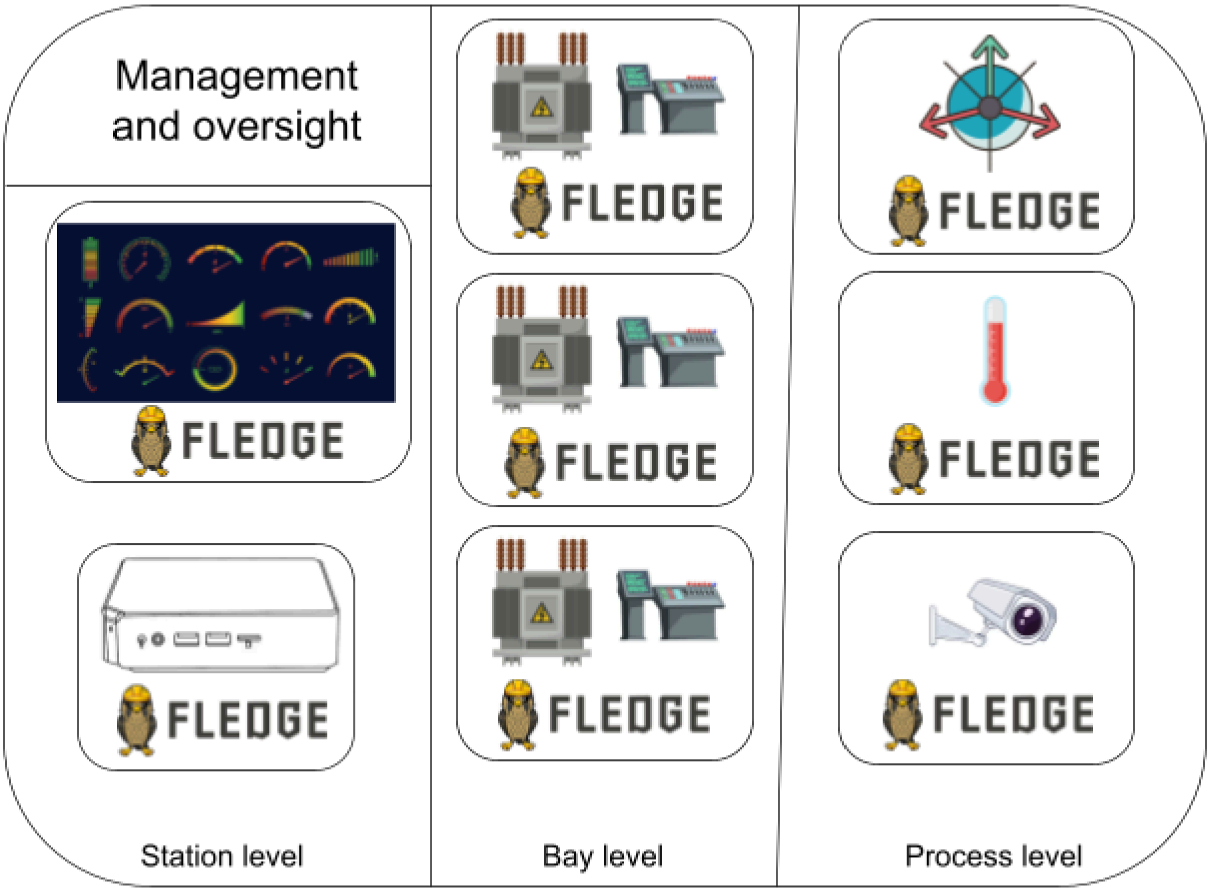

Figure 4: A sketch showing the architecture that can be created usingLF Energy FledgePOWER to facilitate semantic modelling in line with IEC 61850

By combining the data collection and management in an intelligent edge solution, RTE are able to address these issues, because they can subvert the Purdue model. By connecting intelligent electronic devices(IEDs) such as those found in a substation together, insights from this can be developed and deployed locally. Additionally the control functionality of intelligent data pipelines allows for control of the substation to be managed remotely, or even automatically.