



SkyFoundry Insider

Issue No. 32

January 2019

In This Issue

[The Edge and the Cloud](#)

[Challenges Addressed by Computing at the Edge](#)

[It's About More than Data Transfer Speed](#)

[Data Reliability](#)

[Isolation of Fieldbus Networks](#)

[Data Transfer Costs and Performance on Constrained Networks](#)

[Reduced Engineering Effort and Cost.](#)

[Application Reliability and Process Continuity](#)

[Does this Mean the Cloud is Dead?](#)

[SkyFoundry Recognized as Leading IoT Software Platform by Harbor Research](#)

[SkyFoundry Turns 10!](#)

[Haystack Connect 2019](#)

SkySpark Everywhere, Compute at the Edge, Distributed Computing and the IoT

To Support the Next Generation of IoT Solutions Analytics Need to Work at Every Level of the Architecture – from the Edge to the Cloud

One of the topics that has garnered a lot of recent press is the “Edge”. While the edge may become one of those over-hyped terms, the concepts are very simple and important to understand for IoT applications that serve the built environment.

When we talk about the “edge” or computing at the edge, we are referring to performing essential data acquisition and computation functions as close to the data source as possible. For example, running advanced analytics on small IoT devices mounted directly on equipment systems or embedded within the equipment controller itself.

The advances in data analytics for equipment and device data have transformed the operational efficiency for organizations that have embraced the technology. The first generation of analytics for the sensors, equipment systems and IoT devices focused on the cloud as the compute platform. The vast majority of products and services were designed with the requirement to transmit all data from equipment systems up to the cloud (or other centralized server) to produce analytic findings and assemble them into views and reports for the operator.

This approach may have been a natural way to start, but it is not viable for the realities of the IoT. The full benefits of data-oriented applications such as analytics can only be achieved with a software architecture that provides for “computing at the edge” as well as the cloud. This means that data acquisition, storage, processing, analytics and the generation of visualizations must be able to be performed at the edge, without any dependence on the cloud.

The Edge and the Cloud


The concept of edge computing is especially relevant given all of the focus we see on “the cloud”. Cloud computing provides many benefits but isn’t a panacea for all of the requirements encountered in applying IoT technologies to the built environment. The first generation of analytics applications for data produced by sensors, equipment and IoT devices relied on “computing in the cloud” – meaning that they were based on a requirement to transmit ALL data from equipment systems up to the cloud where analytics and the generation of visualizations, reports and notifications would then be performed. That can be a major hurdle in many applications.

“The reality is that it is not possible, cost effective or desirable to transmit every piece of data from every IoT device to the cloud in order to gain value from that data.

“An IoT technology platform needs to embrace the highly distributed and innately non-hierarchical nature of the IoT and support that with a corresponding software architecture.”



Challenges Addressed by Computing at the Edge




Let’s start is by thinking about the IoT in general. The IoT is actually a distributed computing challenge. As noted by Harbor Research, the reality is that it is not possible, cost-effective or desirable to transmit every piece of data from every IoT device to the cloud before we can gain value from that data. In fact, the world we experience every day is based on distributed computing. Think about it for a moment...

Bring up your browser. On your PC, your phone or your tablet...

Go to Yahoo or Google or your favorite site. Look up a subject of interest. Do a search. Boom! There it is. The information you wanted.

How did that happen? Did you upload and store all of the information to your computer first? To your cell phone? To your tablet? No.

Is all of the information aggregated and stored on a single server or location in the cloud? Did someone have to assemble and store it ahead of time in order for it to be to be searchable, accessible, viewable? Obviously, the answer is no.



Data analytics solutions need to embrace the highly distributed nature of the IoT and support the needs of diverse applications with a software architecture that enables computing to occur where it is most efficient, cost effective, and reliable. That means an edge-to-cloud-software architecture.

You request what you want when you want it. You search for what you want when you need it. You subscribe to news feeds that interest you. But you don't try to aggregate it all in one place. Because you can't. And, there is no need to. Nothing on the web works that way. Search doesn't work that way. When you type in a search, that request is dispatched to hundreds or thousands of computers. They all respond and then their results are shown as if they came from a single server. That is accomplished via a technique known as "map-reduce".

Yet most first-generation IoT data applications require all data to be sent to the cloud to be aggregated ahead of time in order to perform analytics and visualization. This conflicts with the reality of the IoT. You cannot bring every piece of data from hundreds, thousands, millions or billions of devices to a single server in order to be able to use that data, visualize it, analyze it, present it, and gain value from it. As the industry moves to more and more deployment of IoT devices and use of data-oriented applications, this limitation has become very clear.

Consider the example of a self-driving car. We can't be dependent on sending data to the cloud before deciding to activate the brakes. That data analytics process needs to occur in the vehicle – **at the edge**. Yet other applications are better served by aggregating data on a central server. Consider how mapping applications collect and analyze GPS data from mobile phones to identify traffic jams and direct us to the best route to our destination. That application is better served by the cloud. As this example shows, it's not as simple as "either or".

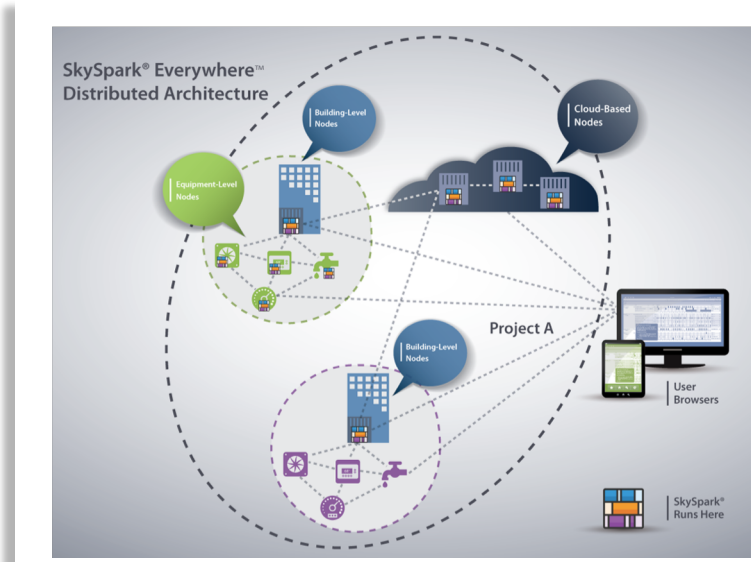
Data analytics solutions need to embrace the highly distributed nature of the IoT and support the needs of diverse applications with a software architecture that enables computing to occur where it is most efficient, cost effective, and reliable. That means an edge-to-cloud-software architecture.

It's About More than Data Transfer Speed

The self-driving car example highlights the "**data latency**" reasons we need to perform data analytics at the edge, but there are others that are just as essential to the next generation of the IoT.

Data Reliability

Data collection close to the end devices, whether it be a sensor or controller, increases reliability versus having to connect that data over the Internet to get it to the cloud. Networks do experience outages. Having an edge device that can store data for a short period of time provides a limited solution to that problem. While the data may be stored for transport to the cloud when the network is restored, no actual value (analytics, visualization, control decisions) is generated from that data while the network is down. With a distributed architecture that fully supports computing at the edge, those processes continue locally even while the edge is not connected to the cloud. That means in-building personnel have access to the full capabilities of their systems even if the connection to the cloud is unavailable.



Isolation of Fieldbus Networks

In many cases sensing and control devices communicate via networks that are not designed to go over the Internet or cannot do so in a way compatible with modern IT security requirements - for example, serial networks like RS485 and RS232, local wireless networks, and others. This means that some type of gateway node needs to be installed to isolate those networks and act as a data translator/gateway/forwarder. By supporting true computing at the edge, the edge node can do more than act as a gateway. It can perform the full stack of functions for data analytics, presentation and control **at a similar cost**.



Data Transfer Costs and Performance on Constrained Networks

This is one of the hidden costs of centralized cloud solutions. Data transfer to the cloud is not free. You typically see two areas of cost. First, most cloud platforms have charges related to data transfer – they charge based on the amount of data you send to the cloud. Equally important are the costs associated with transmission of data over cellular networks.

Increasingly, IoT devices are being connected via cellular networks. Sometimes this is done to avoid the challenges of integrating with corporate networks and IT security requirements. In other cases, it is done because no hard-wired network is available, for example, remote monitoring sites, agricultural applications and the like. The costs associated with transferring high volumes of data over cellular networks can significantly impact the economics of a project.

The capabilities of “computing at the edge” changes that equation. By computing at the edge, the data is collected locally, analytics are performed locally. The only data that goes across the cellular network are the results. ***This can reduce network data usage by a factor of 100 to 1 or even 1000 to 1.***

Reduced Engineering Effort and Cost.

In most systems where data is collected at the remote site and sent to the cloud, some type of gateway device is used on site. There are two costs associated with those devices. First, of course, is the cost of the device. What is typically more important is the cost to set up the device to connect to and collect equipment data and then send it along to the cloud. Then at the cloud application an additional engineering effort is required to bring that data into the cloud software platform. By applying the analytics application at the edge, the duplication of engineering effort is eliminated resulting in significant cost savings.



Application Reliability and Process Continuity

We have spoken about the reliability of data collection, but there is another aspect of reliability. Consider a remote site with local users of the analytics results. Perhaps they depend on those results to optimize a chiller plant or a production process. With an edge-computing solution they still have access to their data **and analytic results** even if the connection to the cloud or central server is lost. In many mission critical applications this is essential.



So Does this Mean the Cloud is Dead?

While the examples we have just reviewed present some of the very real factors that are driving analytics and similar data-oriented applications to the edge, processing analytics at the edge does not mean completely abandoning cloud or central servers – they have their place and will continue to do so. It would be a mistake to think of this as one **or** the other. What computing at the edge means is that the power and benefits of data analytics can be brought to the place in the architecture where they can most effectively deliver value. The Cloud is not dead, **but the Edge is now definitely alive**. And with SkySpark Everywhere's distributed architecture you get the full realization of a true edge-to-cloud IoT architecture.



SkyFoundry Recognized as Leading IoT Software Platform Provider by Harbor Research

Harbor Research

Analytic results are the most visible end result of the functionality provided by SkySpark's Informatics platform. They form most people's view of what SkySpark is all about. But the capabilities of the SkySpark Everywhere platform go far deeper. The layers of technology that make SkySpark's advanced analytics possible form a comprehensive end-to-end software platform for working with IoT devices and their data. From connecting to diverse IoT devices using multiple industrial and IT communication protocols, normalizing multi-structured device data, visualizing data, performing analytics to monitor device health as well as the performance of the end application, SkySpark provides a comprehensive end-to-end IoT architecture.

In December's **Harbor Research Technology Insight, THE PROMISE & PROGRESS of SMART SYSTEMS & IoT PLATFORMS**, Harbor Research reviews the marketplace for IoT software platforms as part of their assessment on the state of the market. We are excited and honored for SkySpark to be recognized as a leader in the IoT space by Harbor.

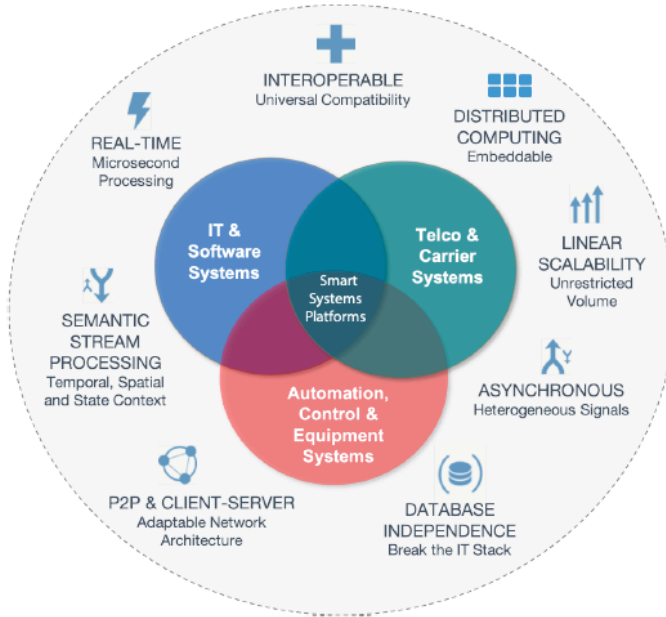
Here is a quick excerpt from the report:

"A handful of startups are accelerating the pace of development, catching larger established rivals off-balance and threatening fundamental shifts." SkyFoundry is recognized as one of these companies. They go on to say - "Their strategies reflect the increasing importance of the convergence of distributed computing with open information architectures; the combination of these capabilities is a fundamental requirement to effectively enable large scale collaboration networks and ecosystems. These players understand enabling completely fluid information, interoperating devices, people and systems requires a simple, flexible, and universal schema that will make information truly portable in both physical and information spaces."

You can download a copy of the full Harbor Research report at this link: <http://harborresearch.com/smart-systems-iot-platforms/>

Try typing the link if a click does not take you to it

A handful of startups are accelerating the pace of development, catching larger established rivals off-balance and threatening fundamental shifts



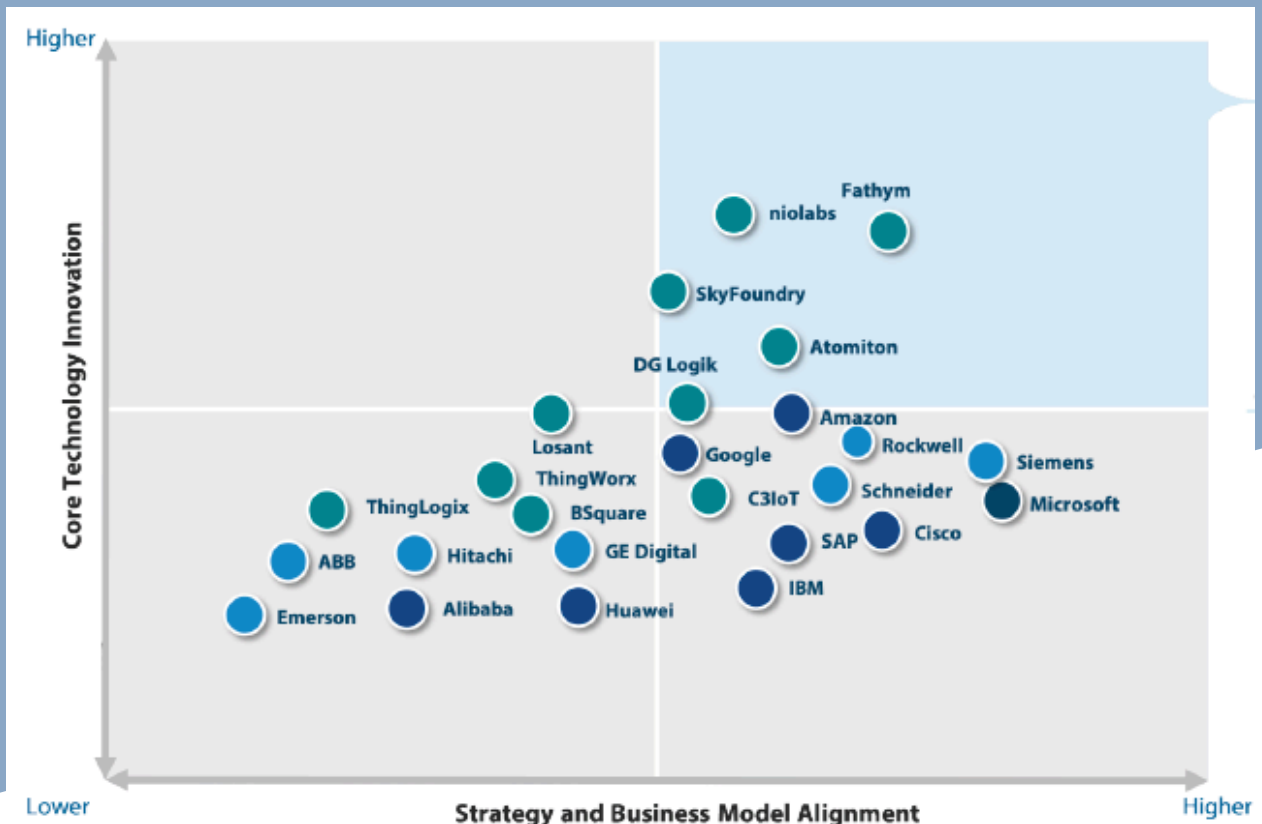
Their strategies reflect the increasing importance of the convergence of distributed computing with open information architectures; the combination of these capabilities is a fundamental requirement to effectively enable large scale collaboration networks and ecosystems

niolabs

SkyFoundry

fathym.

Atomiton



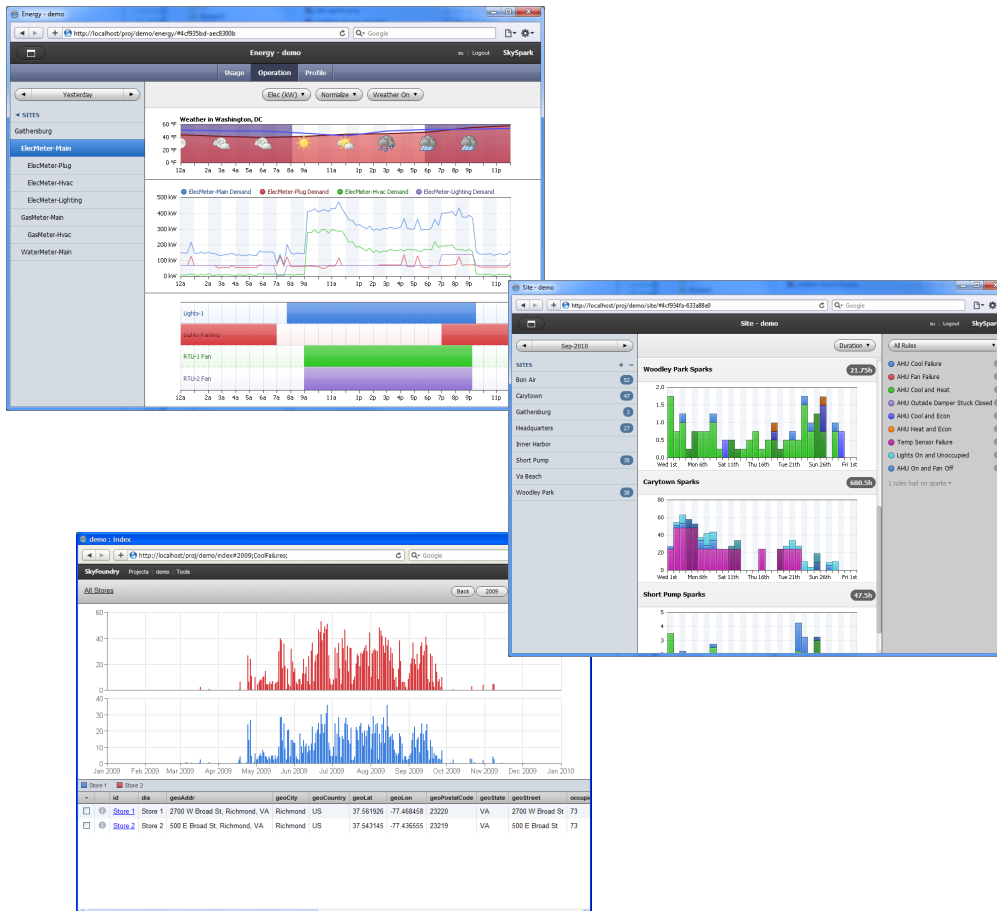
SkyFoundry is 10 Years Old!

January 2019 marks SkyFoundry's 10th birthday.

After years of experience in hardware and software for controls and automation we asked ourselves - what is the next frontier? The answer we settled on was DATA – that the data produced by smart devices and equipment systems was the key to the future of automation, control and the IoT, which was more commonly referred to as M2M in those days. We made a bet that a data-centric software platform with advanced analytics capability would provide benefits across a wide range of applications – and frankly be a lot of fun. That set the course for SkyFoundry.

Ten years on we want to sincerely thank all of our customers and partners that have made this first decade possible, fun, and successful. Here's to 10 more!

...check out some of the screens from SkySpark's 10-year data analytics journey below



SkyFoundry Events
Meetings & Conferences ▾
🛒 🔍

SkyFoundry Events

Where you can meet, share ideas, learn, network and so much more!

USA

SkyPosium

2018

The Worldwide SkyFoundry Community Event

For All SkyFoundry Partners and Users



November 13-14, 2018 | Denver Colorado, USA

EUROPE

SkyPosium - EU

2018

The Worldwide SkyFoundry Community Event

For All SkyFoundry Partners and Users



October 2-3, 2018 | Amsterdam, The Netherlands

SkyPosium 2018 Recap

Our inaugural 2017 “SkyPosium” Community event was a resounding success and this year we brought SkyPosium to Europe as well. Our 2018 SkyPosium user community events were a great success, with combined attendance up 50% from 2017.



SkyFoundry’s SkyPosium events are designed for the entire community of SkySpark users—our reseller partners, end users, engineering consultants, SaaS providers and OEMs — everyone that uses or applies SkySpark. With the majority of sessions delivered by community members, SkyPosium’s unique format provides attendees with the ability to learn best practices and share experiences among actual users and practitioners. SkyPosium will be back for 2019 – watch for announcements in Q1 of 2019.



Haystack Connect 2019

We are excited to again be part of the Haystack Connect 2019 Conference. This year's event will take place at Paradise Point Resort & Spa in San Diego, CA 92109

Haystack Connect 2019 is organized and produced by the Project Haystack Organization—the open source community of people and companies who share the vision that a connected, collaborative community can move the industry forward in ways that no single supplier can. The event builds on the inspiration and mission of the community to address the challenges of making smart device data work seamlessly across applications of all types.

The building automation, energy efficiency, and IoT markets have a variety of events. Some are under the control of a single manufacturer. Some are large conferences where the needs of systems integrators, technology vendors and their customers are lost in a sea of unrelated products and content that are not relevant. Haystack Connect directly serves the needs of the community that is on the leading edge of applying smart data, smart devices, smart equipment and smart building technologies to create a more efficient and sustainable world.

The work of the Project Haystack community addresses a critical gap that has limited the flow of data from smart devices to value-added applications, by making that data self-describing and easy to use, and the biennial Haystack Connect conference is the keystone event for the community. It provides a great opportunity to meet face-to-face with representatives from this unique market and attend an extensive educational program delivered by leaders in the use of device and IoT equipment data. The conference also hosts an Exhibitor Showcase which features the latest and greatest equipment, software and technology in the field. ***SkyFoundry is proud to support Haystack Connect 2019 as a Platinum Sponsor.***

Find full details on haystack Connect 2019, including registration and sponsorship information here: <https://www.haystackconnect.org/>

SkySpark® – Creating Value from Device Data

The past decade has seen dramatic advances in automation systems and smart devices. From IP connected systems to support for web services and xml data schemas, it is now possible to get the data produced by the wide range of systems and devices found in today's buildings and equipment systems.

Access to this data opens up new opportunities for the creation of value-added services to reduce energy consumption and cost and to identify opportunities to enhance overall facility operations.

Access to the data is just the first step in that journey, however. The new challenge is how to manage and derive value from the exploding amount of data available from these smart and connected devices. *SkySpark directly addresses this challenge.*

The new frontier is to efficiently manage and analyze data to *find what matters™*

Proven in Applications of All Types

Applying analytics to the built environment is about more than just energy efficiency and commercial office buildings. SkySpark is used successfully in all types of facilities with deployments across well over **1 Billion square feet** (over 92,903,040 m²) of space on 6 continents. Applications include:

- Commercial office buildings (owner occupied, REITs)
- Utilities (demand response, load management)
- Government and Military facilities
- Data Centers
- Industrial facilities
- Multi-site Retail and Quick Serve Restaurants
- Higher Education
- Indoor Agriculture
- Laboratories (research and universities)
- Entertainment/Hospitality (casinos, shopping centers, hotels)
- Smart Cities
- Facility management service providers

The common theme to all of these applications is DATA. The systems that enable these facilities to support their diverse missions contain vast amounts of data – more data than humans can possibly review and assess. SkySpark analytics is the tool that operators and managers use to quickly *Find What Matters™* in their data, enabling them to optimize efficiency, streamline maintenance, ensure occupant satisfaction and maximize production and profitability.