VANTIQ Digital Twin Architecture

Paul Butterworth – December 2017

Introduction

The Digital Twin is a key enabling concept in business digitalization efforts. A Digital Twin is a digital representation of an asset. The asset may be a physical asset such as a sensor or vehicle or it may be a logical asset such as a business process or service. The Digital Twin can be used as a proxy for the asset making it convenient for integration into digital control systems and business analysis systems.

This document describes the VANTIQ Digital Twin Architecture, a pattern for rapidly building high fidelity Digital Twins on the VANTIQ platform.

Digital Twins

Definition

A Digital Twin is a model of an asset that exists in the digital domain. The asset may be a physical asset such as a sensor or vehicle or it may be a logical asset such as a business process or service. Driven by instrumentation of the asset the Digital Twin behaves as a high fidelity, digital simulation of the behavior of the asset. This capability is used to produce a digital profile of both the historical and current behavior of the asset with the resulting profile used to:

- Detect asset behavior that is less than optimal
- Predict future behavior of the asset

Digital Twins can represent many different kinds of assets at varying degrees of abstraction and sophistication. For example, a single temperature sensor may have a digital twin that reflects the current temperature reading of the sensor while a complete vehicle may have a digital twin that reflects the current behavior of the vehicle in terms of speed, location, heading, power consumption and engine temperature. The engine temperature might be supplied to the vehicle Digital Twin by the temperature sensor Digital Twin, illustrating the construction of complex Digital Twins from the Digital Twins representing their constituent parts or the environment in which they operate.

Digital Twins have a wide variety of uses:

- Observe the current behavior and status of the asset by applying current sensor readings to the Digital Twin and observing its behavior.
- Observe the historical behavior and status of the asset by applying historical sensor inputs to the Digital Twin and observing its behavior.
- With appropriate visualizations, display a virtual version of the asset annotated with its current status and other relevant information.

- Analyze the behavior of the Digital Twin looking for opportunities to optimize the behavior of the asset. This may involve real-time analysis of the current behavior of the Digital Twin or offline analysis of its historical behavior.
- Simulate the behavior of the asset on synthetic stimuli. Simulation might be used to find the optimal operating configuration for the asset, study its failure modes or evaluate the assets behavior under extreme operating conditions.
- Suggest possible optimizations that might be applied to the asset by first applying them to the Digital Twin.
- Predict future states of asset by running simulations on its Digital Twin.
- Optimize the assets behavior by applying real-time analysis to the Digital Twin and then modifying the behavior of the asset through its actuators.

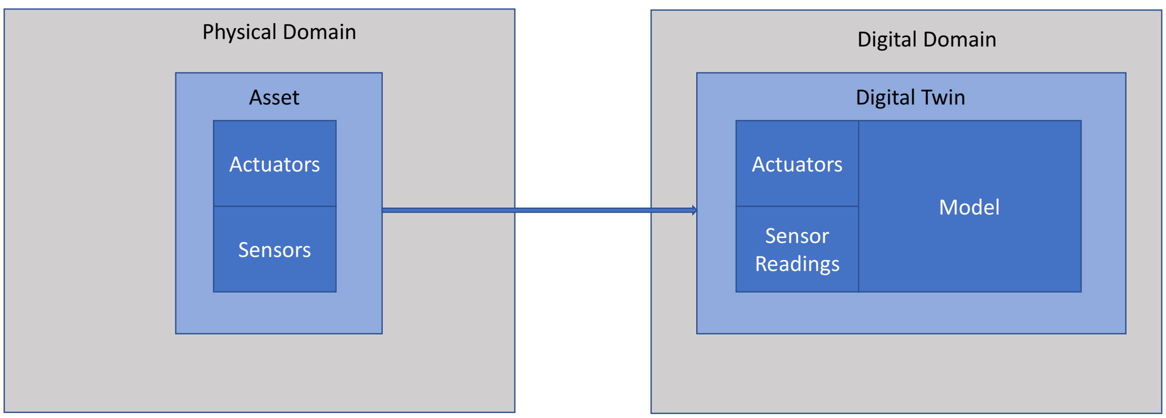

Digital Twins

Examples

Digital Twins are applicable to a very broad range of assets. A few examples:

- Sensor - with the Digital Twin reflecting the current value of the sensor and its state (reporting or offline)
- Vehicle subsystem - whose values may be produced by other Digital Twins such as the sensors mentioned in the previous bullet illustrating how Digital Twins can be composed into larger collections and composites via connections. Examples:
  - Engine
  - Brakes
  - Steering
  - Lighting
  - Entertainment
  - Climate control

    - Environment sensors – lidar, radar, bluetooth proximity sensors
- Vehicle – a composite of vehicle subsystems representing the entire vehicle.
- Other examples related to vehicles but not a component of the vehicle might be Digital Twins representing the environment in which the vehicle operates:
    - Road segments
    - Roads
    - Traffic lights
    - Traffic sensors
    - Intersections
    - Pedestrians
    - Local weather
    - Traffic conditions

Logical Architecture

As a model of an asset in the digital domain, the Digital Twin may perform one or more of the following functions with respect to the asset:

- Collect sensor data from the asset and its environment and reflect the asset's current state in the digital domain. The state may then be visualized in the digital domain and/or recorded for later analysis. A typical monitoring scenario.
- Collect sensor data from the asset and its environment and apply it to a digital model of the asset in order to issue control actions to the asset's actuators to optimize its future behavior. A typical control scenario.
- Apply simulated sensor data to the model to perform what if analysis on one or more control strategies applied to the asset when it is in the simulated state. A typical prediction scenario.

These functions are supported by the various logical components of the Digital Twin architecture:

- Data collection
- Data enrichment
- Actuation
- Modeling
- Situational awareness
- State management
- Analysis

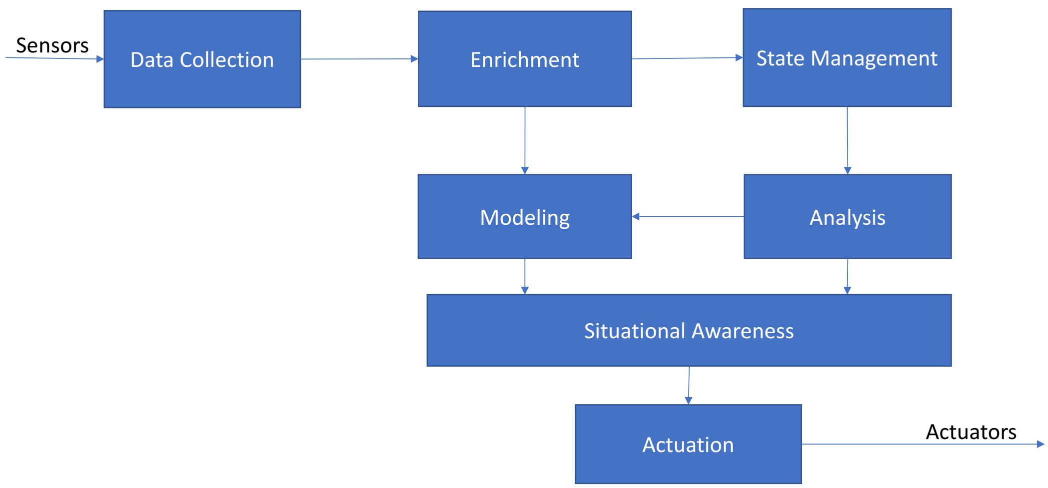

Digital Twin Logical Architecture

The Digital Twin collects data from sensors as event data (aka streaming data) reflecting the current state and behavior of the asset and the environment in which the asset operates. Data may be collected directly from sensors or more indirectly through data feeds. For example, a water temperature sensor in an engine directly reports temperature while current weather conditions are acquired as event data from weather forecasting systems. The data may be enriched or augmented with additional data from non-streaming sources such as the vehicle's operating history or detailed surveys of known road hazards.

Collected and enriched data is presented to the Digital Twin modeling subsystem to produce a real-time model of the assets current state and, possibly, projected state. Simple models may be algorithmic in nature. More complex models may involve deeper analysis using statistical methods or machine learning to produce the asset's current state and/or projected state. For example, engine sensors stream engine operating parameters to the Digital Twin, with the digital twin's model converting these parameters into current speed and fuel consumption and converting its current location, maps of the current road and the destination of the vehicle into its projected location for the next 15 minutes.

Collected and enriched data is also routed to the state management component of the Digital Twin where it may be recorded for later offline analysis by the analysis component. For example, the twin may record several years of operating data for a particular vehicle model and then use statistics and machine learning methods to gain insight into the behavior of the vehicle as it ages. Note that the result of such analysis may be trained models that can be referenced by the modeling component to assist in operating the vehicle.

With the Digital Twin faithfully representing the asset, situational awareness is applied to determine if control actions or notifications to interested parties are warranted. For example, if historical data analysis shows that the vehicle should not be driven at high speeds for more than 20 minutes to maximize the vehicle's lifetime and real-time modeling indicates that the vehicle has been operated at high speed for 20 minutes, situational awareness would detect this situation based on the operating model and the historical analysis and advise the vehicle operator that optimal operating ranges are being exceeded.

If direct action is warranted, the actuation component is invoked to deliver control commands back to the asset. Notifications may also be delivered to other systems such as a notification of the vehicles speed and projected future location reported to a regional traffic management system. For example, the projected location of the vehicle can be compared to known road hazards and the vehicle advised to slow down or warn the driver of a hazard that will be encountered in its projected path.

Simulation

If the Digital Twin acquires data from the physical world, the digital twin models the current state and behavior of the asset. However, if the digital twin acquires data from a synthetic source, the digital twin models the state and behavior of the asset as if it had been either subjected to the synthetic inputs or operating as implied by the synthetic inputs. For example, the vehicle twin might be driven by synthetic engine temperature and vehicle speed.

Once the model is accurate it can be used to simulate possible future states by driving it with fabricated sensor data designed to explore the assets operating boundaries. Integrated with an optimization engine, more efficient operating parameters can be explored and the optimal behavior recommended to the user.

This implies Digital Twin can be used both for modeling the physical domain and exploring simulated alternatives. However, to make this work the Digital Twin must be able to distinguish between physical data and simulated data and the results of applying models to the two types of data. This restricts the actuator component from sending commands based on simulations to the asset and so that simulated state does not get comingled with actual state. The mechanics of maintaining this separation in an advanced implementation will be discussed in the section on the VANTIQ Digital Twin architecture.

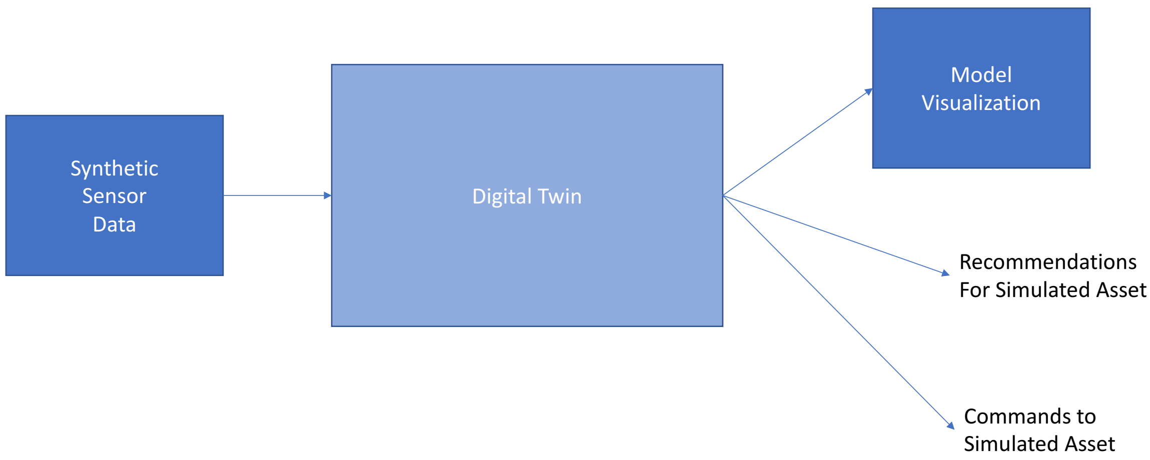

Simulating Scenarios on the Digital Twin

Visualization

Once an accurate model of the asset is available in the digital domain, its status can be considered as a proxy for the asset's status. If a suitable visualization is available the status can be displayed on a virtual version of the physical asset. For example, if the asset is an automobile it might be possible to build a 3D visualization of the car and then zoom in to the driver's seat to view the simulated speed and direction of the car or view a visualization of the actual instrument cluster in the simulated car.

Big Data Analytics

Once enough historical data has been accumulated, machine learning can be applied to the data in an attempt to find previously undetected trends in the operating data. This can be used to boost efficiency or for the early detection of potential failures.

Through the assets actuators the behavior of the physical asset can be modified by the situational analysis component of the digital twin.

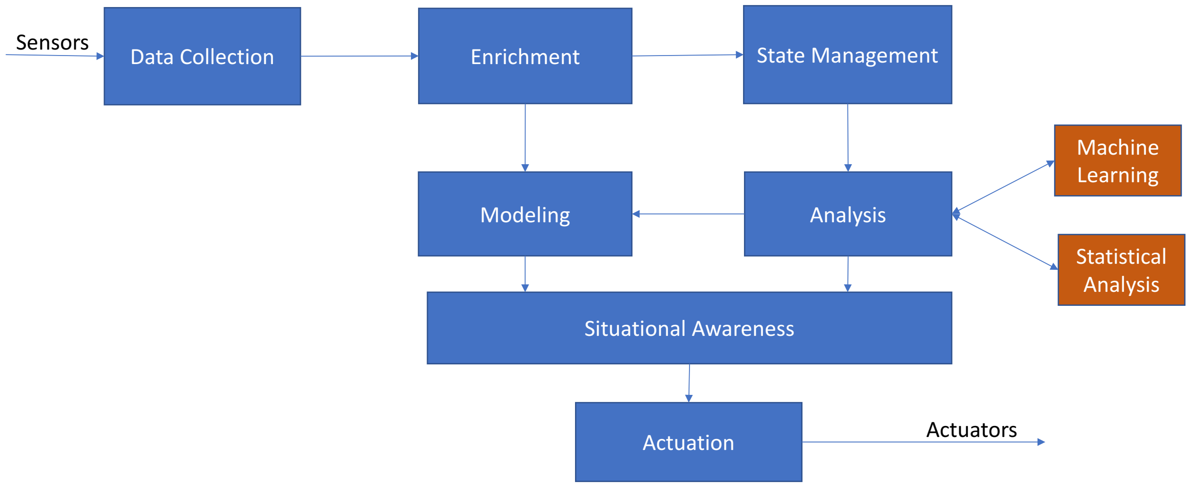

Applying Machine Learning to Digital Twins

VANTIQ Digital Twin Architecture

VANTIQ is an excellent platform for hosting Digital Twins. They may be built in a custom fashion or constructed on the VANTIQ Digital Twin Architecture. The Digital Twin Pattern exploits advanced features of VANTIQ to reduce implementation costs for Digital Twins. The Digital Twins can be built using VANTIQ's low code capabilities in the App Builder and Collaboration Builder and platform features such as distribution.

Overview

The VANTIQ Digital Twin Architecture defines a common model for describing a digital twin type along with App Builder patterns for implementing the components of the Digital Twin and Collaboration Builder patterns for interacting with the physical world and its occupants.

Types

The Digital Twin is represented by a VANTIQ type or set of types that provide state management. The Digital Twin pattern prescribes the type define a standard set of properties common to all Digital Twins that are used by the Digital Twin component patterns to manage instances of the twin. Any types representing the digital twin or the data from which digital twin state is computed must contain these properties.

Specifically, types representing a Digital Twin must have the identifying properties:

- name – the name of the digital twin instance. Each twin of a given type must have a unique name. Names may be human readable or encoded. The only requirement is that they be unique.
- twinType – the type or class of the digital twin. For example, vehicle, engine, steering, speedSensor, engineTemperature. This determines the set of relevant components applied to an instance of the digital twin. Each twinType must have a unique name.

The combination of **name** and **twinType** uniquely identifies each digital twin instance.

Relationships

Digital Twins can be related to form larger aggregates that represent a higher level digital twin. For example, automotive subsystems such as engine, breaks, steering, entertainment are aggregated together to represent a vehicle. When placed in a vehicle setting the subsystems interact in unique ways. For example, if the road sensing subsystem detects an obstacle in the vehicle's path it advises the breaking system which then applies the brakes. Note that the interactions may be direct service invocations such as recognizing an obstacle resulting in the application of the brakes or they may be event notifications such as recognizing an obstacle and then the braking system decides to apply the breaks or the data may be transmitted to a higher level system that takes the notification of an obstacle, the vehicles current speed, the direction of the road immediately in front of the vehicle and determines if the brakes should be applied. All these options are available from the VANTIQ digital twin pattern although the event-driven responses are best optimized by the VANTIQ platform.

Note that such relationships can take other forms. For example, we may have a digital twin for the vehicle and another that represents the road with the road twin made up of several subsystems of its own. In this case the relationship is not an aggregate relationship but a cooperative one. For example, a vehicle intends to go from A to B using the road R. A third type of relationships is aggregation such as the collection of all vehicles traveling on a given segment of the road at a specified time.

These relationships can be modeled in the basic Digital Twin data model in a common fashion by declaring a role and a set of twins that play that role. Note that this representation works for composite relationships (the subsystems of a vehicle), the uses relationship and the collection relationship.

Loosely-coupled relationships

Digital Twins may also represent data that is only loosely connected to the Digital Twin of interest – the vehicle in our example. For example, weather conditions may be important to the operation of the vehicle and a digital twin can be established that represents the weather with each instance of the weather digital twin representing the weather in a different location.

Local disasters, points of interest, shopping, food, etc. might all be represented by digital twins depending on the fidelity needed to represent the physical asset. For instance, if all we need to know about is its location, type and hours of operation it can be represented by a very simple digital twin or maybe no twin at all.

A standard schema is also needed for linking digital twin instances as they are combined in various relationships. The Digital Twin Architecture represents relationships as name, value pairs associated with each instance. These relationships are represented by three properties associated with a twinType:

- component – relates a digital twin to its component digital twins. That is, the components are considered digital twins that are "part of" this digital twin. For example, a digital twin's representing a vehicle has the components engine, steering, brakes, etc. that make up the vehicle.
- connection – relate digital twins that have a relationship that is neither a component relationship nor an aggregate relationship. For instance, vehicles are related to the road on which they are traveling.
- aggregation – a collection of related digital twins. For example, a road might maintain a collection of the vehicles that are travelling on it. Note that there is some overlap between connections and aggregations as they may be two different views of a one to many relationship. For example, many vehicles travel simultaneously on a single road while a single vehicle is travelling on exactly one road at any time.

The various relationship properties contain sets of key-value pairs that represent specific relationship roles. For example, for a vehicle

- component: engine – reference to engine twin
- component: brakes – reference to brake twin
- connection: road – reference to road on which vehicle is travelling
- aggregation: linked – the set of nearby vehicles linked to this vehicle by active telematics connections.

Components

Data Collection

Data is acquired from a variety of sensors through sources and added to the twin's readings. Generally, the readings are not the raw readings from the device but readings that have been transformed into a more usable form. For example, a voltage may be read from a temperature sensor. It needs to be converted from a voltage to the corresponding temperature. The voltage readings may dither due to noise in the sensors electronic circuits. The transform can filter out minor temperature changes due to noise. In some scenarios, it may be more efficient to track significant changes to the temperature rather than all temperature readings and a filter can be specified to suppress events that do not exhibit the desired magnitude of temperature change.

The readings may be processed immediately through situational awareness to drive real time reactions to changing conditions and produce commands to actuators to change the behavior of the asset to better satisfy its operational goals.

Eventually, relevant readings are stored for future processing by machine learning, model visualization and simulation. These readings are stored in either VANTIQ digital twin compatible type (see above) or in a related type defined by the user to be optimal for the data in the domain represented by the Digital Twin.

Once stored, the readings may be accessed for historical analysis and for visualization activities.

Enrichment

Collected data may need to be augmented or enriched with other available data. For the vehicle examples, data from enterprise systems such as maintenance history may augment the sensor data providing richer data to drive modeling and analysis.

Enrichment usually involves accessing information in systems external to the Digital Twin.

Modeling

The modeling component is the proxy for understanding the behavior of the asset. It simulates the behavior of the asset to produce the same results the asset would produce but, in the case of the digital twin, in the digital domain.

It may be used to produce instructions to the asset, delivered via actuators, to change its behavior based on the model's prediction of future optimal states for the asset. For example, in our vehicle example, detection of an object in the direct path of the vehicle may turn into braking commands or steering commands to cause the vehicle to avoid the object.

In all cases, the model is used either to mimic the asset, predict future states of the asset or to augment the asset's behavior (as in the example above) based on the results of predicting future states of the asset and acting to modify those states.

Actuation

The actuation component closes the loop by sending commands from the digital twin to the asset it models. Actuators are implemented via sources that communicate with the asset or a controller that acts as the assets external interface. For example, in IoT systems it is likely control is mediated by a gateway. In industrial control systems, control is likely mediated by the SCADA system.

Analysis

Modeling activities are complemented by analytics activities. While some modeling is generally performed in real-time as part of processing inbound real-time data, more complex analytics may be applied on a lower duty cycle. For example, the system may be updating machine learning algorithms in the background with the models produced by the machine learning system subsequently being used in the real-time processing. The vehicle digital twin may be sending is sensor readings to improve the training of a neural net that detects obstacles. Periodically the improved neural net is installed in the real-time processing flow to improve the real;-time detection of obstacles. The updating of the net occurs in the background while the identification of obstacles using the net occurs in real-time as part of the modeling activity.

Situational Awareness

Situational awareness is the real-time decision-making component of the Digital Twin. The model produces a representation of the current state and, possibly, the future state of the asset given its current inputs. Situational awareness inspects this state and determines if this state warrants intervention either via actuators to change the behavior of the asset directly or via collaboration with relevant users to determine if some change is required. For example, detection of an obstacle may warn the driver of the issue and expects the driver to help avoid the obstacle.

State Management

State represents the current state of the asset. It is intended to be a high-resolution model that represents all relevant characteristics of the asset as it is operating. It is the target of visualization since the visualizers attempt to display a high fidelity virtual asset that entirely mimics the physical asset. This state model may also be the target of situational awareness to drive lower duty cycle decisions than those driven by the more immediate needs of responding to real time sensor values.

The core state is generally represented in the type(s) that model the Digital Twin (see above). Ancillary types may be used to model other aspects of the asset.

Some state information is transient and not recorded persistently as it is needed only for the real-time processing (see data collection above).

Historical data is used to drive analytics. Typically, such data is stored in big data systems where it is easily accessed by standard analytics and machine learning systems.

Some short term history may be stored in the Digital Twin (see state management above) but long term history is generally stored in an external system optimized for analytic processing.

Enhancing Digital Twins

VANTIQ provides a number of powerful facilities for enhancing the Digital Twin experience. We discuss three such facilities:

- collaboration
- distribution
- isolated simulation
- system integration

Collaboration

As with all event-based applications, Digital Twins can benefit from advanced collaboration with the users of the assets modeled as Digital Twins. Continuing with the vehicle model, collaborations may be ongoing with the vehicle operators or, in advanced situations, with traffic managers managing vehicle traffic in a region. Air traffic controllers would be a current example with the twinned vehicles being aircraft.

We can easily imagine Digital Twins giving operators advice on the optimal approach to operating their vehicles in the current situation. Of course, vehicle operators quite often ignore such advice and do what they want. An advanced collaboration system allows the Digital Twin to adapt to the operator's actual behavior producing new suggestion for the operator and perhaps driving actuators on the vehicle to better optimize the behaviors actually chosen by the operator.

VANTIQ's collaboration builder provides a powerful mechanism with which to build such adaptive user interactions.

Distribution

In many use cases Digital Twins are best implemented as distributed systems. Data collection occurs throughout the network and data is most efficiently aggregated close to its source. In addition, real-time decision making and actuation is best implemented close to the actuator as networks delays will quickly combine to make remote actuation too slow to be effective or, in some cases, dangerous.

VANTIQ offers a transparent distribution model that simply integrates edge processing, regional processing and centralized processing so that each participating system performs those functions that are best located at that level of the processing hierarchy. VANTIQ's fully dynamic implementation allows the distribution of the application to be changed at any time including in response to real-time changes in the environment.

Simulation

Simulation can be applied to the Digital Twin to observe the behavior the asset would exhibit if the simulated stimuli were applied to the physical asset. Simulations are applied through simulated stimuli presented to the Digital Twin via sources that produce the simulation inputs.

In most environments simulations are expensive to implement because they must be built into the Digital Twin architecture from the start so that actual data and simulated data can be operated upon separately and simulated results routed to a system other than the actual actuators for the physical asset.

VANTIQ supports simulation applied to any VANTIQ Digital Twin through VANTIQ's built-in user isolation architecture. Essentially, a separate space is established and the digital twin is loaded into that space. It is now completely isolated from the instance of the Digital Twin that is collecting data from the physical asset. In the isolated space, the data collection is configured to ingest simulated data and the actuators are replaced with simulated actuators without changing any of the code in the digital twin.

Visualization

Simple visualizations are easily constructed with VANTIQ but more sophisticated visualizations such as 3D models can be driven by VANTIQ by incorporating the visualization system as another sync for data being produced by the Digital Twin.

System Integration

Digital Twins may need to be integrated with a wide array of enterprise systems. By default, Digital Twins built in VANTIQ are automatically advertised as services that can be used to deeply integrate the Digital Twins with other systems.

VANTIQ Enables the Digital Twin

The VANTIQ Digital Twin Architecture not only supports, but truly enables the concept of the Digital Twin. With VANTIQ, developers can build high fidelity Digital Twins more rapidly and effectively than ever before. Please contact VANTIQ (info@vantiq.com) to further discuss this opportunity.