OPEN INDUSTRY 4.0 ALLIANCE

THE DATA FOUNDATION FOR MODERN PROCESS AUTOMATION

HOW2 STANDARDIZE DEVICE DATA



SHORT CUT

WHAT

Reliable product and device data is essential for efficient plant operation. This paper shows how technical documents, nameplate information, calibration records, maintenance data, and other lifecycle information can be made consistently available across vendors, systems, and the entire installed base. The goal is a shared digital foundation that gives operators access to the right information when and where it is needed.

WHY

Because fragmented device data slows maintenance, creates documentation gaps, and becomes harder to keep complete and trustworthy as equipment is serviced, calibrated, updated, or replaced.

WHO

For plant operators, maintenance teams, instrumentation engineers, and asset managers working with multi-vendor device fleets who need reliable access to product and lifecycle data.

DR. CARL-PHILIPP DING, GLOBAL PRODUCT OWNER DIGITAL DATA CHAIN @BASF

WE ARE TAKING ADVANTAGE OF ESTABLISHED INDUSTRY STANDARDS SUCH AS ID-LINK, VDI 2770, AND THE ASSET ADMINISTRATION SHELL TO STANDARDIZE AND STREAMLINE OUR DIGITAL PROCESSES ACROSS COMPANY BOUNDARIES.

DOES YOUR TEAM STILL SPEND TOO MUCH TIME SEARCHING FOR THE RIGHT DEVICE INFORMATION?

CAN YOU ALWAYS TRUST THAT THE DATA YOU FIND IS COMPLETE, CURRENT, AND CORRECT?

WOULD YOUR DAILY WORK BE EASIER IF EVERY DEVICE CAME WITH ONE RELIABLE DIGITAL SOURCE OF INFORMATION?

IF YOU ANSWER ONE OF THE Q'S WITH "YES" - YOU SHOULD CONTINUE READING.

CONTENT

01 PROCESS AUTOMATION STARTS WITH A SCAVENGER HUNT

02 HOW CAN THIS BE SOLVED?

03 HOW DATA IS KEPT FRESH

04 LET'S TALK CROSS COMPANY

05 SHOULD YOU START NOW?

Authors

Bartłomiej Ahmida | Marketing Manager, Digital Integration | Endress+Hauser

Marcel Auer | Research Assistant / Doctoral Candidate | IRS / KIT

Thomas Weinschenk | Expert Business Integration | Endress+Hauser

Editor

Lucas Wolf | Technical Content Manager | Open Industry 4.0 Alliance

Publisher

Open Industry 4.0 Alliance
Christoph Merian-Ring 12, 4153 Reinach, Switzerland
https://openindustry4.com
info@openindustry4.com

OPEN INDUSTRY 4.0 ALLIANCE

01 | PROCESS AUTOMATION STARTS WITH A SCAVENGER HUNT

A MAINTENANCE EXAMPLE

A plant moves through six stages before anyone thinks about maintenance: planning, development, construction, commissioning, production, and finally usage and optimization. Also our Members, for example Endress+Hauser, are already involved from the first stage, when a customer is still deciding how big the production line will be and which regulatory requirements the equipment needs to meet. That early exchange is where use cases around instrumentation, engineering, and documentation actually start.

But once the plant is running, the picture changes. Picture a maintenance technician who needs to replace a flow meter. To do that job, they need information from the warehouse team, information from the engineering team, and the technical specification of the device itself. None of these live in the same place. The warehouse gives one part of the answer, engineering gives another, and the calibration certificate is often still a paper document. In regulated industries such as pharma + food & beverage, these certificates must be retained for audits and liability purposes. Even when they are available, however, they are often disconnected from the digital systems used for maintenance. A technician doing this job in a multi-vendor plant can end up carrying five different tablets for five different vendor systems, just to look up information that should take one scan.

This gets worse with scale, not better. Two sensors are easy to keep track of. A hundred thousand sensor endpoints across a real plant are not. Every audit turns up the same finding: the same piece of information has to be manually re-entered into five different systems, because none of them talk to each other. The core problem is not that the data does not exist. It is that it is missing structure, and structure is exactly what does not scale by hand. This does not only apply to the maintenance use case, but to many more.

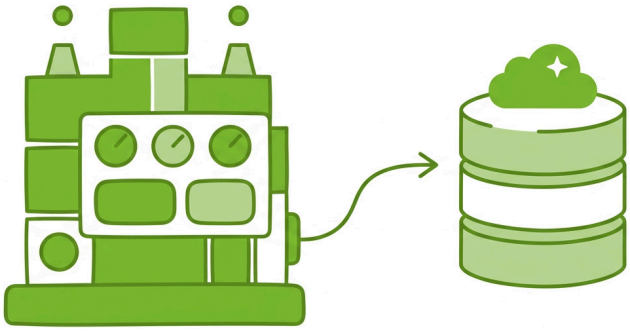


- Instrumentation Engineering
- Identification & Documentation
- Technical Data
- Automated as Built
- Instrumentation Management
- & Many More

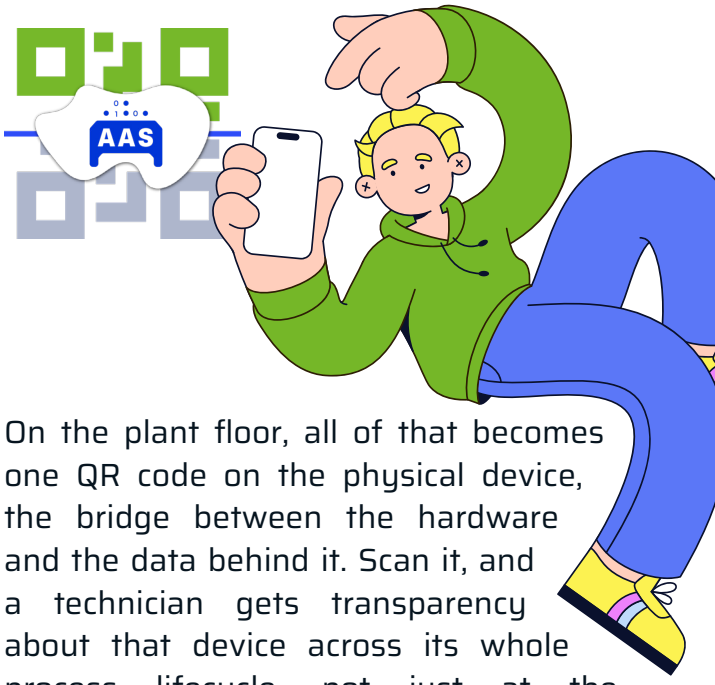
02 | HOW CAN THIS BE SOLVED?

FROM SCATTERED FIELD DEVICES TO ONE STANDARDIZED SHELL

The fix is to give every device one product data repository, for example in the form of an Asset Administration Shell. Instead of technical specifications, manuals, and calibration certificates living in five disconnected places, they live as submodels inside one shell, such as a digital nameplate and a digital calibration certificate, each with its own set of linked properties.



None of this comes from a single company deciding on a format and shipping it. Companies bring their requirements and proposals into the Industrial Digital Twin Association, and the IDTA is the one that actually builds and publishes the resulting submodels, the digital nameplate, the digital calibration certificate, and whatever comes after them. The Open Industry 4.0 Alliance is where those submodels then get tried out against real use cases. It is the playground: companies test what works, what does not, and share the practices that come out of it, instead of every company finding this out on its own. The Digital Data Chain Consortium covers a third piece: getting industry leaders in the process industries to standardize and automate how they exchange data with each other, so systems integrate and run more efficiently instead of each company keeping its own format. Three associations, three different jobs, and none of them trying to do the other two. That is the power of many.

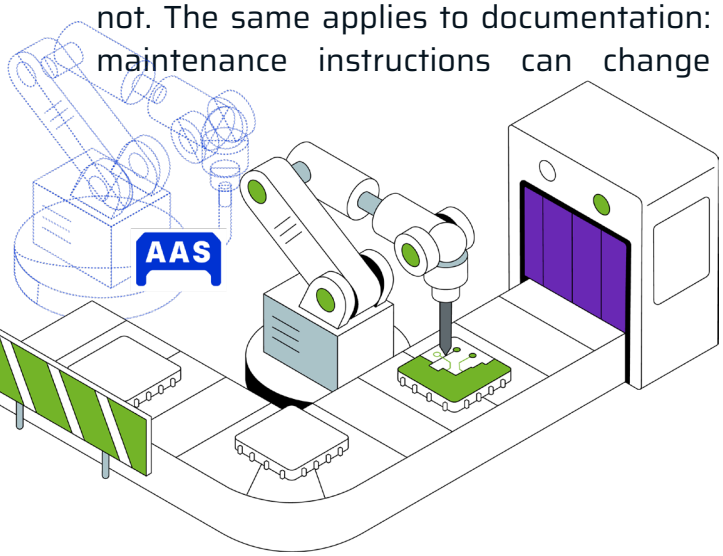


On the plant floor, all of that becomes one QR code on the physical device, the bridge between the hardware and the data behind it. Scan it, and a technician gets transparency about that device across its whole process lifecycle, not just at the moment it was installed. In practice, this already works across three layers of access behind that single code. Scan it with any phone, and you reach the publicly available information already published on the manufacturer's website. Scan it through the connected IoT platform, and you can access not only the documentation assigned to that specific asset, but also additional device-related data and services. Relevant information can then be made available to ERP and other enterprise systems through appropriate interfaces. The code itself never changes. What it unlocks depends on who is scanning it.

But a shared standard is only as good as its ability to stay accurate after the fact. A shell that is correct on installation day and wrong six months later is not actually solving the problem it was built for.

03 | HOW DATA IS KEPT FRESH
FROM UPDATES TO HANDOVER

What the AAS actually is, once it exists for a real device, is a digital twin: not a static file, but a live representation that is supposed to always match what is physically sitting on the plant floor. A useful digital twin is two things at once. It is the information itself, and it is the services, databases, and control logic running on top of that information, the part that actually does something with it rather than just storing it. A maintenance app that pulls a calibration value from six months ago and presents it as current is worse than having no app at all, because it looks trustworthy right up until it is not. The same applies to documentation: maintenance instructions can change



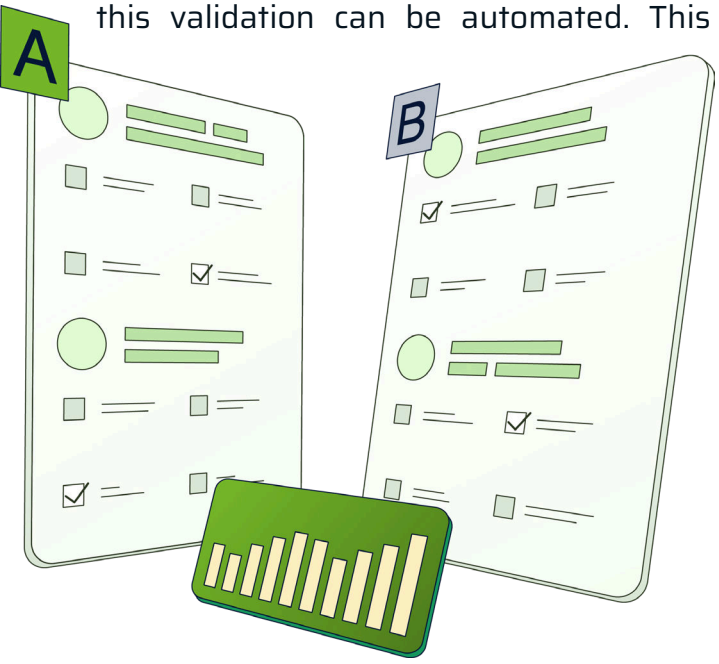
over time, and critical updates need to reach the operator reliably. Handover documentation and change notifications therefore also have to remain up to date throughout the device lifecycle. Keeping that information current is where the AAS splits a device's data into two separate artifacts that have to stay in sync with each other: an asset type and an asset instance. They are not the same thing recorded twice.

An Asset Type AAS describes a device type: the functions, technical specification and properties shared by every unit of that type. Because device behavior is increasingly defined by firmware, this description changes over the product's lifetime. A single firmware release can, for example, add encrypted communication to an entire device family. Each such release requires a new version of the Asset Type AAS. Only the manufacturer can publish it, since only the manufacturer knows what the firmware changed. The previous version must remain available and unchanged, because devices in the field still run the old firmware and their Asset Instance AAS still refer to it.

An Asset Instance AAS describes one physical unit: its serial number, its current parameter settings, and its calibration and maintenance history. Type and instance follow different schedules. The manufacturer releases new type versions, and the operator updates devices during maintenance windows. One device can receive a firmware update today, after which its Asset Instance AAS references the new type version. An identical device elsewhere in the plant continues to reference the previous version until its own maintenance window. This requires every type version to be uniquely identifiable. The AAS metamodel does not enforce this on its own, because references point to an identifier rather than to a version. Manufacturers therefore have to assign a distinct identifier to each type version or agree on an equivalent convention.

The separation also protects operator data. Parameter settings can reveal how a process is run. They remain in the Asset Instance AAS and are not shared with the manufacturer. Responsibility for the instance data follows the device lifecycle. The manufacturer creates the Asset Instance AAS at delivery, including serial number, factory calibration and

nameplate data. During engineering and installation, the integrator adds as-built information. After commissioning, the operator maintains calibration, maintenance and operational records. Responsibility for the Asset Type AAS stays with the manufacturer throughout.
In a pharmaceutical plant, this separation has regulatory consequences. When a failed device is replaced, the operator has to check and validate that the replacement meets the requirements of the instrumentation specification. If the instrumentation specification, the Asset Type AAS of the replacement and its Asset Instance AAS are available in machine-readable form, a large part of this validation can be automated. This



presupposes that all three use the same semantic identifiers and that none of them is outdated. Storing calibration results in the Asset Instance AAS also makes the calibration history accessible. Each result is stored with reference to the asset and can be compared with previous calibrations and reference data, so that changes in measurement behavior become visible over time. Metrological traceability to reference standards is still established by the calibration certificate itself. The AAS makes that certificate available across systems.

04 | LET'S TALK CROSS COMPANY

NICE BLABLA. MY FACTORY CONSISTS OF MANY VENDORS. LET US EXPLAIN.

Keeping AAS instances consistent within one company is already a governance challenge. Making that consistency survive across manufacturers, engineering companies, integrators, and operators is considerably harder. The Asset Administration Shell provides the technical framework for exchanging information, but flexibility is part of its design. It does not, by itself, ensure that two companies describe the same device, property, or engineering requirement in the same way.

Cross-company interoperability therefore depends on more than adopting the AAS format. The participating organizations need a shared understanding of the models, submodels, workflows, and data structures used at each step of the asset lifecycle. These conventions have to be tested against real engineering scenarios and developed with the parties that will ultimately exchange and consume the information. Otherwise, the result may be technically valid AAS data that still requires manual interpretation before it can be used.

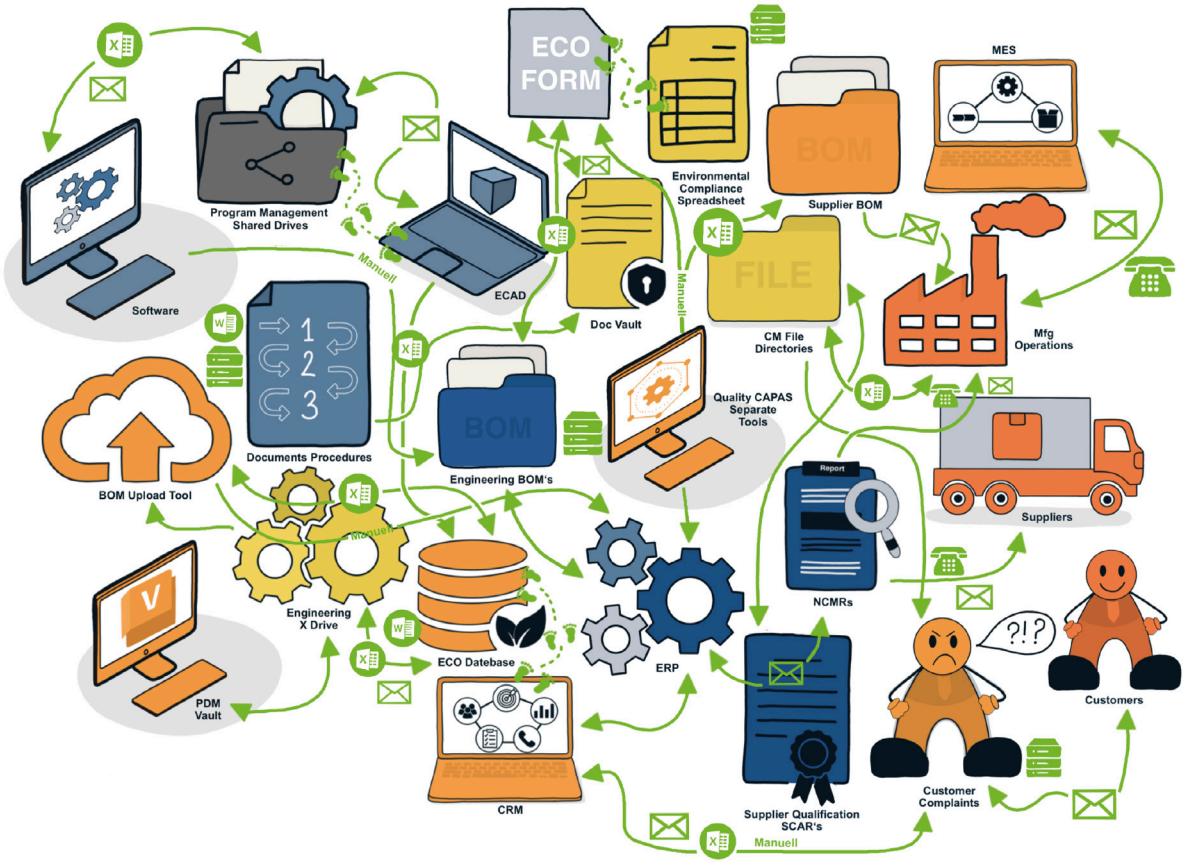
This becomes particularly important in process plant engineering. The lifecycle of an instrument typically progresses from functional requirements and P&ID development through functional design and specification to device selection, implementation, onboarding, and finally maintenance during operation. Especially from specification onward, several AAS-based use cases are already mature enough to show where standardized, machine-readable information can remove manual work.

Digital calibration certificates already show how standardized asset information can create value across company boundaries. Device specification is another strong example. Today, engineering companies and manufacturers still exchange requirements through PDFs, spreadsheets, emails, and direct coordination. Requirements must be interpreted, matched against vendor portfolios, clarified, and returned, a process that can take days or even weeks.

A cross-company AAS workflow can make this exchange structured and machine-readable from the start. Based on NAMUR NE 159, engineering requirements can be captured digitally and extended with additional specification submodels where needed, for example for PROFINET communication. Manufacturers can then compare these requirements automatically with their device portfolios and return matching device types much faster.

The key benefit does not come from the AAS alone. It depends on companies agreeing on how models, submodels, and semantics are used at their interfaces so that information can continue downstream without repeated manual interpretation. This end-to-end workflow has not yet been proven at industrial scale. Current joint research with KIT is therefore focused on identifying where interoperability already works, where gaps remain, and where standardization can deliver measurable gains in engineering speed. We increasingly see requirements for standardized digital data appearing in RFQs and procurement specifications. For example, BASF is already requesting Digital Nameplates and Digital Handover Documentation from suppliers for one of their plants. This is a good indication that the industry is gradually moving from pilots and discussions toward practical implementation.

The industrial environment behind these workflows is rarely built around a single system or data source. Engineering tools, document repositories, ERP and MES systems, supplier data, spreadsheets, emails, and manual handovers all coexist and exchange information in different ways. This diversity is exactly why cross-company standardization matters: the goal is not to replace every existing system, but to create a common structure that allows information to move between them without being repeatedly translated, reinterpreted, or entered by hand.



05 | SHOULD YOU START NOW?
AN HONEST OPINION

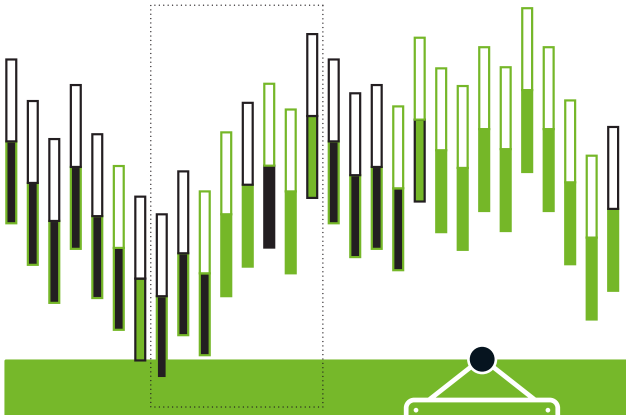
No one can say with certainty that AAS will become the dominant standard. It may evolve considerably, another standard may emerge, or the industry may move in a different direction. What can already be addressed today is the quality and structure of the underlying data. Clear semantics, consistent models, machine readable information, and well-defined interfaces will remain useful regardless of which standard eventually becomes dominant.

That preparation creates flexibility. A company with structured and accessible data can adapt much faster when customers, partners, or regulators begin to expect standardized digital information. A company that waits until those expectations are already established may first have to untangle years of inconsistent data, manual processes, and disconnected systems before it can participate.

AAS provides a credible basis for starting that preparation because it is grounded in industrial requirements and is already being evaluated in real use cases. This stage offers the chance to experiment, identify gaps, improve internal data structures, and build experience while the direction is still taking shape.



Major technological shifts often look uncertain at the beginning and obvious in hindsight. There is no need to predict the winner today. There is a strong case for making sure your data is ready when the market decides.



HELLO, WE ARE OPEN

No vendor and operator gets a consistent, cross-company standard by building it alone behind closed doors. The moment an AAS has to mean the same thing to a manufacturer, an integrator, and an operator at once, that only works if all three built it together. The OI4 & our workgroups are exactly that room. Bring your catalog data, test the prototype, help shape the standard now, before it becomes something you are simply told to comply with. As a Swiss association, OI4 is vendor-neutral and pre-competitive: we sell no product, but bring together the perspectives of more than 100 members to identify what truly works across industry. The strongest standard or technology wins on merit — and our members see what is coming before the market does.

START WITH OI4

OPEN INDUSTRY 4.0 ALLIANCE

VERSION HISTORY

VERSION	DATE	DESCRIPTION	EDITOR
1.0	November 5th, 2026	Initial release	Lucas Wolf