Invented for life

BOSCH

Data Homogenization in the Age of Industry 4.0

Digital Twin System as the basis for ease of communication between software, hardware and people in highly efficient production processes

A White Paper from Bosch Connected Industry

July 2019

Content

1. Management summary

The number of components that produce data and applications that consume it is steadily growing in the manufacturing environment. More and more frequently the stream of data does not stop at plant boundaries: IIoT software developers and machine manufacturers are just as interested in production data as machine operators on site. A common basis is essential in order to ease communication within this growing ecosystem of hardware, software and people.

The Digital Twin System addresses the needs of all those involved through comprehensive semantic data homogenization and by conveying manufacturing data along with contextual information. The principle behind the Digital Twin System is to link raw data to context. A single digital twin is a digital replica of a physical asset, such as a machine. The Digital Twin System groups the data produced by this asset into generally comprehensible information based on aspects, i.e. information groups such as machine faults or condition data.

Authorized users can easily find and use the underlying semantic models, the digital twins themselves or parts of them. This also allows external machine manufacturers or programmers to access the relevant production data. The Digital Twin System thus forms the basis for comprehensive digitization of production and logistics. This leads to new opportunities to increase efficiency in production: data can be compared and used by different systems, since the required context is always included; every recipient receives precisely the information he or she needs.

Once production is digitized using the Digital Twin System, it is possible to exchange contextualized information about the state of entire production lines. The external machine manufacturer or software developer can offer innovative services or develop new solutions without having to be on site. The Digital Twin System not only helps to optimize production processes, but also creates synergies across plant boundaries.

2. The challenge of increasing data and software heterogeneity

Cost, time and quality: companies are under increasing pressure to provide services to all corners of the magic triangle. The intelligent use of production data, however, makes it possible to turn competing targets into allies. For example, those who use machinery and production data for predictive maintenance can detect wear earlier. On the one hand, this saves time due to fewer machine breakdowns. On the other hand, fewer defects or rejects lead to an increased quality of the end product. Simultaneously, costs are reduced.

The user needs to adapt valuable data from different sources and make it applicable for his purposes.

However, the user is currently facing a problem: he needs to adapt this valuable data from different sources and make it applicable for his purposes. Two interrelated trends are increasingly gaining momentum: the rising quantity of data-producing assets, whether they be sensors, machines, gateways or entire systems, and the thirst of innovative software for precisely this production and machine data. The more complex this digital ecosystem is, the more difficult communication between people, hardware and software becomes. The enormous range of different solutions and globally located companies, programmers and machine operators need a common basis in order to meet market demands and improve production efficiency with regard to cost, time and quality.

Data homogeneity, i.e. a consistent semantic form of production, machine and status data that is comprehensible to all users, addresses the needs of all participants equally – as long as it makes data easy to find and at the same time provides the possibility of defining selective access to particular information. This need manifests itself at all levels of modern production and logistics: from the day-to-day operation of a single system to the development of a cross-sector Industry 4.0 solution.

3. The industrial ecosystem and its data requirements

The idea of comprehensive connectivity is not new. As early as the beginning of the 1990s, the Purdue model[1] developed by Theodore J. Williams led to different levels of industrial connectivity: from individual components such as valves to business operations, i.e. management of a company. Communication and the exchange of data between different parties and levels was already a huge issue. Almost 30 years later, the complexity and "scope" of this network, as well as the speed at which physical assets and software have to be implemented and new ideas have to be put into practice, have increased massively.

Industry 4.0 transforms data into a valuable currency that can be easily exchanged between different stakeholders.

Today, more than just a few applications and people in the immediate operating environment must be able to understand and interpret data. Industry 4.0 transforms data into a valuable asset or currency, that can be easily exchanged between the stakeholders. Everyone should be able to identify what is relevant to them at a glance. However, even within a single company or plant, this can be a major challenge: the number of software applications in use is tremendous, and components transmit data in different formats and in different degrees of detail. Process understanding and contextual knowledge decrease quickly as the "distance" from production increases:

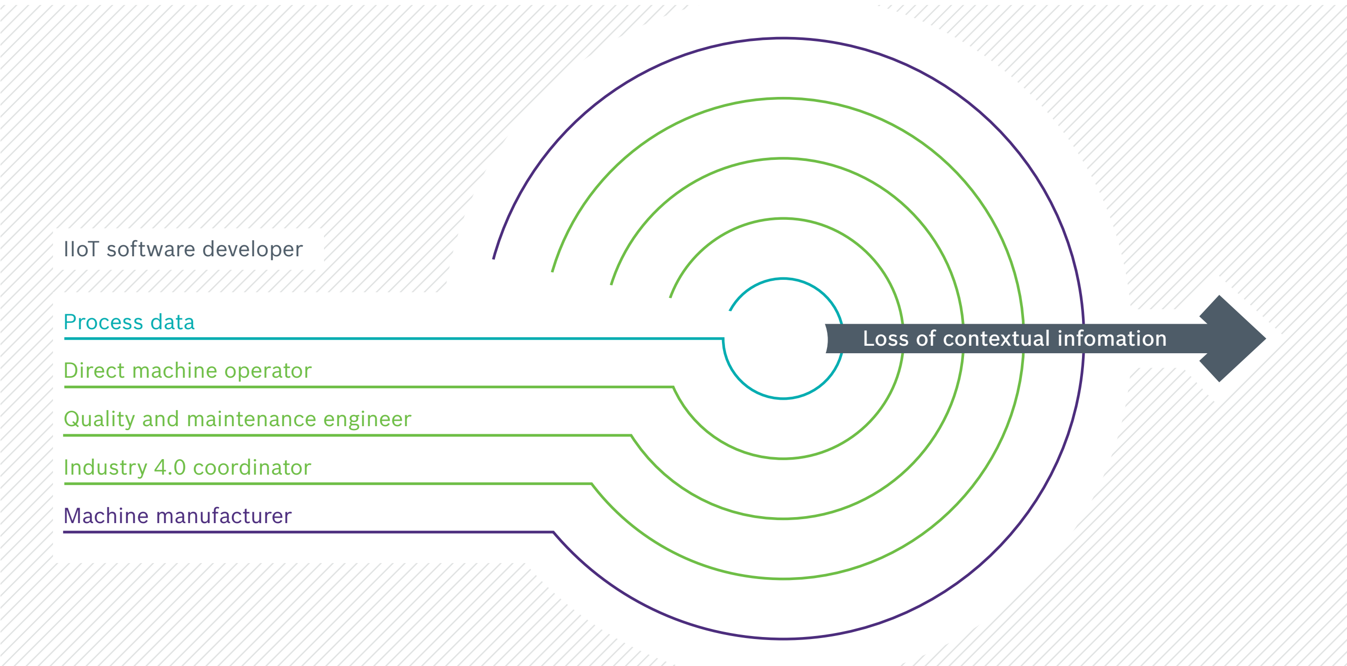

1) https://www.researchgate.net/figure/Purdue-Model-for-Control-Hierarchy18_fig2_293811556

Take the example of a senor that produces data such as a temperature value of 100 degrees: the direct process participant can put this value in the context of what he sees, i.e. a material being heated. However, based purely on the data, a machine manufacturer might not see whether this high value was related to a heating process or whether there was a malfunction. In contrast, an IIoT developer with little knowledge of how the sensor is used cannot even estimate whether the value should be viewed as particularly high or low in the context of this process; he may even think the 100 degrees is in Fahrenheit instead of Celsius.

This example shows that making data available and placing it in context both for software applications and people is a mandatory prerequisite for the successful digitization of production processes. In addition, every participant – from the machine operator to the machine manufacturer to the IIoT solution developer – has further individual requirements for a data homogenization system.

3.1 Machine-operator level

The machine operator is interested first and foremost in the maximum OEE. All processes must run smoothly and must be optimized on a continual basis. Most importantly, data must be available as consistently as possible across all assets, i.e. connectable objects such as sensors, machines, gateways and systems. Various software solutions use this data to generate warning messages and maintenance orders, uncover potential for optimization or provide recommendations for action. Condition and monitoring data, messages and the like must reach all relevant software solutions quickly and with as little difficulty as possible. This alone is reason enough to establish continuous data homogeneity and contextualization. These properties are even more critical for the outward expansion of the network.

The machine operator must maintain control over who receives particular information and when they receive it.

After all, improving efficiency within the factory's own walls is only possible up to a certain degree. Once this point has been reached, the machine operator must rely on outside support. For example, if the operator shares his production data with the machine manufacturer, the latter can provide the operator with services that improve efficiency. If in turn the machine manufacturer shares this data with developers, they will be able to adapt software solutions or updates to real production conditions. The cross-company exchange of data can also offer advantages to both sides.

With all these data transfers, one thing is crucial: the machine operator must maintain control over who receives particular information and when they receive it. A secure tool is required to manage this access.

3.2 Machine manufacturer level

The machine manufacturer would like to sell as many of his systems as possible. However, he is facing increasing international competition and must develop new business models accordingly. Services in the area of predictive maintenance, for instance, require a smooth exchange of data between the operator and the machine manufacturer who is dependent on field data.

The machine manufacturer can access the data released by the operator directly via the cloud and can work with it immediately.

If the operator provides his data in a homogeneous and contextualized way, there is no need for expensive programming or follow-up services. The machine manufacturer can access the data released by the operator directly via the cloud and can work with it immediately. Thanks to contextualization, he is able to interpret data generated by the machine or system correctly. Likewise, he can access important peripheral information, such as upstream and downstream processes, environmental data or the quality parameters of the final product.

3.3 IIoT solution developer level

The solution developer relies on data from day-to-day operations. This is the only way he can program applications that fulfill real requirements. His understanding of the process is usually not particularly deep, which means that he has to rely on the comprehensive contextualization of each data package.

The most important requirement for the IIoT developer is a model for arranging the data into specific categories with several semantic levels.

What is self-evident to a machine operator or machine manufacturer, e.g. the fact that a workpiece becomes unusable above 100 degrees Celsius, is just a number and means nothing more to the programmer without reference to a specific production process. He also needs to know that this value is a critical temperature exceeding the upper limit and is not just a temperature reading related to condition monitoring. In addition to easy data access, the most important requirement for the developer is a model for arranging the data into specific categories with several semantic levels.

4. Approach: data homogenization with the Digital Twin System

To meet the requirements of all parties, we need as system that provides comprehensive semantic data homogenization, conveys contextual information and arranges data into specific categories, such as warning or status messages, while also providing full access control to each data package.

Bosch Connected Industry, the business unit established by Bosch in 2018 for Industry 4.0, has developed Nexeed, a comprehensive portfolio of software solutions. As a leading user and provider of Industry 4.0 solutions, Bosch understands the individual challenges that developers, manufacturers and operators face with regard to data homogeneity. Measures bundled under the Nexeed Open Integration project facilitate communication between people, software and machines. The Digital Twin System is currently being developed to remove “language barriers” and to provide an easier way of finding and accessing data. This system ensures that the data of each physical object in production or logistics is presented in a way that can be understood by everyone and found quickly. Hence, it addresses the aforementioned needs of machine operators, machine manufacturers and IIoT developers.

The term “digital twin” is not precisely defined and is currently interpreted in different ways in the industry. It is primarily used to describe the digital replica of an asset, e. g. a tooling machine, with all the information relevant for a specific application, e.g. the simulation of exceptional loads or as a component in a digital plant design.

The approach presented here goes much further: the Digital Twin System from Bosch Connected Industry forms the basis for comprehensive digitization of production and logistics by gradually creating consistent data homogeneity and interoperability. Each asset has a digital representation with consistent semantics and contextual information: the underlying semantic models (aspect models), the digital twins themselves, as well as parts of them (aspects, information packages clustered by category) are independent of each other and can be used and found by the specific user.

The Digital Twin System gradually creates consistent data homogeneity and interoperability.

4.1 Creating semantics

The basic semantic principle behind the Digital Twin System from Bosch Connected Industry can be summarized as follows:

Information = raw data + context

A digital twin is an information medium or digital representation of a physical asset, such as a sensor. It consists of a series of aspects, while each individual aspect is a collection of related raw technical data with contextual information. The aspect model defines which information is collected under an aspect. This information may be status data or warning messages. Information that is frequently needed together is summarized in one aspect, providing the recipient with precisely the knowledge he needs.

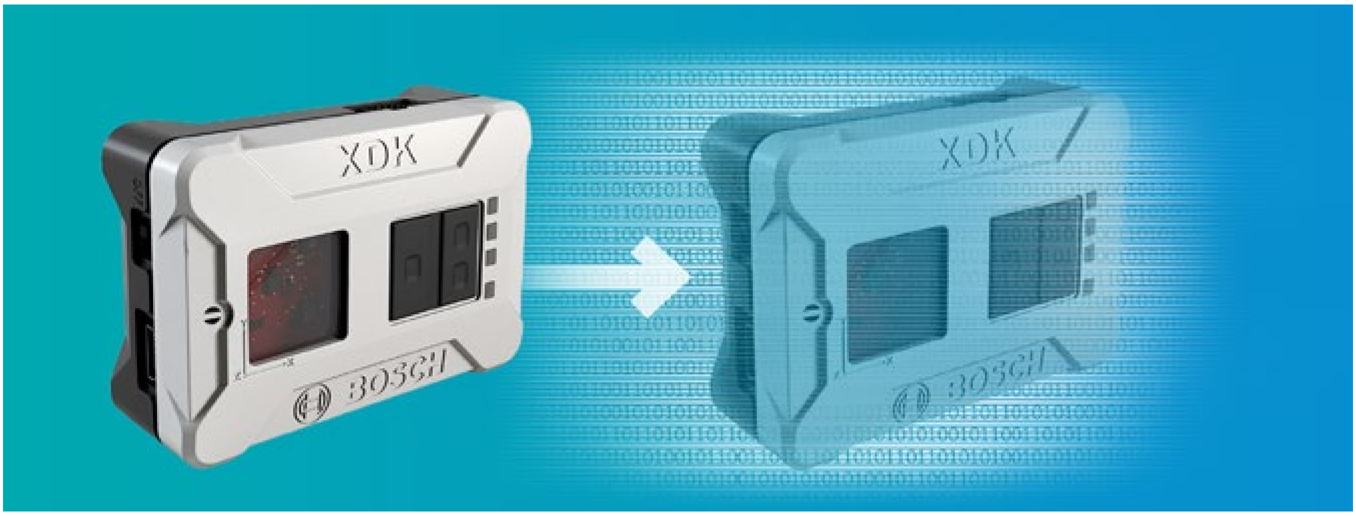

A raw data package from a machine is a good example. This data is usually easy to understand for the staff involved in the process. However, all other people and the software that will further process the data, require contextual information, particularly when a software system receives raw data from different assets and has to display the data in a uniform way, e.g. for diagrams.

Raw data
1
“Critical heat quantity”
P0557
2018-09-25T08:15
0
2018-09-25T08:14
104.5

This is where semantics comes into play. The user can now read the data better: the figure 1 stands for the type of message; “critical heat quantity” is text that people can read. The error ID is P0557, the error status is 0, etc.

Raw data	Some Semantics
1	Type of the error
“Critical heat quantity”	Human readable text
P0557	Error ID
2018-09-25T08:15	Date and Time
0	Status of the error
2018-09-25T08:14	Date and Time
104.5	Temperature

For some employees and software applications, this semantic information is all they need. For example, they know, or the processing software already recognizes that “1” represents a warning. However, what is still unclear is why there are two different time stamps. This requires even more semantics:

Raw data	Some Semantics	More Semantics	
1	Type of the error	... Error(0), Warning(1), Info(2)	
“Critical heat quantity”	Human readable text	... of the message in “DE”	
P0557	Error ID	... in the user’s list of errors	
2018-09-25T08:15	Date and Time	... when the error occurred	
0	Status of the error	... “Active(0)”, “Inactive(1)”	
2018-09-25T08:14	Date and Time	... of a the sensor value	... at the cooler outlet
104.5	Temperature	...of the cooling water	... in °C

Now it is clear why there are two dates and times. The raw data belongs to two different aspects, i.e. information groups. The first one groups data under machine faults. The second aspect, which consists of a time stamp and temperature, aggregates pure condition data. A maintenance technician or a software program that generates tickets for maintenance technicians would, in this case, only be interested in the first aspect. A maintenance engineer, in turn, would also be interested in the progression of the temperature over time, since he can extract a model for predictive maintenance from this information. Different software applications and different parties are thus interested in different aspects.

Different software applications and different parties are interested in different aspects.

Since the aspect models must only be modeled once, it is possible to quickly establish transparency and access when connecting additional assets. Thanks to the uniform semantics, it is further possible to search for information such as particular error codes or time stamps and identify correlations easily.

4.2 Providing access and discoverability

The uniform semantic model opens up a wide range of new possibilities for discovering correlations within production, but also for using synergies across company boundaries. However, to fully exploit the potential of this data homogeneity, easy access to data must be ensured. The challenge is that many different assets, applications and groups of people enter and retrieve data.

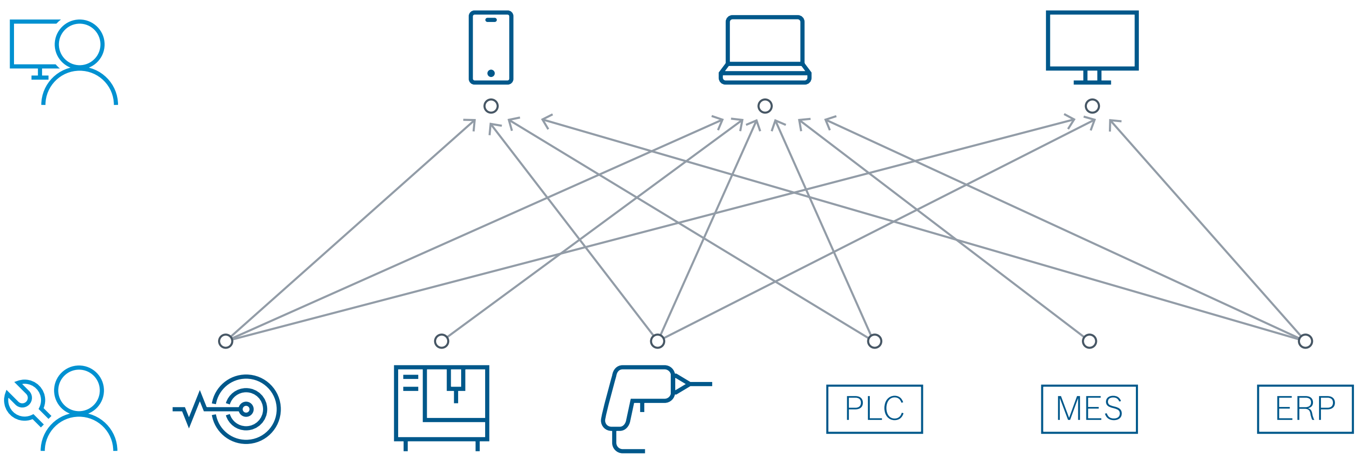

A central starting point is becoming more and more important considering the increasing number of data queries by internal and external users and software solutions. The Digital Twin System therefore features a powerful administrative center, which is a shared interface for the storage and easy discoverability of digital twins and aspect models.

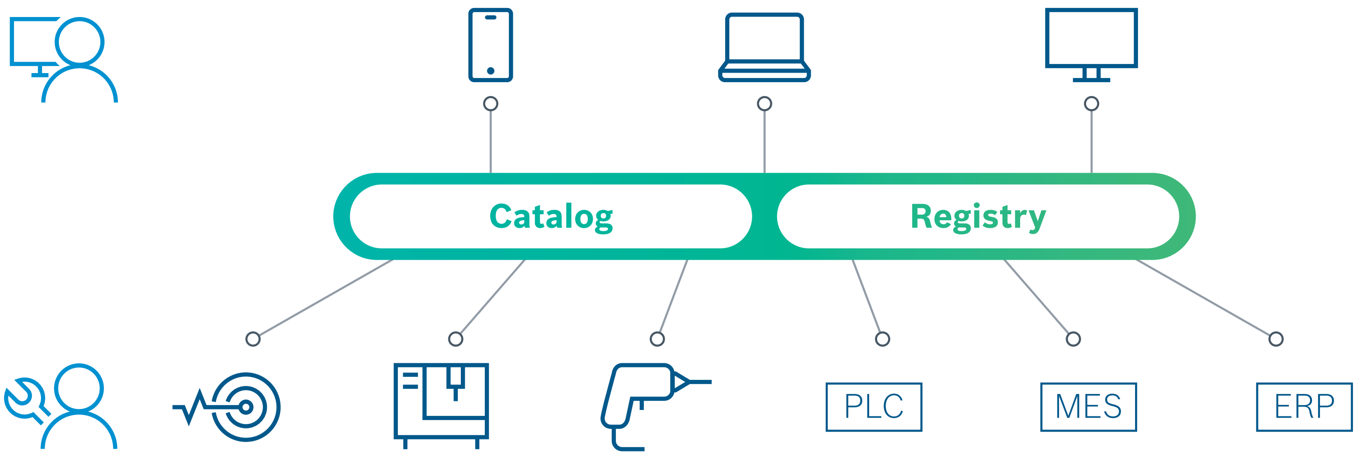

This administrative shell consists primarily of a “catalog” (a type of database for semantic aspect models) and the “registry” (so to speak a telephone book of all digital twins). The registry helps the user discover which digital twins exist and which aspect models they can service. The user also learns where and how to retrieve information about them.

4.3 Defining access groups

Users with direct or indirect access to the catalog and registry (for example through software applications) have a comprehensive insight into a wide range of different production data. This easy access presents two challenges:

- Each user only wants to view and access the information that is relevant to him and could be overwhelmed by the quantity of additional information he does not need.
- Since access for internal and external users and software applications is possible via the cloud, security-related information or trade secrets could end up in the hands of unauthorized users.

The Digital Twin System solves both problems through comprehensive access management. The data owner decides who receives access to particular information. This is handled by the Tenant Relation Management, which is an application that manages access by different groups. Tenants are defined in order to minimize the cost and reduce the number of multiple accesses set up for individuals who want to access the same assets or aspects. A tenant is usually an entire organization or organizational unit with several user groups, such as production planners and maintenance technicians. This makes it possible to provide data easily across an organization.

Access to the catalog or registry provides the user with comprehensive insights into production data.

5. Components of the Digital Twin System

Bosch has been addressing the challenge of data homogenization and utilization across machines, systems, software and enterprises for a long time. Now, individual solutions and implementation approaches are consolidated, harmonized and expanded under the same roof of the Digital Twin System for the first time. In the following, the elements of the Digital Twin System are explained in detail.

5.1 Aspect Model

Chapter 4.1 explained why and how raw data is theoretically supplemented with contextual information. To generate this type of semantic model for an aspect of a digital twin, a concise yet powerful modeling language is needed. Bosch has developed the Bosch Aspect Meta Model (BAMM), which is based on a preexisting, proven modeling framework for precisely this use. BAMM defines which information particular runtime data contains. The core elements of BAMM are shown in the following figure.

The Bosch Aspect Meta Model makes aspects comparable and reusable.

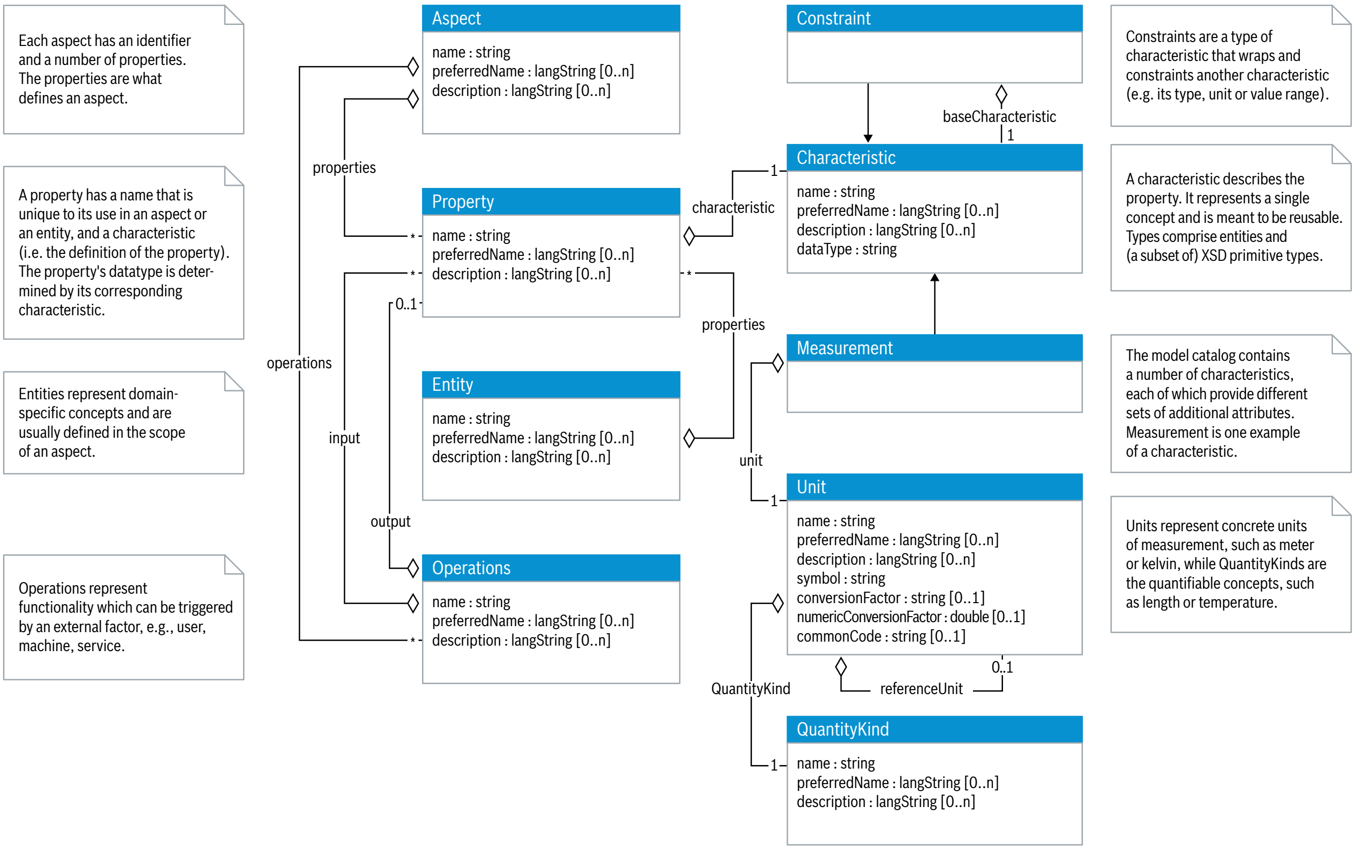

This already shows that the main focus of BAMM is not just on homogenization at the data level, but also on comparability and reusability across aspects. A simplified example of an "Error" aspect model illustrates how BAMM is used:

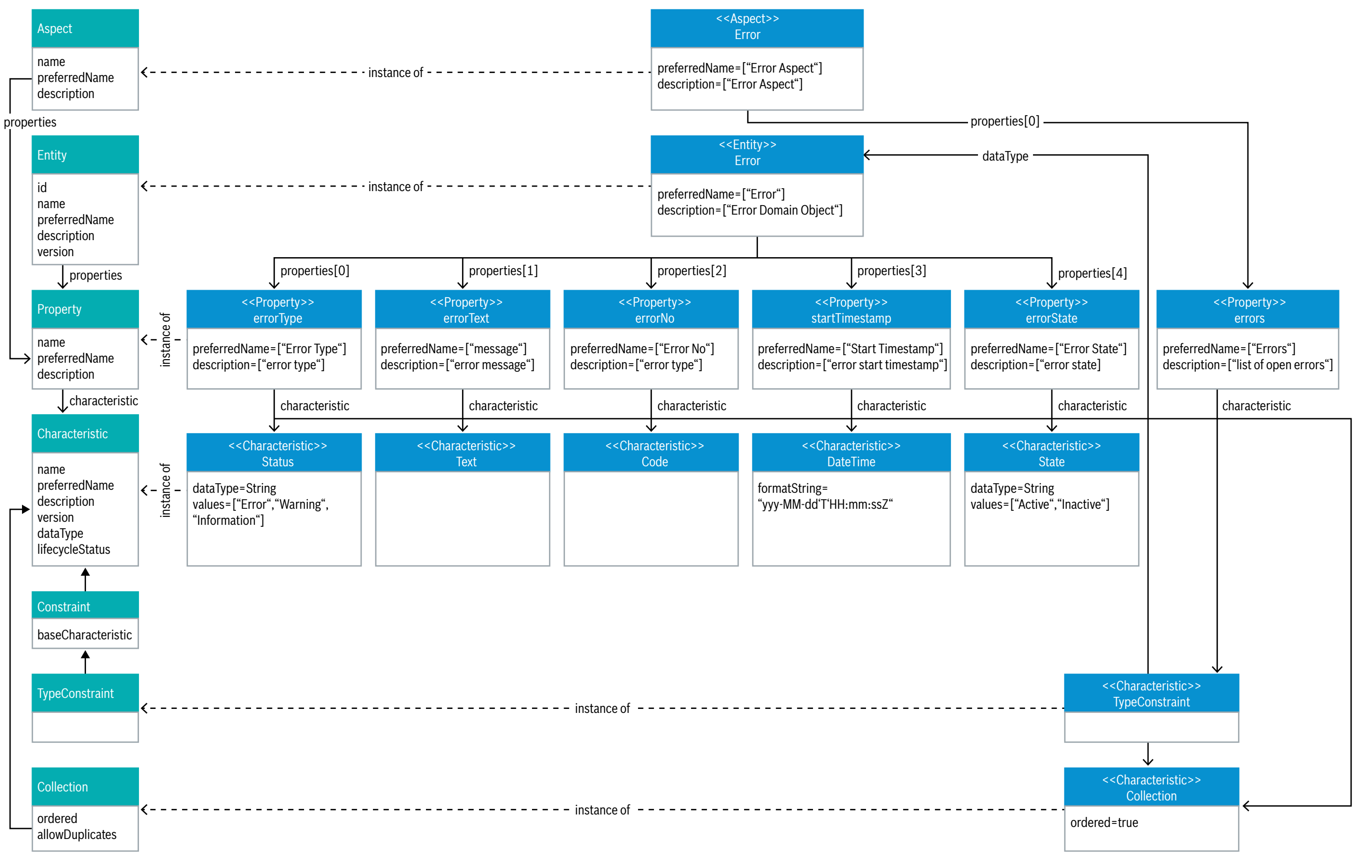

The name of this aspect model is “Error”. It describes a set of error properties, for instance of a machine.[2]

This model makes it possible to recognize the following error data properties: error type, error text, error number, time stamp and status. Each of these properties has specific characteristics, which describe all possible values: for example, “error”, “warning” and “information” are possible values of the error type, described in the status characteristic.

Thanks to BAMM, the data descriptions are homogenized across all aspect models, independent of the asset. If, for example, several identical machines feature different monitoring sensors as a result of gradual modernization, a uniquely modeled aspect model in BAMM can ensure that the data from different machines ultimately results in error messages with the same semantic logic.[3] Cross-system problems, such as those resulting from the use of a batch of defective raw material, can be identified immediately.

BAMM makes it possible to homogenize data descriptions of all aspect models.

The serialization of aspect models is facilitated by TURTLE (Terse RDF Triple Language), a text-based syntax for presenting data in a Resource Description Framework (RDF) model. This makes it possible to easily share aspect models created by domain experts in the Aspect Model Repository of the Digital Twin Catalog with the approved users. Aspect implementations provide source data in a uniform format. TURTLE is used to interpret this data, i.e. to augment it with context and turn it into information. Not all aspect models need their own aspect implementation. The Bosch Machine Twin Services described in chapter 5.6 provide a single aspect implementation for five aspect models.

2) An individual error with its properties is defined as an entity. In the simplified illustration, an aspect consists of only one entity.
3) To do this, it may be necessary to create different aspect implementations.

5.2 Tenant Relation Management

With the Tenant Relation Management application, the aspect model offers a tool for assigning access rights securely. The machine operator has full control over the data and can access all aspects of an asset as well as all aspect models. The following figure shows the operator as “tenant1”, i.e. the user in our machine example who is viewing process, product, machine, error and quality data, each summarized in one aspect.

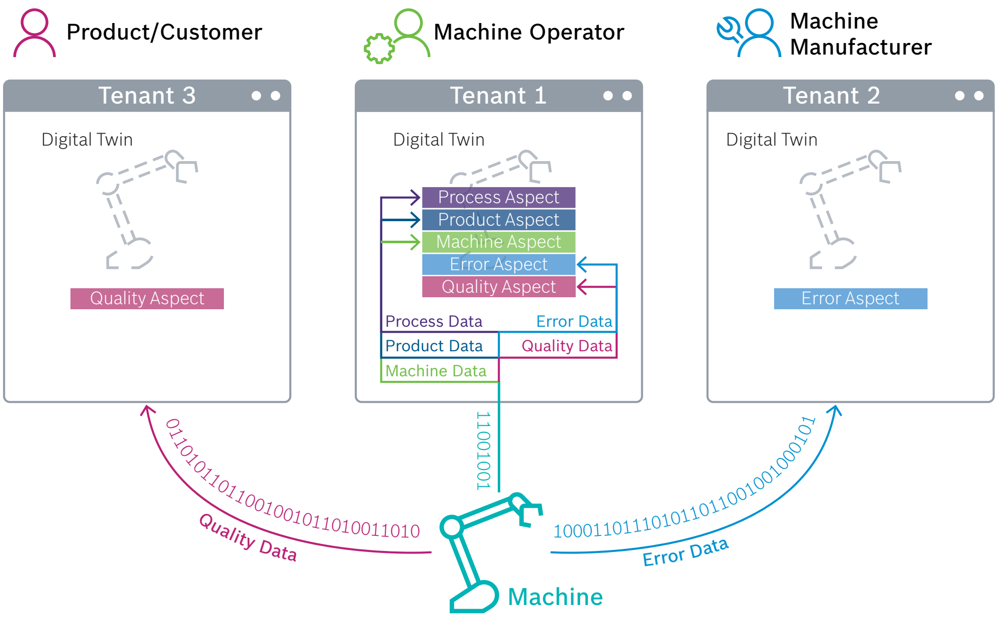

For example, if only the error messages are to be shared with a machine manufacturer, a new access group can be defined, “tenant2” in this figure. This group, consisting of any number of users, only sees the “error” aspect assigned to it. At the same time, the access group “tenant3” can be used to share the data grouped under the quality aspect with authorized customers.

Easy access and rights control leads to a more open way of handling data.

The Tenant Relation Management application thus makes it easy to control access and rights. This results in a fundamentally more open way of working with data. If the data owner has full control and is always aware of who has access to what, he will also be more willing to share certain information across enterprises. If all parties involved use this system, data can be easily transferred in both directions, and synergies can be used optimally to increase efficiency and develop solutions.

5.3 Digital Twin Cockpit

The success of every solution depends on its user-friendliness and intuitive accessibility of the available information. The Digital Twin Cockpit, which is the central management point for digital twins, is therefore designed to be as clear and uncomplicated as possible. The web-based application makes it easy to set up digital twins in combination with the aspect models, aspect implementation and choice of data sources. The Digital Twin Cockpit extracts the necessary data from the Digital Twin Registry and the Digital Twin Catalog depending on the user's access level. When new digital twins are added, this information is added to both the registry and the catalog.

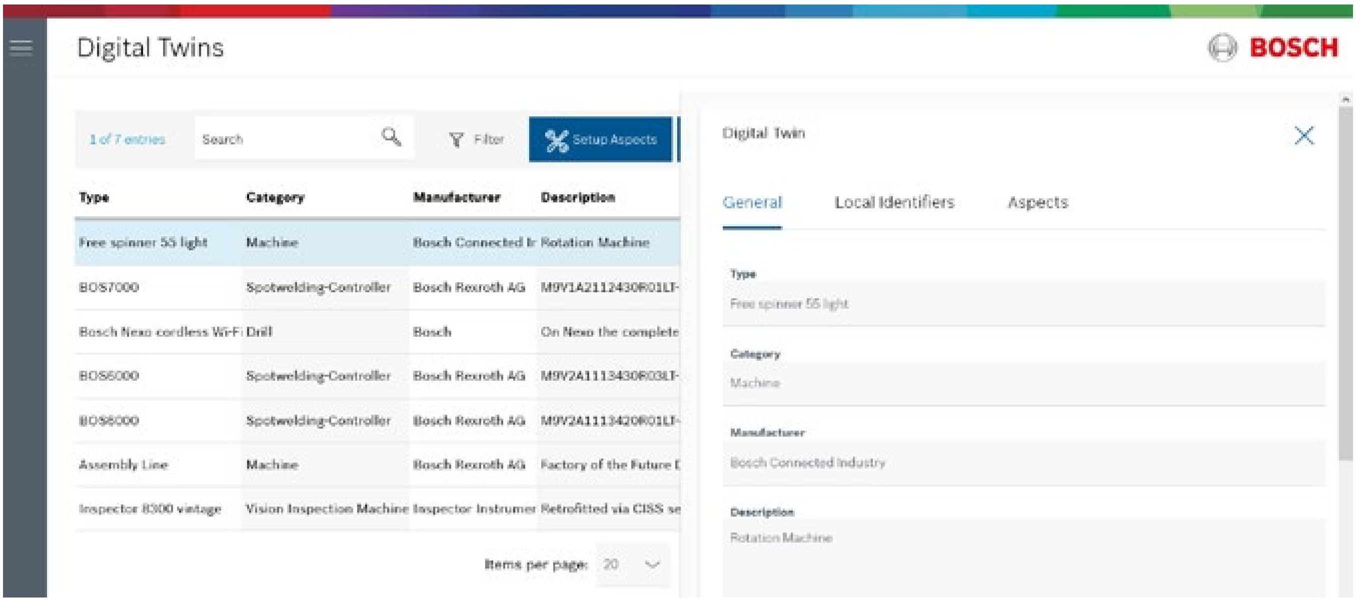

5.4 Digital Twin Registry

The Digital Twin Registry is the "telephone book" of digital twins. All digital twins and their aspects are listed here, together with basic information about the underlying asset, asset manufacturer and access options.

As soon as a new digital twin and its associated aspects have been set up in the Digital Twin Cockpit, this information is transferred automatically to the Digital Twin Registry.[4] The registry also contains information about "endpoints", such as REST APIs. This makes it possible to access the data of individual aspects directly – of course only if the user or the retrieving software solution is authorized to do so.

4) A twin and its aspects can also be added to the registry directly via an interface. The cockpit is intended for human users; the registry interface is for software applications.

DT	Type	Aspects	Endpoint
A	Machine	Errors	https://Machine Twin(A).../Errors
		ErrorHistory	...
		Condition	https://Machine Twin(A).../Condition
		MonitoringHistory	...
B	Machine	Errors	https://Machine Twin(B).../Errors
		ErrorHistory	
		Monitoring	https://Machine Twin(B).../Monitoring
		...	...
C	Control	Diagnostics	...
		...	...
D	Sensor	Acceleration	...
		Frequency	...
...	...	...	...

5.5 Digital Twin Catalog

The Digital Twin Catalog is the “semantics and context” database of the Digital Twin System. It contains the BAMM (Bosch Aspect Meta Model) as well as all aspect models, selected aspect implementations, information about solutions for specific requirements and reference information.

This is the backbone of all aspect models. After a new aspect model has been added to BAMM, it is transferred to the Aspect Model Repository of the catalog, validated and released for use. Once this is done, the aspect model can be used via the cockpit to add new digital twins.

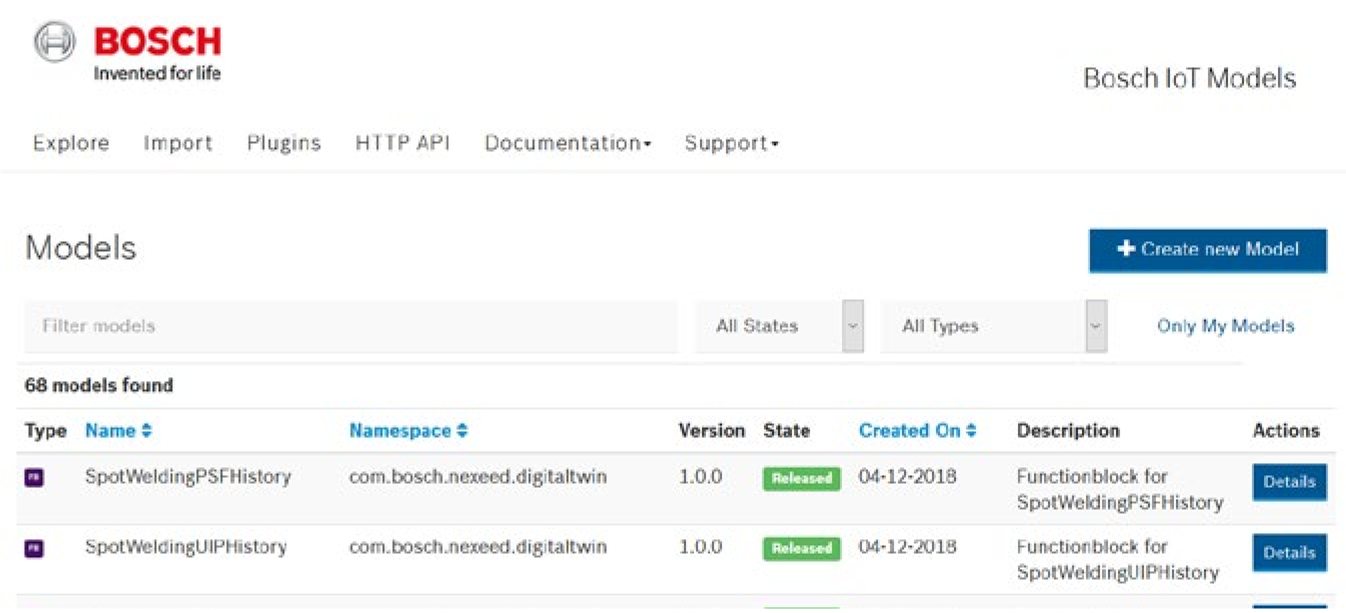

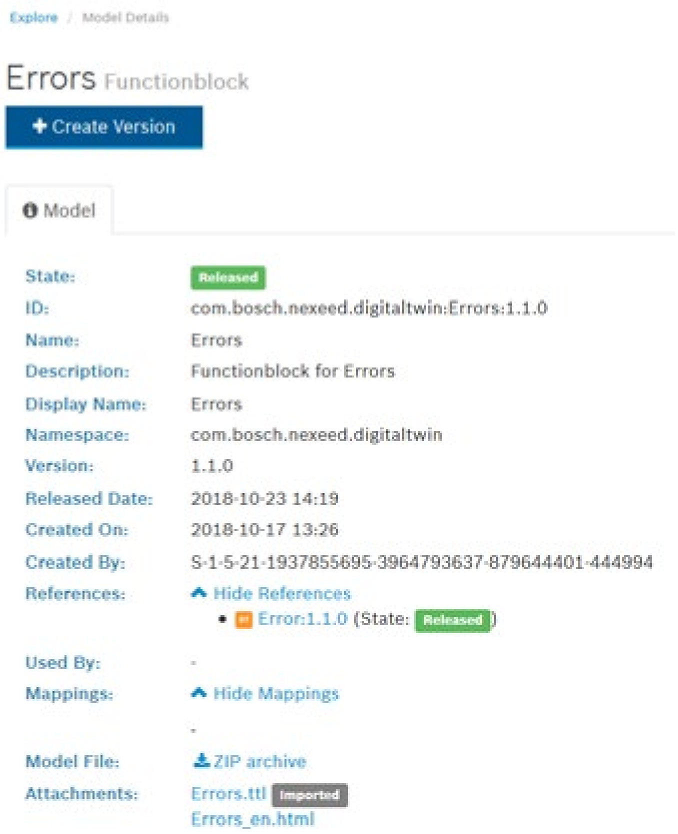

This is what the already familiar error aspect model looks like in the Digital Twin Catalog, where basic information, the TURTLE code file (.ttl) and the model diagram (.html) are displayed.

Aspect implementations and their linked aspect models are listed in the Aspect Implementation Repository.[5] Since the user is not familiar with all aspect models and thus cannot search for them specifically, the catalog also contains a solution directory. Here the user has access to information, the particular aspects he needs for a specific solution or which ones he would have to set up. The repository also manages the status of aspect models, such as “draft”, “in review” or “released”, allowing the user to see immediately which models are already in use and which are still under development.

However, the Digital Twin System user will usually not access the Digital Twin Catalog directly, since all of the information required for creating a new digital twin is available to him via the cockpit.

5.6 Machine Twin Service

The Machine Twin Service provides convenient access to a variety of aspects of a machine, as well as aspect implementations for typical core aspects. Hence, the Machine Twin Service offers a fast, easy and proven starting point for integrating machine data, particularly for situations in which machinery is retrofitted with the Bosch Rexroth IoT Gateway as the data source. The Machine Twin Service is updated regularly with new aspect models and aspect implementations.

5) The user can also develop aspect implementations and assign them to one or more twins in the registry without the catalog having knowledge of the implementation. This can be useful in the case of non-reusable aspect implementations, or if there is already an endpoint at which data can be retrieved.

6. Practical use

Data homogenization cannot take place overnight. Preparation and implementation are both long-term processes. However, even small milestones such as a machine implementation already show visible success regarding improved efficiency. Each additional milestone becomes easier to attain, since they are built on existing structures such as preexisting aspect models and aspect implementations. This way, data homogenization increasingly picks up speed. An isolated project, such as a single machine or production line, can be the starting point.

6.1 Implementation phase

When an enterprise decides to homogenize its data with the Digital Twin System, the process begins with an implementation phase. First it is necessary to determine the type of information the individual assets can provide and the type of information the different stakeholders require. At best, solutions already exist that can immediately utilize the digital twins or aspects profitably. At the same time, data sources must be available that provide the required data; retrofitting with IoT gateways may be necessary.

If there is a machine with different sensors and a higher-level software solution for production monitoring, the implementation process begins with the modeling of aspect models. Depending on the available data and planned user groups, different aspects are created based on BAMM and are added to the Digital Twin Catalog.

An aspect implementation is then developed for these models in order to provide the data from the data sources for the aspect models.

The user can now log in via the Digital Twin Cockpit and generate the first digital twin for an asset, such as a temperature sensor. A unique identification number (e.g. the sensor serial number) is used to link this physical asset permanently to its digital twin. The digital twin is initially "empty".[6] Next, the relevant, pre-modeled aspects are assigned to the digital twin from the Digital Twin Catalog. We will use the already familiar "error" and "status", which are needed for production monitoring by the software solution in this example. Now the previously developed aspect implementation is assigned. It connects the formerly purely digital aspects of the data twin with the physical assets.[7]

If necessary, the last step consists in configuring the data sources. Depending on the age and manufacturer, this can be done via device bridges or IoT gateways.[8] The twin and its aspects are now fully integrated in the Digital Twin Registry. This marks the transition to day-to-day operations.

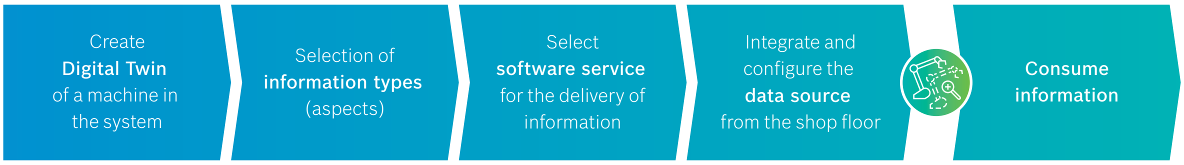

6.2 Day-to-day operations

The Digital Twin System is invisible during day-to-day operations: it handles the semantic matching and provisioning of data and controls data access in the background. In our example, the link matrix with an asset and software program is initially easily manageable and straightforward. The value created by the system really only reveals itself when additional assets and software solutions are integrated.

6) Depending on the asset, general information such as manufacturer details can also be linked to the digital twin at this stage.

7) The aspect implementations form the interface between physical assets and digital twins. The defined aspect model is not able to extract the necessary data from the data source. A specifically assigned aspect implementation is required.

8) This step can be omitted if the endpoint already provides data for retrieval. In this case the endpoint is only assigned to the twins in the registry.

The machine operator can now create a digital twin step by step for each additional temperature sensor in the factory and assign preexisting aspects to it. New aspect models need to be modeled in the case of additional assets. This creates a continuously growing network of digital twins during production and an ever-broader base of aspect models, which in turn provide a generally understandable information base for a wide variety of software solutions – from the company's own production monitoring system to cloud-based data analyses using artificial intelligence. Since the data of all production systems is provided in a semantically uniform format, cross-production diagrams can be generated, such as heat maps showing error rates or time-correlated fluctuations in quality.

Moreover, the now homogenized data can be shared easily and conveniently with all users. If a sister plant has the same system setting, the plants can share the existing aspect models and aspect implementations using the Tenant Relation Management application. Sharing data beyond company boundaries is of course also possible, and may create new data-based business models or competitive advantages. If, for instance, the manufacturing company offers to share its quality data with their customers who are manufacturers of high quality products, this can be a decisive reason for the client to choose this company as a supplier.

6.3 Partnering with external machine manufacturers

Machine manufacturers and manufacturers of industrial components are very interested in how their products perform in different production environments. This knowledge helps them not only to adapt innovations better to requirements, but also to offer better support for operators as the efficiency of their machines improves. Thanks to the Digital Twin System, connectivity has advanced significantly and has already uncovered some optimization potential.

Condition data can be shared with machine manufacturers to increase overall equipment effectiveness.

As an additional measure to increase OEE, condition data is now also shared with machine manufacturers, who in turn can use the data as the basis for customized services such as condition monitoring or predictive maintenance. Based on the transferred data, critical developments, premature wear and unusual error clusters can be discussed and solutions developed. The machine manufacturer is no longer only a hardware vendor. Instead, he becomes an active partner and service provider of individually tailored services.

6.4 Development of new IIoT solutions

In the world of industrial software there is no “one size fits all” solution. Every production environment is unique and software must often be implemented at great effort. Many issues only show up during day-to-day operations when theory and practice clash. The solution developer rarely has an opportunity to test his software under real operating conditions in advance.

Should the factory based on the Digital Twin System decide to implement new software, the expense and effort can be minimized in the early stages of development thanks to previously shared production data. Once the solution developer has the required access rights, he can use the registry to access the endpoints directly with an API query and download aspect-specific data packages in the background. With this semantically augmented information, the developer can understand the processes that take place during production and thus develop more tailored solutions. In principle it is even possible to develop a solution based only on the aspect models and without real production data, which could process all data that corresponds to these models.

Thanks to semantically enriched information, a programmer can understand the manufacturing processes and develop appropriate solutions.

By knowing the underlying semantic structure, these solutions are also easy to implement; they can process digital twins or aspects directly and profitably, or even support modeling or the creation of new systems. The machine operator benefits from improved, innovative solutions, and can profitably share his production data as a knowledge resource.

7. Outlook

This paper has shown how the Digital Twin System can establish a homogeneous database for connected production. The more assets a digital twin has, the more comprehensive the digital representation of the complete production environment becomes.

The ease of use of this homogeneous data pool, as well as the newly acquired ability to compare data of the most diverse assets shared even beyond company boundaries thanks to clear semantics, provide the starting point for innovative software solutions and new ways to collaborate. The possibilities are almost endless:

Thanks to regulated access to data from several locations, a maintenance service provider can generate dynamic overview maps for his technicians. This helps the technicians choose their routes most efficiently and promptly reach machines showing signs of wear.

Machine operators can freely select a variety of services via the cloud, for instance a one-time data analysis of a specific production line that is inexplicably producing rejects. The service provider receives scheduled access to particular aspects that he can process immediately thanks to the semantic contextual information.

When setting up a new production line or a brand-new site, hardware and software combinations that are already running efficiently along with all of their existing aspect models and aspect implementations can simply be transferred to the new solution. Aspect models and aspect implementations for particular assets can be shared across companies, and new business areas can be established with modeling as a service. And these are just some examples of the benefits the future may hold when implementing a Digital Twin System.

Visit our website to learn more.

Bosch Semantic Stack

A portfolio for a product-centric digital transformation to leverage your data across the entire product life cycle.

www.bosch-semantic-stack.com

Robert Bosch Manufacturing Solutions GmbH
Bosch Connected Industry

PO Box 30 02 20
70442 Stuttgart
Germany
+49 (711) 811 – 0

Industry4.0.bci@de.bosch.com
www.bosch-connected-industry.com