

Understanding the Smart Manufacturing Innovation Platform

The Smart Manufacturing (SM) Innovation Platform (SMIP) is an essential part of CESMII's strategy to enable, accelerate and scale the research and innovation that's so necessary to solve both, our energy productivity and overarching Smart Manufacturing challenges here in the US. The SMIP is a combination of strategic technologies useful for organized collection and transmission of manufacturing data, which work together to enable open access to a variety of applications, helping the industry identify constraints, predict challenges, and optimize decision-making and operational orchestration. The objective is to enable every researcher, vendor and System Integrator that engages with us to improve how they create value for manufacturing. At a high level, there are three main components to the SM Platform: the Edge, the Core Services, and the Integrated Applications – although the architecture will support many different deployment and scale options, as a general rule the demarcation between components is as follows:

- The SM Edge resides close to the machine or system that is emitting data, and is used to collect that data in-context and transmit it to the Core Services.
- The Core Services reside in a hyperscale Cloud platform and receive the transmitted data, organize and store it, and make it available to constituent applications and services.
- The Integrated Applications are consumers of the data, but also augmenters and modifiers, that can work together through Workflow orchestrations in a data-centric fashion.

All of these components use or benefit from, the Smart Manufacturing Profile – a key enabler to democratizing data access in manufacturing. The SM Profile provides a “decoder ring” for the source equipment or system, a mechanism for establishing and maintaining the intent and context of the data as it flows through the system, and an interface definition language for programmatic access when the data is at rest in an information system.

The following document will explore the current state, and future goals, of each of these components. A key enabler for each is the Smart Manufacturing Profile, an understanding of which is considered a pre-requisite for this document. For more information on the SM Profile, please see the earlier companion White Paper “Expression and Application of Smart Manufacturing Profiles.”

Smart Manufacturing Innovation Platform Edge

As with all components of the SM Innovation Platform, the Edge is both a real and available piece of application code and an ideal concept that has yet to be fully realized. The difference between the current state and the ideal end state is a continuously collapsing delta, as today's software is iteratively improved toward the completed vision. At the time of this writing, one might summarize the current state as representing perhaps 20% of the final state. But that

doesn't mean the current version isn't useful – only that its uses may be constrained, or that work-arounds might be necessary while the component evolves.

Since its earliest implementations, this Edge component has supported Historical data Access. Typically brokered through a specialized piece of software, Historical data, in the form of time series samples, represents a useful, but slightly delayed view of an operation – when the correct data has been chosen for Historization. That is, assuming a usable selection of Plant Floor data is available in a Historian, such as OPC HDA, OSI Pi or Wonderware InSQL, the SM Edge component can readily make use of such data for insight and innovation. The importance of this first source of data cannot be understated: OSI Pi, and similar Historians, are often the first piece of information infrastructure that a manufacturer will acquire. A Historian is as close to a common denominator as we're likely to find in the highly fractured and siloed ecosystem that is U.S. manufacturing. Connecting to these data sources is a critical first step.

However, there are a number of challenges presented with using a proprietary Historian as the primary data source. The aforementioned dependency on human-configured data selection often renders the Historian as a poor proxy for the operation – one that emits only an incomplete view. Depending on the Historian that's in use, or how its configured, compression used for long-term storage of samples, may further obscure that limited set. Changing the configuration, either to expand the dataset or modify the compression settings, is likely to impact brittle, or custom code that was written specifically for a given configuration – creating cultural or operational resistance to changes that could limit the value the SM Innovation Platform can provide.

Obviously, given these challenges, it would be desirable to be able to include other data sources, or provide for automatic configuration of data collection – which is one of the goals of the SM Profile. In many cases, if it's available as a data source, the Historian may provide a useful source of meta data, or curation, that would be useful – as a part of a greater whole. Taking OSI Pi, as an example, part of that product is a library of over 300 different protocol adapters, that allow Pi to communicate with raw plant data sources. Replicating this connectivity is a non-trivial effort, that only a very few vendors in the industry could support. Furthermore, making sense of those raw points of data, once connected, is not something that has ever been done without human curation. The very selection of points of data for Historization, requires a naming convention that re-labels what is sometimes a raw memory location or OEM provided tag name into something with semantic meaning. Since no standards have been adopted on these labels in the industry, both the original source name “0x00FF12” or “t_01_temp_1” and the human assigned name in the Historian “RO7000-SideA-TTEMP” are useless for establishing context. Typically, this ambiguity is solved by establishing *yet another* ontology at the information software layer. *Sometimes* that semantic layer is published or re-used, through templating, or private libraries. Occasionally, some industry standard, such as PackML, can be relied upon through-out (or at least at the Historian layer.)

The SM Edge, and the accompanying SM Profile, attempts to solve this problem by requiring either the originator of the source data, or some expert that can speak with the same level of

authority, to publish the semantic meaning in the form of the SM Profile – a format not specific to a particular implementation. In the case where a complex, multi-level system can only be articulated in the aggregate, through a Historian or other, higher level system, the SM Profile acquisition directives can be bound to the Historian’s API. In the case where a system was built by an OEM, or to some standard or spec (see: VDMA) that acquisition can be bound directly to the source equipment communication protocols.

In either case, the SM Profile separates the task of declaring the data model from the task of building the protocols specific bindings, allowing for maximum re-usability across otherwise incompatible systems. Once a Profile is applied to a system, either automatically, through some detection system yet to be developed, or manually through human selection, the pre-identified data can begin flowing. Importantly, each Profile (or Profile fragment) specifies only one selection data from the whole – with no intent for that selection to ever be exclusive. Other Profiles can simultaneously be applied at the Edge, each with a different selection of data. If Profiles duplicate a data selection, a function of the Edge will be to marshal the selection and optimize data collection, so as to avoid the problems of multiple requests from the same (potentially bandwidth constrained) piece of equipment.

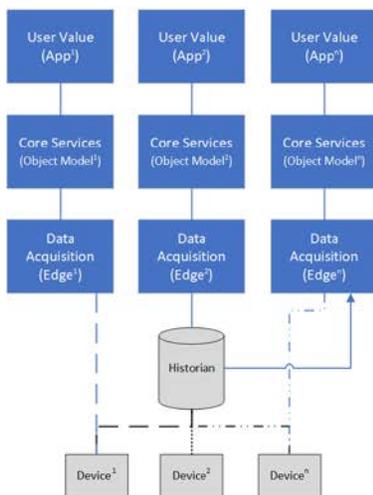


Figure 1a – Without the SM Innovation Platform

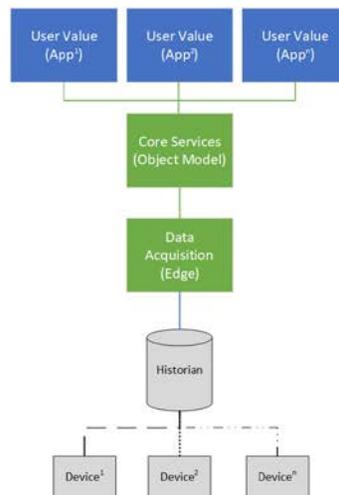


Figure 1b – SM Innovation Platform Current State

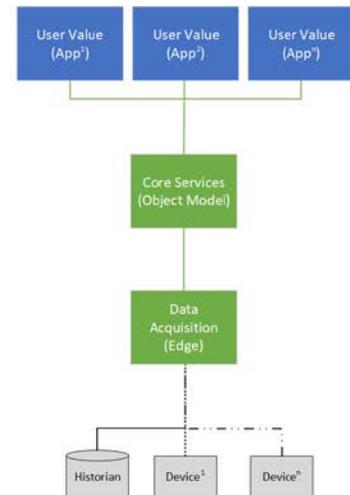


Figure 1c – SM Innovation Platform Future State

Even in its current state, the SMIP presents a significant improvement over the current landscape, but the delta between vision and reality remains largest here: Smart Manufacturing in the U.S. will require vendors to work together to solve the protocol problem, effectively normalizing the data to a common format. The SMIP will support this by embracing OPC UA, and other open standards, as a protocol for data access: regardless of the inbound protocol, the resulting objects in the SMIP will be available in a common and open format

Finally, it’s the responsibility of the SM Edge to securely and efficiently transmit the selected data to SM Core Services, typically running remotely.

Smart Manufacturing Innovation Platform Core Services

It may be simpler to identify what the SM Core Services are not:

- The SM Core Services are not an infrastructure platform for hosting arbitrary code (although providing an execution environment for specific use-cases is a capability that will be available). Instead, SM Core Services rely on Hyper-Scale Cloud providers, such as AWS or Microsoft Azure, which can host executable code that can interact with the SM Core Services, once granted appropriate authorization and permissions.
- The SM Core Services are not a Big Data platform (although the SM Core can make use of, and interact with Big Data platforms.) Big Data storage and exploration exists in a relatively mature ecosystem, typically requiring significant infrastructure to support. Instead Innovation Teams can use the SM Core Services to test specific theories, or deploy their findings at scale, at a fraction of the cost of a hosting a Big Data platform.
- The SM Core Services are not another proprietary information layer that requires another arbitrary semantic labeling scheme (although proprietary information layers can interact with the SM Core). Rather, the SM Core Services are a type-safe, semantically strict end-point, that embraces open standards to allow secure brokered access to contextual data that application developers can rely on. It provides a contract for data that separates the task of developing new insight (eg: an analytic or visualization) from the task of getting access to data in the expected format. The SM Profile dictates the semantic labels both at the Edge and in the Core Services.

A corresponding list of capabilities the SM Core Services *will* provide includes:

- A Graph database of structured data, in the form of objects and relationships, that expose, to an authorized consumer (or producer) everything collected by the Edge (or other, orthogonal data sources.)
- An Application Programming Interface (API) that allows for the discovery of data, and type information, for every object stored in the Graph database.
- An Application Programming Interface (API) that allows for the insertion or mutation of data, to support both orthogonal (non-traditional, or non-linear – that is not flowing directly up from the plant floor) data sources, and net new data or objects in the Graph.
- Secure access brokered through an authorization model that allows an administrative user to grant or revoke application and user permissions at a level of granularity that is appropriate for security and safety. In choosing this granularity, CESMII and its members will attempt to create an easy-to-understand experience, but err on the side of caution, so that the manufacturer is always in control of their data.
- A User Interface for organization and basic visualization of data as it flows into the SM Core Services. This includes the ability to structure Graph objects into a purely organizational model, such as S95.
- A database for “warm” storage of data useful for visualization, development and run time of applications for visualization, exploration or prediction of manufacturing data.

The definition of “warm” will be finalized through iteration but will stop short of “Big Data” as noted above.

- Facilities for instantiation and management of SM Platform resources, including the Edge, Core and Apps.

In short, the SM Core Services are the organization, management and brokering layer of the SM Platform – the literal core where data (or links to external data) are stored and made available. Built on Open Standards, where available, the SM Core Services in a scalable fashion are a uniquely CESMII offering – but one that others could implement, using the specifications that will be published during development.

Of the components of the SM Platform that are in early-stage development, the Core Services are the most mature. CESMII member ThinkIQ has provided:

- An implementation of the object model that uses Postgres as both the Graph database and the “warm” data storage.
- An Edge connector that supports historical data sources and transmits sample values to the Core where it is stored against type-safe object instances that are implementations of strict type definitions (the latter provide a direct analogy to the SM Profile)
- Secure programmatic data access over an encrypted channel. At the time of this writing, this is provided as a GraphQL-compliant API in “alpha” quality.
- A user interface for organization and reviewing (via Trend) streaming and historical data stored in the Core.
- Administrative tools for UI and data level restriction of user data access.
- Rudimentary tools for instantiation and management of SM Platform resources.

For these Core services to continue to evolve, it is critical that to continue to invest in efforts that test, utilize and enhance these capabilities – everything else in this document is a constituent of, and dependent on, a healthy set of core capabilities.

Smart Manufacturing Marketplace, Applications and Orchestration

The final set of capabilities of the SM Platform become possible once data is flowing from the Edge and brokered via the Core. A vibrant ecosystem of Applications that consume the data produced by the SM Platform is enabled when application developers can de-couple their development from the problem of getting access to data. While there are near infinite possibilities, for the sake of exploration, this document will divide application development into two broad categories: “Visualization and Insight” and “Analytics and Predictions” -- both of which will be marketed and available for acquisition in the SM Marketplace.

Briefly stated, the first category is about liberation of the data, empowering the U.S. manufacturing workforce to understand and react to data flowing from an operation, to make better decisions. This could come in multiple forms:

Descriptive analysis: rendering data in a more human readable fashion or delivering that data to a surface where it can be more readily used (such as connecting a WiFi-enabled scale to a dashboard or presenting a metric in a smart phone app).

Diagnostic analysis: automating primitive logical evaluations in order to report actionable guidance (for example, replacing a manual created alarm with an automatic alert driven by an OEM's own equipment performance profile.) Applications in this category aren't particularly innovative in the sense that they are solving a previously unsolvable problem, rather, they are innovative in that the SM Platform makes them more instantly accessible. Instead of each new app needing its own path to data (that is, its own set of 300 connectors to support the various protocols that have proliferated in this space), its own data storage, and its own data model, all apps that build on top of the SM Platform have the bottom 75% of their architecture already complete, and available for them to rely on. The developer needs only to focus on providing a great user experience and useful outcomes. The users of the Marketplace will determine which of those experiences and outcomes are worthwhile, as the good ones float to the top of the rankings. Note that of the three parts of the platform described here, the Marketplace is the least mature – existing primarily in concept, but as one with patterns and tools to leverage that have been developed for similar offerings in other spaces (eg: Apple AppStore, Azure Marketplace.)

The second category of Marketplace offerings has all the same challenges of the first but explores a relatively untapped direction of development. Most attempts at building predictive models in manufacturing use an export from potentially a wide variety of data sources (the classic .csv file of columns and rows of labeled features and sampled values) – no real-world implementation could rely on such a data source! Predictive and prescriptive models depend on a continuous flow of contextual (structured, labeled data) require something like the SM Innovation Platform to deliver it – and to host the result of the computation. This category of apps are the fullest participants in the SM Platform. While the data science exploration can use any data source, the final set of variables used in the selected algorithm are themselves a Profile! Articulated as such, the researcher can define criteria necessary for the prediction and expect the Innovation Platform to fulfill those requirements as part of the API. As in the first category, productization is divorced from the problem of getting the data and can focus on a user experience that is suitable for a manufacturing persona that may not be versed in data science themselves. Furthermore, the Graph nature of the Innovation Platform allows a predictive or prescriptive application to mutate the object, which was the source of the data, expanding it for all other platform constituents to make use of.

The primary API for both categories of application development is expected to become the GraphQL API, which preserves the structure of the source data in the form of strongly typed objects, the manifold relationships between those objects, the extension of those objects through new relationships or mutations, and the application of security and authorization of data access. Other APIs, including the more verbose OPC UA client/server format, or the Open Manufacturing Platform's Common Data Model (CDM) may be exposed in the future as well. Applications that use one of the supported APIs can be made available in the SM Marketplace,

and easily connected to a particular instance of the SM Platform through an instantiation process that is invoked (but not controlled by) the SM Marketplace upon the completion of whatever commercial transaction is deemed appropriate. Each application vendor will be responsible for selecting their own execution environment, and a business model that is suitable for their run-time costs and revenue expectations, but importantly, the application will not have access to SM Innovation Platform instance data without an Administrator-involved permission and authorization process.

The final step in leveraging the SM Innovation Platform stack (Edge, Core and App) will be the orchestration of multiple applications toward some industry, or user-specific goal. This important orchestration capability is mature and available today, through CESMII member Savigent’s Workflow engine. While the SM Core API’s are still in an alpha release state, Savigent can already consume and respond to data from this API and use values from that data to invoke sophisticated automated and user actions in response. As an example, Savigent could respond to a raw data value change from any member attribute of object in the Innovation Platform, and evaluate that value to determine if an operator should be directed to take corrective action, an external system should be notified to update the status of a customer order, or set-point should be changed in a control system – or all of the above, in a specific sequence. As of this writing, such orchestrations need to be manually assembled, potentially re-using the library of templated actions provided by Savigent. In the future, SM Profiles can be linked to common Workflow steps to accelerate deployment. Similarly, Workflows could be attached to apps, such that an energy prediction app, could include a Workflow template for maintenance steps useful in optimizing equipment to reduce energy consumption.

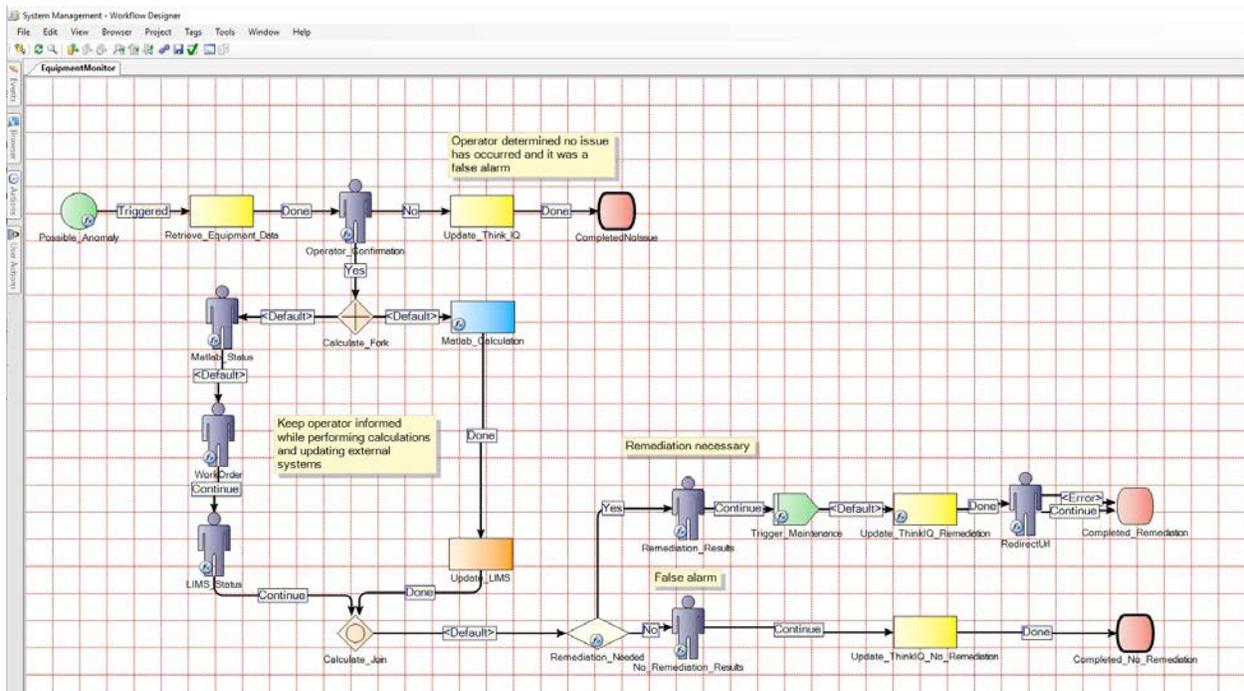


Figure 2 – App and Operator Orchestration Workflow

Deconstructing our Legacy for a Better Future

For decades, the control system has been either the execution or data source platform for “Operational Technology” information systems – although, due to the proprietary and fractured nature of that environment, usually brokered or proxied through a Historian (or some other proprietary) interface. Technologies like OPC UA or DDS that have the potential to normalize these heterogeneous environments to a reliable interface leave it up to the implementer to provide the protocol adaptation, and stop short of provide a universal namespace for type definitions – essentially moving the problem up a layer, but not actually solving it. The SM Innovation Platform by itself is not the solution – the solution is the SM Profile and a common distribution system (such as a Marketplace.) But the SM Profile needs something to implement and enforce it at each layer of the architecture; the SM Innovation Platform does just that. By implementing, demonstrating and offering a low-cost way to solve the problems of multi-protocol interaction, structured data storage, and a secure and reliable programming interface, the Innovation Platform leads the industry toward a future where innovation is focused on solving real manufacturing and energy problems – instead of on copying and labeling data from one system to another. This will require change.

New applications, that can take the SM Platform as an assumption, are free to prioritize *real* innovation. Existing applications that have their own mechanisms of extract, storing or organizing data will initially need to be retrofit. Compatibility layers, translators and middleware will temporarily be necessary; some legacy vendors may be reticent to change, because of the sunk cost they have in their own stack. However, security and scalability concerns will eventually win out – it is not tenable to have 10 different information systems with 10 different plant floor connections and 10 different user data bases; consolidation and standardization is inevitable. The opportunity for CESMII and its members is to ensure that consolidation happens around open standards and specifications – that no one vendor can monopolize manufacturing information, but that all can participate in solving the pressing problems of energy consumption and manufacturing optimization.