

# Connected Complexity

## The Padi Platform: Open-Source Tools For Open Data

---

We are at the cusp of a “perfect storm” of smart systems innovations. Multiple parallel technology developments, including data modeling and machine learning are accelerating and enabling more complex and adaptive systems such as digital twins. Along with the value these innovations bring, however, so too grows the complexity of connecting and integrating machines, equipment and data in a meaningful context. A new open-source software tool, Connection Profile, enables simple, durable and context-sensitive integration between complex systems without wasteful custom development.

---


WHITE PAPER

Harbor  
Research

# Multiple technology developments are accelerating and enabling more complex and adaptive systems

## THE REALITIES OF “CLOUD COMPUTING”

Back in the 1990s, when the web (as opposed to the internet) was just getting going, people used to represent the internet with a little cartoon of a cloud. Countless presentations to venture capitalists featured little pictures of desktop computers, modems, hard drives, etc., all connected by wiggly lines to the cloud-cartoon. When technologists needed investment, they would just draw a diagram where everything was connected to a cloud. Here’s a classic example<sup>1</sup>:



Despite the naiveté of these pictures, they portrayed an important truth: There was one cloud, not many. Microsoft Cloud did not exist back then, nor did Apple Cloud, Google Cloud, Amazon Cloud, Salesforce Cloud or ABC Widgets Cloud. There was simply “the cloud”—the one and only shared cloud of information, the common property of humanity. Nobody owned, or could ever own, “the cloud” of the 1990s dot-com diagrams.

Further, that singular cloud embodied a huge idea about connectivity: When you connected to the one-and-only cloud, you were connected to everything else that was connected to the cloud. Like all peer-to-peer

<sup>1</sup> <https://www.pcmag.com/news/what-is-cloud-computing>

(P2P) architectures, the original design of the internet embodied a radically decentralized view of computing. In the internet depicted in that early cloud, your personal computer wouldn't have "information in it." Rather, your personal computer would be "in the information."

In a mere twenty years that idea has been completely inverted. Even though cloud providers would argue that URLs and Linked Data (LD) exist to link everything, the private clouds of today have no inherent ability to interoperate significantly with other clouds, private or public. These private clouds are all based upon proprietary designs built to compete, not cooperate. Cloud-computing is in fact the same old corporate client-server computing we had before, just dressed up in cloud-clothing.

### **DATA ISLANDS ARE THE ENEMY**

Contrary to popular practice today, the "big cloud" world of the 21st century will see little benefit from keeping device data and its resulting insights sequestered in corporate clouds with little or no interoperability. That situation may serve the narrow, short-term interests of corporations seeking to protect market share by hoarding data and locking customers into their platforms. But for people interested in furthering the cause of connectivity and shared information, the "data islands" that result from this practice are a constant source of frustration and missed opportunities.

What is wanted instead is the integration of data creation and data processing, regardless of manufacturer or operating system, and the liquid flow of data streams among all sensing and computing devices, wherever they may exist in the network topology.

The internet was the result of simple, inherently extensible concepts that have served the world very well in the years since it was designed. No matter what people chose to build on top of its basic architecture, the internet dutifully connected clients and servers and delivered their packets of information gracefully. That was what it was designed to do.

However, there are other equally valuable low-level network services that the internet was not designed to perform. Possibly the most important of them is the vendor-agnostic, automated integration of internet-connected devices and their data streams. During the internet's lifetime, a number of simple and extensible concepts for doing this surfaced, but for a variety of reasons—notably that they were ahead of their time and met with significant corporate resistance—none of these solutions caught on.

**The purpose  
of integration  
platforms is to let  
system developers  
work more  
collaboratively**

**Digital twins  
are combining  
networking,  
software, data  
and analytics  
innovations  
in new and  
unprecedented  
ways**

## **ENTER THE PADI PLATFORM**

The latest integration platform to cross our radar is called Padi, which is the brainchild of IoT and automation pioneer Anto Budiardjo. Padi means “rice plants” in Indonesian, and the rice-bowl in the company’s logo represents the thousands of “grains” of device information that are of interest at any moment, compared to the billions of grains that are irrelevant to a given user.



The purpose of integration platforms is to let system developers work more collaboratively by simplifying the bewildering number of endpoints and devices they have to deal with on a daily basis. Budiardjo realized some years ago that most integration platforms, though they used the internet to some degree, were not genuinely internet-centric. They kept solving the “integration problem” over and over again but often with too much customization.

Existing technology has proven cumbersome and costly to apply with many conflicting protocols and incomplete integration tools. Padi, by contrast, integrates data flows in a reusable way with generalized, extensible concepts similar to those that underlie the internet itself and uses an innovative open-source concept called Connection Profiles to allow the connection and integration of any internet endpoint to any other.

## **THE PROMISE OF MACHINE DATA and DIGITAL TWINS**

Before delving further into the new thinking that makes this story possible, let’s talk about why it’s necessary at all. Machine data can offer extraordinary business advantages to both the companies that manufacture and support machines as well as the users of machines. The ability to detect patterns in aggregated data is the “holy grail” of smart systems. New machine data and learning technologies enables not only data patterns but a much higher order of intelligence to emerge from large collections of ordinary sensor and machine data.


The fact that a rapidly expanding range of devices have the capability to automatically transmit information about status, performance and usage and can interact with systems and people anywhere in real time, points to the increasing complexity of adaptive machines and systems. The emergence of digital twins is a very good example of this.

What is a digital twin? The Digital Twin Consortium (DTC) defines a digital twin as “a virtual representation of real-world machines and processes,

synchronized at a specified frequency and fidelity.” Digital twins of complex machines or processes actually consist of many models of component parts and sub-systems, each one pinned to measurement data collected by sensors. Digital twins can include current actual operating information, as well as analytics capabilities and a recommendation engine to enable real-time operations. Whether driven by historical or live data-feeds, they provide users with visibility into how the machine or entity is performing at any given moment, as well as what to anticipate in the future.

While digital twins are cleverly combining networking, software, data and analytics innovations in new and unprecedented ways, development of these solutions are not without challenges.

### DIGITAL TWIN OF TURBINE



**Digital Twin  
models are  
reaching beyond  
simple standalone  
machines to  
more complex  
applications and  
systems**

### INTEGRATION ISSUES CREATE DIGITAL ORPHANS

All of these new machine innovation opportunities presuppose technologies, standards and tools that make connecting equipment and integrating machine data, cheap, easy and seamless.

How will the physical and digital integration aspects of these new opportunities actually take place? How will virtually any device, right down to a lowly light bulb, become a peer that connects at will to the global data network? Obviously, billions of devices of wildly varying types cannot each

**The Padi  
platform's  
magic lies in  
how it achieves  
connecting  
and integrating  
"anything to  
anything"**

receive individual attention and configuration, or conform to elaborate a priori specifications. If it literally takes a network engineer to screw in a smart light bulb, digital systems and twins are never going to work.

Many schemes and so-called "standards" for device connectivity and device data integration already exist. But of course, all those "solutions" add up to one big problem. We don't want many standards; we want one.

**LEVERAGING NEW TOOLS TO ENABLE DIGITAL TWINS**

We are at the cusp of a "perfect storm" of smart systems innovations. Multiple parallel technology developments, including data modeling and machine learning are accelerating and enabling more complex and adaptive systems such as digital twins.

From creating a feedback loop for improving product designs, to gaining better visibility on the operation of industrial machines and systems in order to optimize them, to developing new use cases such as predictive maintenance, digital twins have become a topic of wide interest.

Along with the value these innovations bring, however, so too grows the complexity of connecting and integrating machines, equipment and data in a meaningful context.

Applications and use cases enabled by digital twins are very diverse, depending on the nature and characteristics of the machine or equipment involved and the complexity of systems that can encompass multiple devices, machines and entities. To date, digital twin development has largely been focused on stand-alone machines and equipment segments, such as a pump, a wind turbine, or similar, where the digital twin is a digital model based on the 3D CAD and CAE analysis that was developed during the design phase of the equipment.

As technology matures, digital twin models are reaching beyond simple standalone machines to more complex applications involving multiple interrelated machines or devices that are operating as a system, such as a manufacturing production line where the process behavior that's being modeled can span many activities during operation and throughout the system's lifecycle, including its supply chain and other related performance data. When fed with historical or current actual data, digital twins can be used to predict the behavior of a real-world physical system, similar to the way that atmospheric models are used to predict the weather.

The growing complexities involved with developing digital twins cannot be understated. Challenges with connectivity and integration, normalizing diverse machine data types as well as the integration of complex models

and simulation software. In many ways, the tools we are working with today were not designed to handle the complexities of today's digital twins or similar machine data innovation. These challenges are diluting the ability of organizations to efficiently and effectively manage development.

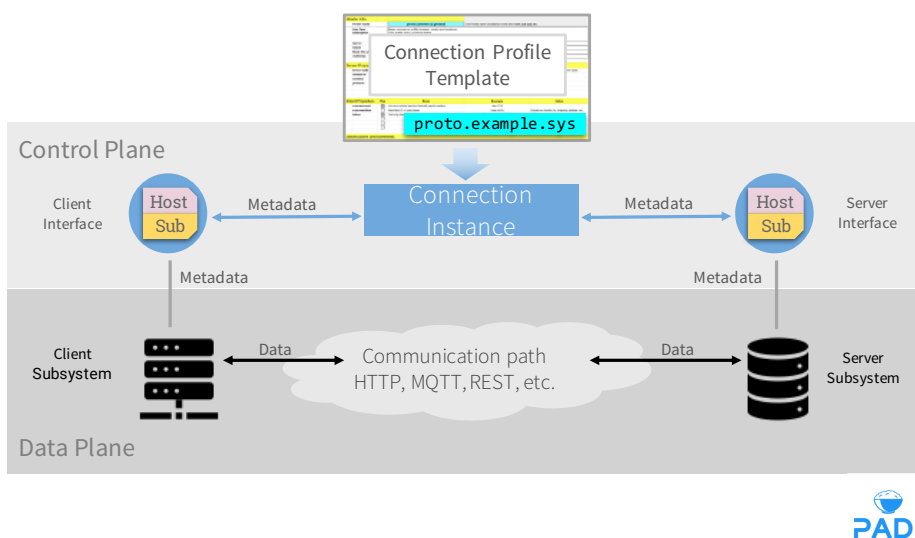
Customers expect evolving platforms and software tools to be functional, ubiquitous, and easy-to-use. Within this construct, however, the first two expectations run counter to the third. In order to achieve all three, a new approach is required -- but what kind of approach?

### CONNECTION PROFILES ADDRESS INTEGRATION ISSUES

For many years now, systems developers and users have faced the rapid proliferation of many random endpoints trying to connect to other endpoints in the absence of any meaningful context. The Padi platform's magic lies in how it achieves its connections among diverse equipment and systems. Its designers set as their goal the Shangri-la of integration—connecting anything to anything, no matter what or where a given “thing” was. This goal came out of Budiardjo's long work in the device integration field, where he spent many years designing specialized drivers and one-off middleware for project-specific issues. As anyone who has done this type of work knows, the solutions to these difficult challenges are not immediately reusable in other contexts, and the engineering also tends to be brittle and break under the slightest stress—for example, a single URL changing.

**The developers of the Padi platform understand that integration platforms need to automate internet-connected devices in a straightforward manner**

### CONNECTION PROFILE



To be successful, integration platforms need to go to a much deeper, more abstracted level. Connection Profiles are not defined for specific projects, but rather for specific use cases. For example, to connect a certain type of client to a certain type of server for a specific purpose. This is done in a



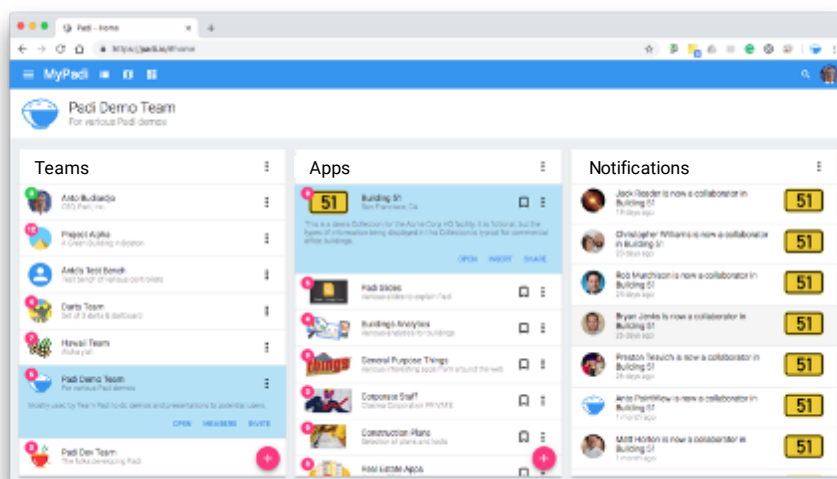
**Users need software tools and methods to automate development and facilitate re-use**

way that allows that Connection Profile to be used for any instance where a client and server of the same specifications are being connected.

Connection Profiles itemize the characteristics of a “client” (for example, a device) or a “server” (for example, a data repository). (Note, however, that “client” and “server” are terms used largely for convenience. Any entity utilizing Connection Profiles could function as a client and a server simultaneously.) Connection Profiles have unique names, just like internet domains, and there is a registrar of Connection Profiles to ensure that they retain the necessary clarity.

Finally, integration platforms need to maintain a connection instance over the course of its lifetime. Anyone with real-world experience integrating internet-connected devices will immediately recognize the importance of this statement. Things in any system change over time—URLs get modified, software gets updated, something physical gets broken or replaced, and so on. With Connection Profiles, the connection becomes dynamic. The server tells the platform that X or Y has changed, the platform tells all clients, and within moments the connection is up and running again—with minimal downtime and no human intervention. You can swap pieces of equipment or modify the software, and the connections will continue to work.

**COLLABORATION TOOLS**



The developers of the Padi platform understood that integration platforms need to unify and automate internet-connected devices in a straightforward manner. The “things” that Padi connects typically have their own native UIs, encoded in HTML, that are presented on a dashboard (via iFrames) as if under a “single pane of glass.” Padi considers these UIs to represent “apps” that users invoke just as they would use them in any other context.



Sometimes these web apps connect directly to the service or thing being depicted. But often in the IoT, devices and services are not on the public internet. They may be in a building on a private network for security purposes, or behind some credentialing mechanism that requires an extra layer of intelligence. One aspect of Padi's design—the Padi App Server—takes care of navigating automatically to those connected things regardless of their public/private status. From the user's perspective, it's the same experience whether they're on the public internet or not.

## **THE BUSINESS CONTEXT**

Importantly, the Connection Profile doesn't limit itself to technical information about how the systems and things are connected. It can also include business information. For example, there may be a property called "Cost" that would indicate the charges per period of time that you would incur to use a specific source of data or information. Specific instances of the Connection Profile may also include license information, security tokens, or other parameters of the business relationship as dictated by the use case for which the instance was created.

Equally important is the fact that the data transferred between server and client does not go through Padi at all. The protocol used by the endpoints could be anything. Padi enables the endpoints to communicate across the internet using parameters provided by Padi. Nor does the UI connection itself go through Padi. Padi simply provides the metadata and orchestration.

Connection Profiles can be thought of as blank contract templates. Specific instances of Connection Profiles function as actual contracts. If you abide by the rules stated therein, the instance will take care of everything covered by that contract. Obviously, this is done with an eye to context, because Padi wants to match things that belong together for a business reason—for example, because they're in the same building, or are owned by the same company.

In the world before Connection Profiles, you'd decide on a project-by-project basis that this needs to talk to that, and you'd begin a very involved and expensive development cycle involving engineers in a room writing specifications, deciding what protocols or standards to use, and then getting custom code written. Then you'd have to install it, with people from both sides of the equation present, and pray that for the life-cycle of that equipment (often decades), the integration would continue to function.

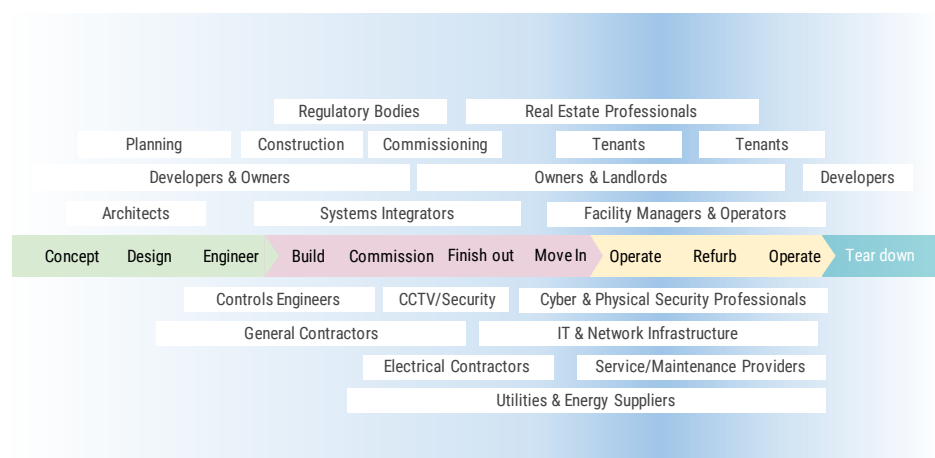
Many consultants make a lot of money doing this type of integration work, and Padi—with its open-source and extensible Connection Profiles—is a disruptive technology in their world. Instead of developing project-specific

**Integration  
platforms need  
to maintain  
a connection  
instance over  
the course of  
its lifetime**

**Integration platforms need to be future-proofed by making the fewest possible assumptions about the nature of networked objects and the data they produce**

customized software, you create a Connection Profile instance that says, in essence, “I want to connect a train-ticketing system with an energy management system.” Any necessary development becomes part of a product-development process, not a project-development process, and all development is reusable and can be made available again and again. At installation time, the systems are connected with minimal if any configuration, such configurations being project-specific information.

### BUILDING LIFE CYCLE



### END POINTS

When telephones first came into existence, all calls were routed through switchboards and had to be connected by a live operator. It was long ago forecast that if telephone traffic continued to grow in this way, soon everybody in the world would have to be a switchboard operator. Of course that has not happened, because automation was built into the systems to handle common tasks like connecting calls. We are quickly approaching analogous circumstances with the proliferation of smart connected devices. Each new device requires too much customization just to get integrated into a system. We must develop software and methods to automate development and facilitate re-use, or risk constraining the growth of this market.

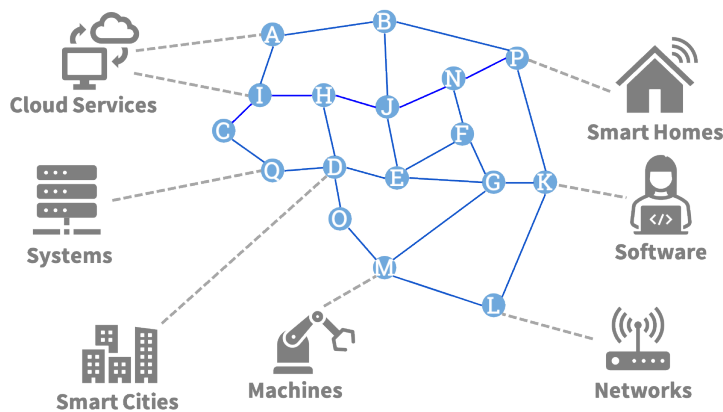
Radical new thinking about information technology must begin at the most basic levels, with new conceptions about the interactions of information with people, systems and devices. The Padi team is future proofing their innovations by making the fewest possible assumptions about the nature of networked objects and the data they produce, carry or process - the company takes a much broader, all-encompassing view of information.

## ADDENDUM

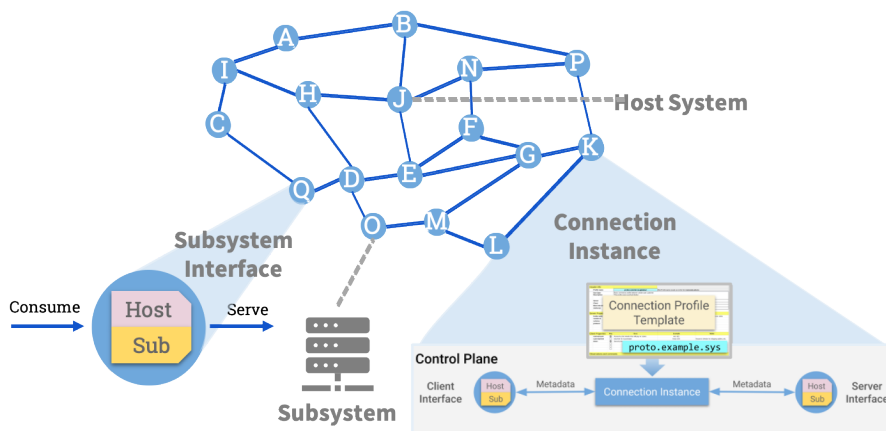
### A DIGITAL FABRIC TO ENABLE A “SYSTEMS OF SYSTEMS” APPROACH FOR DIGITAL TWINS

Recall that the Connection Profile schema are designed to integrate different domains with each other. This is keenly relevant to digital twins, which essentially represent an integration architecture. Connection Profiles, with their ability to robustly integrate connected devices, create a digital fabric for unifying the emerging phenomena of complex digital twins that simulate systems in the real world.

In the diagram below, internet endpoints (labeled with capital letters) are connected with Connection Profile instances. The resulting representation of the digital fabric.

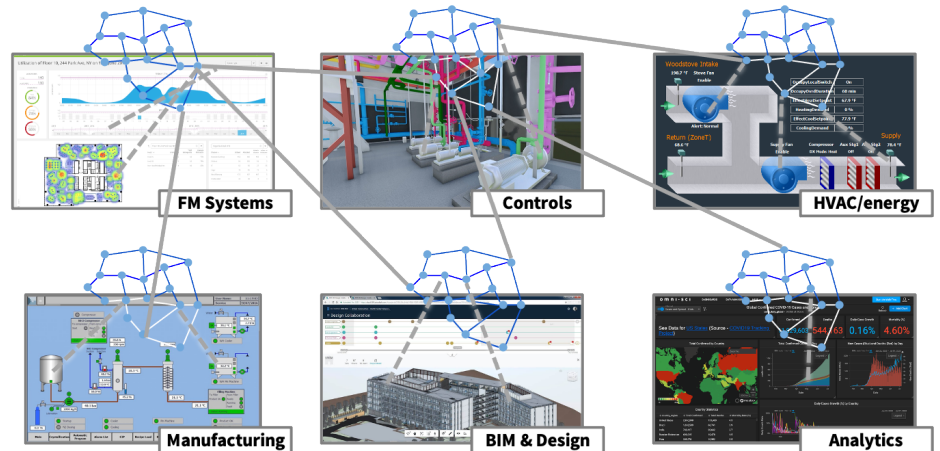


As noted above, Connection Profiles are designed to integrate different domains with each other. The digital fabric can be thought of as “an overlay” connecting real-world entities, thus enabling a “System of Systems.” digital twins are made up of diverse domains (HVAC, Security, transport, energy, etc.), each of which is effectively the systems being brought together in a System of Systems.

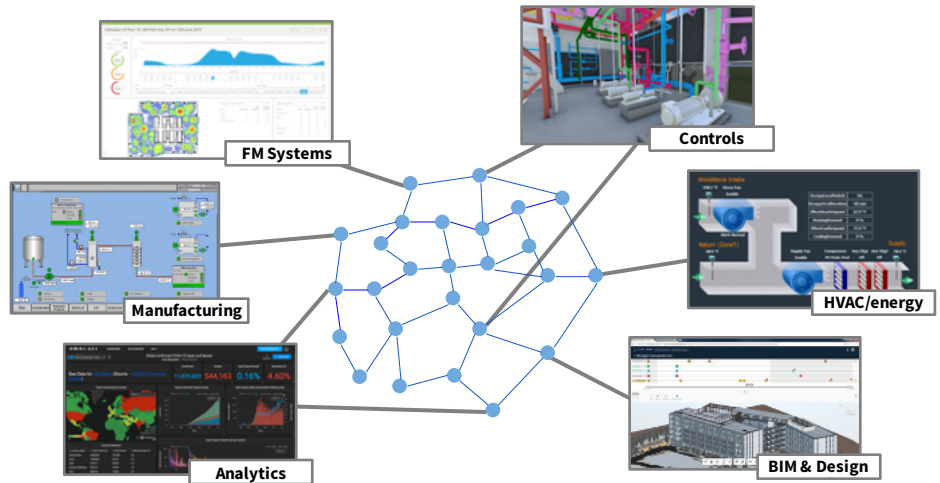


Connection Profiles can be used within a domain, though there are typically other existing connectivity mechanisms within specific domains created by people working in those domains. In the following illustration, the digital

fabric is being utilized in stand-alone systems to make them inherently able to integrate with other stand-alone systems, thus creating a System of Systems:

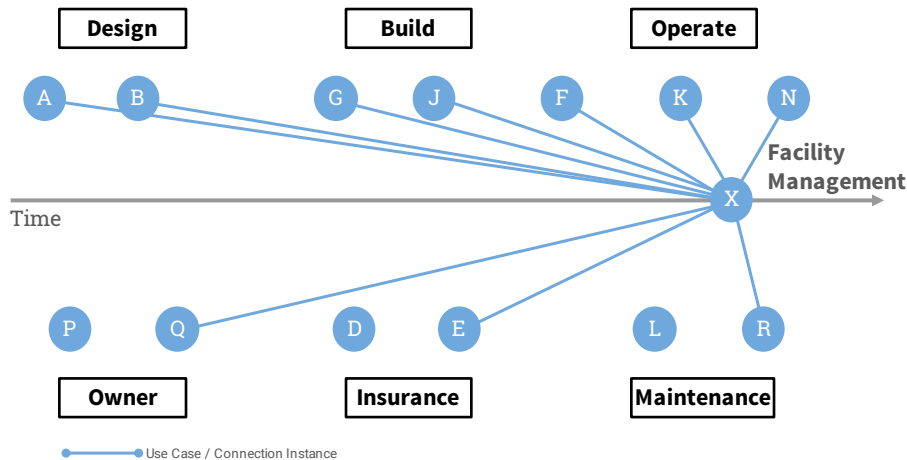


The next graphic illustrates the digital fabric used to integrate various domains in a digital twin:



Finally, we see Connection Profiles implementing the concept of a digital thread. In the diagram below, X represents a connected thing (say, a boiler) that starts at the left side of the timeline and travels (in time) from the design tool (before it exists physically) (A), to a Building Information Modeling (BIM) tool (B), to the beginning of its physical existence when it is added to the Bill of Material (BOM) (G), all the way to the Facility Management system (N). At each stage, X creates a connection instance to whatever tool is being used

with the appropriate protocol/schema. X collects and effectively aggregates them over time so a complete history of X can be created:



Connection Profiles in the Padi platform are effectively the codification of use cases, which lie at the center of digital twins systems and their impact on human understanding. Padi and Connection Profiles enable a simple, durable integration between complex systems, minimizing wasteful project-specific development spending. In so doing, Padi makes a great leap forward to that original vision of the internet as a single, seamlessly integrated cloud.

## IT'S TIME TO ADOPT A STANDARD FOR DIGITAL TWIN CONNECTIVITY and INTEGRATION

The confluence of fast, cheap processors and sophisticated but affordable software has made simple simulations viable for almost everyone. If you consider that the processors in a contemporary smartphone now exceed the power of supercomputers from the 1990s, it's clear that even the smallest firms can benefit from the power of digital twins.

And yet adoption of this emerging technology has been oddly slow to catch on, primarily because the aggregation of the many sophisticated models necessary to represent complex reality has remained outside our grasp. We believe that the absence of a connectivity and data integration standard has been the chief obstacle. Because the Connection Profile mechanism is open-source and in the public domain, we believe it represents an important step towards making Digital Twins a reality.

## Harbor Research

Founded in 1984, Harbor Research Inc. has more than twenty five years of experience in providing strategic consulting and research services that enable our clients to understand and capitalize on emergent and disruptive opportunities driven by information and communications technology. The firm has established a unique competence in developing business models and strategy for the convergence of pervasive computing, global networking and smart systems.

For more information visit [www.harborresearch.com](http://www.harborresearch.com)