EXOSITE

USA Headquarters
275 Market St, Suite 535
Minneapolis, MN 55405
+1.612.353.2161

Taiwan Office
WenXin Road, Section 4
#955, 15F-5
Taichung, 406 Taiwan
+886.4.2247.1623

The Augmented Digital Twin

Combining physical and virtual data to unlock the value of IoT

WHITE PAPER

Released: January 2017
Authors: Ryan Kucera, Mike Aanenson, and Mark Benson
Revision: A

Table of Contents

> "The concept of a digital twin moves manufacturers and operators closer to the ultimate goal of selling outcomes (results) instead of products (machines)."

1. Introduction

The movement towards connecting every device, machine, and sensor to the internet (a.k.a., the Internet of Things (IoT)), takes advantage of two key technologies:

- **Connectivity,** which has become increasingly easier with recent advancements in Wi-Fi, cellular, and local wireless technologies.
- **Web services,** like computing, storage, and networking, which have become a commodity resource and allow easy end-user access to data from embedded devices.

Having a network connection to a device is not a new concept, but the advancements made in these technologies have brought about new challenges for companies and original equipment manufacturers (OEMs) who have little experience in data. It is relatively easy to connect a device—but what does that actually mean?

It is compelling to regard IoT as a world of silicon and signals, sensors and circuits; as an engineering triumph that transmits once-obscure data to the modern databases and interfaces of our connected world. Indeed, without these technologies, IoT would not be the multi-billion device industry it is today.[1] However, as the end users of IoT shift from early-adopters to mass market, and as hardware and connectivity become commodity assets, the value of IoT must also shift.

This shift in the accessibility of technology and data has given birth to the concept of a digital twin. A digital twin is a virtual representation of a physical device that allows us to gain greater insight by combining both measured physical parameters (e.g., vibration and temperature) and other digital information about the asset (e.g., manufacturing born-on date and maintenance history). This combination of physical and virtual information enables a range of new insights about an asset, including those regarding performance and health, and predictive insights about what may happen in the future. The concept of a digital twin moves manufacturers and operators closer to the ultimate goal of selling outcomes (results) instead of products (machines).

In this whitepaper, Exosite will demonstrate the concept of a digital twin and how it can transform an IoT solution from connected data into a valuable view of a meaningful asset. This is enabled through leading-edge developments in IoT cloud platforms, time-tested software development practices, and the integration of data and information ecosystems that derive new asset knowledge. This whitepaper will also discuss best practices to implement a digital twin model for an IoT solution. Finally, a digital twin maturity model is provided that offers a roadmap about how to grow an optimized digital twin over time. Whether you are working in a legacy environment or leveraging Exosite's next-generation Murano platform, these concepts and practices will add value and meaning to connected data.

2. The IoT Asset

Implied in the term "digital twin" is the existence of a sibling; an object, asset, or thing that is being transmitted and replicated into a digital representation. In order to begin an exploration of the digital twin, it is important to first understand this sibling.

To begin, imagine condition monitoring sensors installed on a factory air compressor system; with the exception of signal strength or battery status, the state of the sensor or measurement device is immaterial. The sensor exists primarily to describe the condition of a specific subpart of the asset, such as the temperature of a compressor line. The "thing" we care about is not the wireless-capable thermocouple or the temperature of the compressor line in particular, but rather the holistic integrity and health of the air-compressor system itself. This "thing" is the asset, the sibling, which will become the target of a digital twin. Throughout this document, the term "asset" will be used to reference the physical thing being modeled.

[1] http://www.gartner.com/newsroom/id/3165317

2.1 Observing a Non-Connected Asset

Direct observation or measurement of an asset is often the most familiar, most accessible means of understanding its condition and state. Want to know the condition of a compressor? Chances are, walking past it could provide some sense of its condition through direct observation of its key properties, especially if there is unusual noise, vibration, or heat. Instruments and gages can facilitate this direct observation by quantifying what can often be sensed through observation. Note, however, that concern is not typically given to the state of the measurement device or sensor, but rather for the asset it is measuring. This is an important distinction to make when beginning to think about connected solutions.

2.2 Local Monitoring (SCADA, M2M)

Early condition-monitoring solutions would typically map signals from sensors to channels or memory addresses, with physical displays or user interfaces that would directly represent these values.[2] The need for direct observation of the asset was removed, but it was up to the observer to infer or correlate these measured values with asset state, condition, and maintenance records.

2.3 First-Generation IoT

Many examples of first-generation IoT solutions overlaid this local monitoring design pattern on top of Internet communication protocols and software user interfaces. Database schemas and user interfaces reflected the sensor and hardware values and layout, and similarly relied on user interpretation to correlate data and events, real-world conditions, and maintenance records, which were often stored in a separate, disconnected system.

It may seem logical to scale a first-generation IoT solution by adding sensor types and increasing the range of data captured; however, simply filling databases with streaming data from sensors leads to systems only data scientists are be able to use, as most people have trouble identifying with databases. However, if this information can be presented in a meaningful model, it is possible to enable users to forget about the data and instead react to meaningful knowledge.

This is the purpose of the digital twin—to transform sensor and device data into tangible asset knowledge for business users, application users, automated processes, and more.

3. Asset Modeling with Digital Twins

To begin modeling an asset in an IoT solution, it may be tempting to pick up a development board or measurement device, inspect its array of sensors and outputs, and reflect these attributes into the properties of a digital twin. However, instead of centering the model of a digital twin on the development board, measurement device, or sensor, think instead of the entire asset being measured (for instance, an air compressor system). This is comparable to a blood pressure cuff on a patient in a doctor's office; what is being measured is not the sensor or measurement device itself, but a particular dimension of the patient's overall health and well-being.

In thinking about digital twins in this way, it becomes natural to enhance and augment the model with further meta-data, nearby environmental conditions, maintenance data from similar equipment, service history, account information for the company that owns it, manufacturing born-on data and related configurations, and data from other web services that together can create a rich and comprehensive representation of the physical device. This section provides best practices to consider when developing a digital twin that incorporates additional sources of data.

[2] https://en.wikipedia.org/wiki/SCADA#Third_generation:_.22networked.22

3.1 Model the Asset, Not the Sensors

As introduced above, the digital twin provides optimal value when it can describe the asset, not the hardware. Remember, hardware and sensors are no more than instruments and gauges that facilitate observation. The table below provides questions that can help decouple an asset from its connected hardware. Leading IoT platforms enable asset modeling with a collection of complementary capabilities. First, streamlined product definition and management features enable reflection of the connected device's sensors and digital data, minimizing the need for custom code and datastores to collect and organize data. Second, custom data and event routing capabilities enable the type of flexibility needed to utilize this data in a model that is meaningful to users and their monitored assets. Third, integration with external data sources can be used to augment and enrich the digital twin with more information than exists solely on the asset itself.

Data Type	Asset Questions	Example
Temperature	What is the temperature sensor attached to? What events cause temperature to change? What is affected by changes in temperature?	A temperature sensor connected to a livestock water tank measures the temperature of the water contents, and is affected both by environment and by heating elements. Water that approaches freezing can affect the well-being of livestock animals. In this case, your asset can be either the water supply system or the livestock herd itself.
Pressure	What material (e.g., liquid, gas) does the pressure sensor measure? Is the pressure sensor in a closed environment?	A pressure sensor on a tractor tire measures gauge air pressure. The gas type is known, and along with configured data about tire size and correlated data from temperature sensors, derived data about the load weight and duty cycles of the asset (the tractor) can be known.
Accelerometer	Is this sensor activated by a human, a machine, or an outside force? Can changes in motion be correlated with known events or other sensed data? Who needs to be notified when motion exceeds known thresholds?	An accelerometer attached to a smart mobile asset such as a lawn mower can detect changes in mowing patterns, rotary alignment, and other conditions that may indicate a need for maintenance.
Voltage, Current, Power	What type of power supply is being monitored? Is a user concerned about the cost of power? Does the power supply (e.g., a battery) need replacement or maintenance?	Voltage and current sensors measuring a bank of batteries can indicate the health of the batteries or the charging system, and can detect system performance or safety issues that should be addressed
Lack of Data (offline state)	What does the absence of data from a connected device indicate? Can gaps in data be interpolated?	Without identifying the underlying asset, the meaning of an offline state can be hard to quantify. Using the temperature sensor example above, an offline state may indicate a power failure, in which case an alert to a maintenance worker can prevent an unexpected freeze and breakage of the water supply system.

3.2 Separate Condition from Control

The digital twin reflects the current state of an asset; however, many IoT use cases also require command and control interaction with a connected device. Command and control requests should be separated from an asset model's state. This is often done with separate values for "observed" and "target" values for a data resource. Leading IoT platforms incorporate this separation inherently, allowing a digital twin to define which data resources can accept commands and maintaining the correct state of the asset as commands are received, executed, and confirmed.

When presenting a digital twin asset to users, handling the potential difference in the physical and virtual state is important. Showing two different elements during a command sequence or control session—where one contains what the user wants and the other contains what the device currently has—keeps users informed and empowered to react appropriately.

3.3 Leverage Data to Provide Meaningful Views

IoT brings challenges in representing something that may not be as simple as one physical connected thing to one digitized thing. That is to say that the data is important as grouped topics to different people. Using the same example of a factory and its sensors, most users do not care about the single sensor, although some might. During installation or in a situation where signal issues occur or a battery fails, an installer may want to know about that specific electronic device. However, a different set of users will only care about the sum of sensors in reference to the complete machine being monitored. Another set of users may only need to see data about the factory as a whole, not individual machines or sensors. It is important to be able to provide the concept of a digital twin for different levels of groupings of these physical electronics.

Tactics for enabling a meaningful view of an asset model can range from simple UI layout and graphics, to more complex approaches such as derivative data and forecasting. To fulfill this requirement, consider the balance between development complexity and creative freedom. Dashboard tools are an efficient way to visualize data, but can be constraining when correlating or deriving data from multiple sensors. Visualization libraries can accelerate custom charts and graphics, but require intermediate development skills and the right application hosting environment. Leading IoT platforms offer these options and more within their capabilities, allowing a solution to mature and target specific roles and use cases. The table below provides an overview of the platform capabilities, tools, and skill levels that may be needed to develop meaningful views for the example use cases below.

Use Case	Platform Capability	Tools	Skill Level
Proof of connectivity	Dashboard environment	Freeboard, Domo	Web UI configuration of data source types and options
Trend analysis	API data access, UI application hosting	Highcharts, nvd3	JSON, JavaScript, HTML
Enterprise data analytics	API data access, time-series datastore	Excel, Databricks	Lua scripting
Consumer mobile app	API data access, user authentication, role management	Ionic, Xamarin, Native iOS/Android	JavaScript, Java, Swift

3.4 Oversample First, Optimize Second

As an IoT product develops from concept to completion, a digital twin should also evolve. At first, take a liberal approach to measuring asset and environmental health. After data (and as a result, meaningful insights) are collected and identified, optimize data streams, bandwidth, and sampling rates to remove unnecessary cost and complexity. Use an IoT platform that offers the ability to collect and normalize disparate data sets, integrate enterprise web services, and exposes event listeners to capture, store, process, and display the right data at the right time. More information is provided in the following sections about how to design and evolve an appropriately scaled augmented digital twin.

4. Digital Twin Maturity Model

The following section describes a maturity model for digital twins, beginning first with a way to classify digital twins based on the level of information known about an asset and its environment (maturity) and how the digital twin reflects and anticipates time-based events from the real world (asset reporting frequency). The maturity model shows how a digital twin can improve over time as more data (and derivative knowledge) is accumulated.

The remainder of this section describes the tiers or types of twins that can be utilized, and the qualities and constraints that delimit each type.

4.1 Digital Twin Maturity

The scope of data that is recorded and retained within a digital twin determines what can be known about an asset's state and condition.

- **Partial**
  - The minimal digital twin typically contains a small number of data sources, such as temperature, pressure, and device state.
  - Partial digital twins can be useful to monitor a key metric or state from a low-power or resource-constrained asset, such as a connected light bulb that simply reports its current power state.
  - This level is also seen in proof-of-connectivity development, as it enables quick development of device-to-platform functionality.
  - A partial digital twin contains enough data sources to create derivative data. For example, if pressure is down but temperature is up, and linear regression identifies a correlation, a corresponding inference about the health of the asset can be made.
- **Clone**
  - The clone form of a digital twin contains all meaningful and measurable data sources from an asset.
  - This level is applicable when a connected asset is not power- or data-constrained.
  - This level is useful in prototyping and data characterization phases of IoT development.
- **Augmented**
  - The augmented digital twin enhances the data from the connected asset with derivative data, correlated data from federated sources, and/or intelligence data from analytics and algorithms.

4.2 Asset Reporting Frequency

The frequency of state and condition changes that are reflected into a digital twin determines what decisions can be made using its data.

- **Statistical sampling**
  - For statistical sampling, data is updated at a fixed sample rate that ensures a statistically significant cohort of data within a desired timeframe. There is no reaction to threshold or state-change observations, and

efficiency is prioritized over timeliness (typically 1 minute to 1 day frequency).

- **Condition monitoring**
    - For condition monitoring, data is updated either at a fixed rate that is frequent enough to enable meaningful alerting to condition changes, or data is updated when edge devices detect a state/threshold change (typically 1 second to 1 hour frequency).
- **Real-time control**
    - During a real-time control session, data is updated as frequently as possible within reasonable technological constraints, and low latency is prioritized over efficiency/cost. A persistent socket connection (e.g., websocket) is utilized and enables sub-second latency (100-300 ms).
    - An example is a connected light bulb that requires the ability to respond to commands and report current power state in real time.

The figure below shows the level of analytics maturity that can be achieved with a digital twin based on the data scope and reporting frequency used.

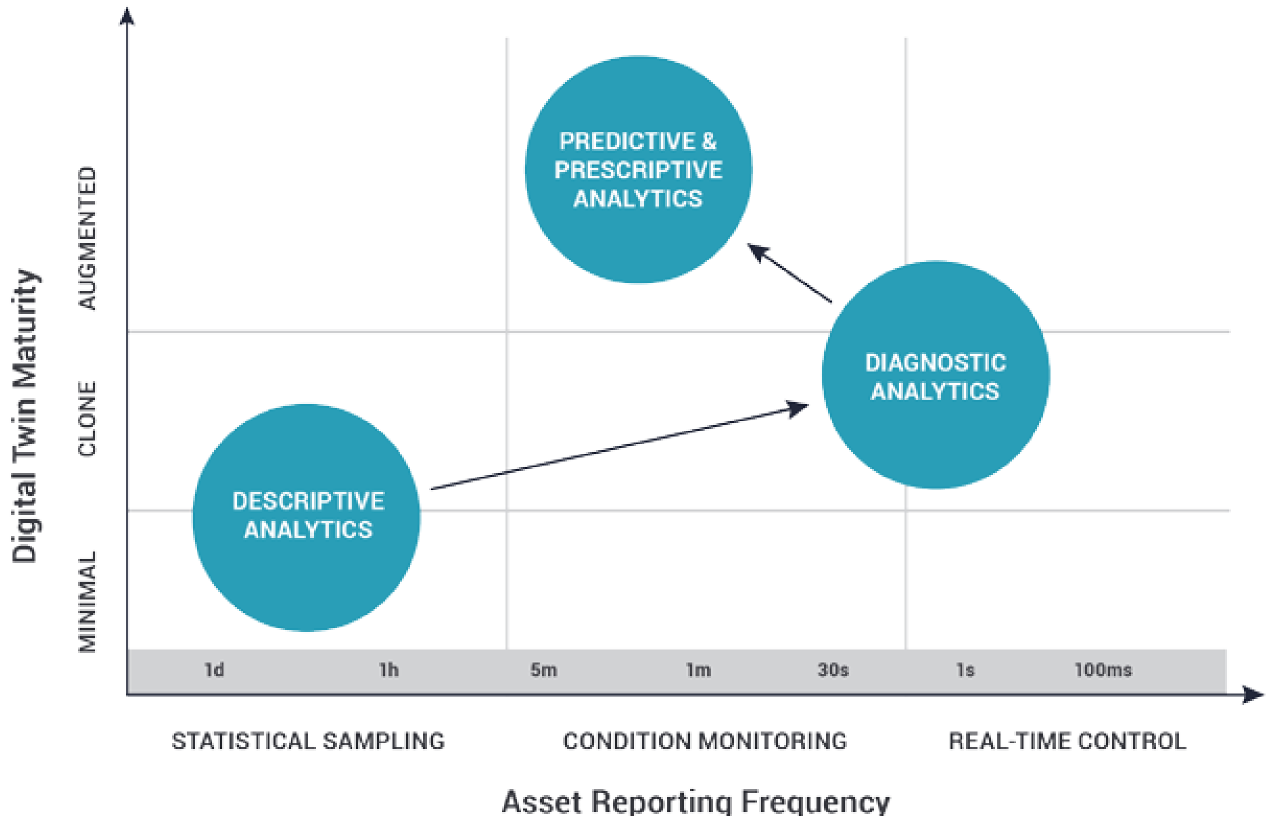

Best-in-class digital twins start with a basic set of sensors and relatively infrequent sampling frequencies to prove the technical feasibility of connecting a physical asset to its digital twin. At this stage, basic descriptive analytics (what the current state of the asset is and what events of interest may have occurred) are possible.

With the technical feasibility of the digital twin proven, additional sensors are often added to the asset and the reporting frequency is increased to provide better granularity and insight into specific performance aspects of the asset and/or more contextual data leading up to error conditions. At this stage, diagnostic analytics (why did certain events of interest occur?) are possible.

With a baseline of asset health, performance data, and the inclusion of additional sources of data such as asset maintenance history (from an enterprise resource planning system), account data (from a customer relationship management system), parts data (from an asset management system), and other adjacent data sources (such as environmental or weather data surrounding the asset), the maturity of the digital twin can be improved to an augmented status where it may be possible to reduce the number of sensors (measurement points) and reporting frequency.

It is this augmented digital twin that becomes, in many ways, more valuable than the physical asset itself and enables a range of predictive analytics (what will likely happen in the future) and prescriptive analytics (what should be done to prevent mechanical failure or reduce downtime) that unlock the true value of IoT.

5. Summary

IoT unlocks the resources of cloud computing and the design patterns of application development. Leveraging these resources correctly evolves connected hardware from gauges and dashboards into augmented, virtual representations of physical assets that enable new, valuable insights about these assets that were not previously possible. However, with this comes the risk of added complexity. Modeling and implementing the correct digital twin can multiply the value and utility of this model, while minimizing the risks. This white paper has outlined the best practices to help develop a useful digital twin, including a straightforward maturity model to build a roadmap for an optimized augmented digital twin that capitalizes on the true value of data.

MINNEAPOLIS
AUSTIN
TAICHUNG
TAIPEI
BUDAPEST

EXOSITE.COM