

# The Future of Smart Systems and IoT Analytics

Harbor Research was recently given the opportunity to examine a new data management and analytics platform that takes a refreshingly new approach to machine data intelligence. SkyFoundry leapfrogs the current market's noise and clutter about big data by viewing core data management and analytics development for smart connected products and Internet of Things systems as a unified challenge that can be addressed by a single, scalable solution from embedded nodes at the edge to the cloud. In so doing, SkyFoundry is re-defining the concept of connected analytics platforms and how value is created from machine data



Smart  
Systems  
Design

Harbor  
Research

The Internet of Things (IoT) and the new world of Smart Systems are ushering in an era where people, machines, devices, sensors, and businesses are all connected and able to interact with one another. The convergence of networked computing and large-scale data management with real time machine intelligence is driving the integration of the physical and virtual worlds and creating unimagined business opportunities. Data management, modeling and analytic tools are the core enablers of these new opportunities. However, numerous hurdles have constrained growth in machine data and IoT analytics; users are just scratching the surface of the real value that the vast stores of collected data from their connected products can unlock. A new generation of machine data intelligence is emerging that takes the value of connected products one significant step further. These tools dramatically increase visibility and understanding of usage and behavior. Perfect information about connected products and machines is the holy grail of the Internet of Things.

---

# TABLE OF CONTENTS

---

01: The Advent of Machine Intelligence and Analytics	Page 04
02: Current Computing Paradigms Only Tell Part of The Story	Page 05
03: The Need For An Architecture That Matches the Real World	Page 07
04: Enter SkyFoundry	Page 07
05: Information Is Not Free	Page 08
06: Cyber-Physical Computing Is Radically Different	Page 09
07: Designing The Future of IoT Data and Analytics	Page 10
08: What's Really Required?	Page 12
09: SkySpark: Next Gen Tools To Make Data Meaningful and Useful	Page 13
10: New Relationships and Ecosystems To Extend Innovation	Page 15
11: New Business Mode Drives New Values	Page 17

# THE ABILITY TO DETECT PATTERNS FROM DEVICE DATA IS THE HOLY GRAIL OF SMART SYSTEMS

Machine communications and the Internet of Things are combining to create new modes of asset awareness, intelligence and support services. In its simplest form, the Internet of Things is a concept in which inputs from machines, sensors, people, video streams, maps and more are digitized and placed onto networks. These inputs are integrated into Smart Systems that connect people, processes, and knowledge to enable collective awareness, efficiencies and better decision making. It is the combination of the data from machines with value added applications that delivers the value of the IoT.

We have now entered the age when everyday objects will communicate with, and control, other objects over networks—24/7/365. The objects are everything from consumer appliances and wearables, IT infrastructure and the elevator you've been waiting for; devices that enable industrial productivity, safe and comfortable indoor environments, consumer entertainment, health monitoring, safety, security and convenience to everyday life.

The convergence of large-scale data management and networked computing with real time machine intelligence is driving the integration of the physical and virtual worlds. The intersection of these trends - the Internet of Things and People - will create unimagined new values. Data management, analytic tools and new skills will be the core enablers of these new values.

Sensors, machines and a wide range of devices generate massive amounts of structured and unstructured data, requiring a whole new class of data modeling, management and analytics tools to uncover and capture value. In the hands of talented users, analysts, domain experts and data scientists, these data can generate productivity improvements, uncover operational risks, signal anomalies, eliminate inefficient service cycles, and even drive enhanced security protocols.

In a truly connected world of smart systems, not only people but all electronic and electro-mechanical products and machines will produce mountains of valuable information, all the time. Consider that today the number of connected devices on the planet has surpassed the number of people - 8+ billion - and depending on your definition of a sensor, there are already many more sensors on earth than people. This will generate phenomenal volumes of data ripe for value creation.

The ability to detect patterns from these devices is the holy grail of smart systems. Machine data analytics, often thought of as part of the evolving "big data" story, allows not only data patterns but a much higher order of intelligence to emerge from large collections of ordinary machine and device data. The implications of mining and analyzing machine data are immense; this is where the real core value creation opportunity lies within the Internet of Things.

## 01

### The Advent of Machine Intelligence and Analytics

#### IoT DATA AT SCALE:

Our research and experience verifies that the scale of data driven by devices will dwarf anything we have previously seen. For example:

- » A single large oil refinery produces more data in a day than all of the New York Stock Exchange and AMEX combined.
- » In a 200 turbine wind farm, each turbine has 50 sensors with over 100 data points collected every 40 milliseconds, producing over 6,000 data points every second.
- » Estimates of data produced by Smart Grid applications could reach between 35 and 1,000 petabytes per year.

---

# THE IT AND TELECOM SECTORS HAVE FAILED TO RE-EVALUATE THEIR RELATIONSHIP TO ADVANCING TECHNOLOGY

Today, significantly better tools are required to organize and integrate the huge volumes of sensor and equipment data for analysis; the conventional tools available for professionals within the IoT or the operational technology (OT) arena still fall far short of real world needs.

Before delving into the new thinking that makes this story possible, let's talk about why it's necessary at all. The IT and telecom sectors have failed to re-evaluate their relationship to advancing technology and to their constituents. The business and technology paradigms to which these industries cling today are far too limiting, too cumbersome and too expensive to foster and sustain new growth. They are like what the health insurance industry likes to call "pre-existing conditions." Something most organizations do not want anything to do with.

From a Telco perspective, today's discussions of IoT systems focus almost exclusively on communications -- the "pipe" -- and very little on the information value. In other words, on things that look good to the carriers. There are many popular visions about wireless monitoring and wireless control. Such as it is, wireless is a fantastic new advance -- no question. But, focusing on the communication element alone as 'first-order' business value amounts to grabbing the wrong end of the technology stick. Wireless communications alone steal the limelight from the real revolution -- utilizing new technologies and processes to liberate information from sensors and intelligent devices to leverage collective awareness and intelligence.

From an IT perspective, today's corporate IT function is a direct descendant of the company mainframe, and works on the same "batched computing" model—an archival model, yielding a historian's perspective. Information about events is collected, stored, queried, analyzed, and reported upon. But all after the fact.

That's a very different thing from feeding the real-time inputs of billions of tiny "state machines" into systems that continually compare machine-states to rules and algorithms and then do something on that basis. In short, for connected devices to be used to their full potential in business, the prevailing corporate IT model has to change.

In the course of the last two decades, the world has become so dependent upon the existing ways computing is organized that most people, inside IT and out, cannot bring themselves to think about it with any critical detachment.

The client-server model underlying today's computing systems greatly compounds the problem. Regardless of data-structure, information in today's computing systems

---

## 02

### Current Computing Paradigms Only Tell Part of The Story

---

## TODAY'S COMPUTING SYSTEMS WERE NOT DESIGNED FOR A WORLD DRIVEN BY PERVASIVE INFORMATION

is a huge collection of information-islands whether on your servers, your service provider's servers or anywhere else. Assuming the islands remain in existence reliably, they are still fundamentally incapable of truly interoperating with other information-islands. We can create bridges between them, but islands they remain. That's what they were designed to be.

IT professionals talk these days about the need for ever-evolving information services that can be made available anywhere, anytime, for any kind of information, yet IT rarely has the experience and background to architect and integrate real world physical systems.

With each additional layer of engineering and administration based on historical IT design principles, computing systems come closer and closer to resembling a fantastically jury-rigged Rube Goldberg contraption with development and maintenance costs that inhibit the growth of IoT applications. The reason is simple. Today's computing systems were not really designed for a world driven by pervasive and diverse information flow and interaction.

The next cycle of technology and systems development in the smart connected systems arena is supposed to be setting the stage for a multi-year wave of growth based on the convergence of innovations in software architectures; back-room data center operations; wireless and broadband communications; and smaller, more powerful client devices connected to personal, local and wide-area networks. But is it?

Artificial intelligence, machine learning and the Internet of Things are all in some way trying to break from today's computing paradigms to enable real-world physical systems, and the data they contain, to become self adaptive, intelligent information systems. But as we look around at even the latest advances, we still find an architecture with its roots on the client server model of the past. Even as the "cloud" becomes the server, a key element of the IoT is being ignored - an architecture that matches the structure of the physical world.

At the same time that we have seen dramatic proliferation of smart devices - devices that operate at the "edge" - in remote sites, buildings, equipment rooms and telecom closets, vehicles and people's pockets, we have been lead to believe that the only way to get value from that data is to bring it all to the cloud.

## THE IoT IS A DISTRIBUTED COMPUTING CHALLENGE

The IoT is a Distributed Computing Challenge. IoT deployments are by their very nature distributed systems. We hear lots of talk about the “cloud” as it relates to the IoT to the point that in many cases it seems like the “cloud” is being presented as the solution to all things IoT. 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 recognize and embrace the highly distributed and innately non-hierarchical nature of the IoT and support that with a corresponding software architecture.

The IoT is not and will not be a well organized, centrally planned development. It will grow in ways that may appear unplanned and disorganized when viewed from a classic central computing vantage point – it will more resemble an organic system. Value added applications will spring to life in environments where the right resources (data) find the right nutrients (software applications that can deliver financial value). These applications will occur at all levels of the architecture, from the edge to the cloud. Acceptance of this reality is essential to the effective design of a data management and analytics platform for the IoT.

It has become evident that the larger software, data and BI players are not likely to become the defacto leaders in Smart Systems and IoT analytics. While the “traditional,” players may have what appears to be a head start they are just not accustomed to the unique characteristics and requirements of physical sensor and machine-based applications within the IoT arena. A growing number of relatively young companies are introducing products that have been purposely designed for IoT analytics. This paper is about an important new IoT data analytics platform and application offering from SkyFoundry.

The SkyFoundry team of innovators understand that conventional IT and business intelligence tools we are working with today to discover, manage and analyze machine data were not designed to really address operational technology (OT as opposed to IT) applications and challenges. SkyFoundry’s SkySpark platform has been designed to provide an end-to-end software architecture that matches the often unruly IoT world and efficiently provides business insight for organizations by aggregating, mining and analyzing data generated from sensors, equipment, machines and products, wherever those devices and their data may be located.

SkyFoundry’s platform is not a simple incremental improvement, or new flavor of what we already do. Their development represents a true shift in thinking about how devices, data, people and physical systems will be integrated and how they will

### 03

#### The Need For An Architecture That Matches The Real World

### 04

#### Enter SkyFoundry

---

## TRADITIONAL APPROACHES TO DATA DISCOVERY AND SYSTEMS INTELLIGENCE HAVE MANY FAILINGS

interact in massive peer to peer distributed systems. SkyFoundry's approach is not about leveraging aging IT technology into a new application context; it's about looking forward to a single, unified platform for device data interactions which liberates information and value creation by abandoning traditional relational database and client-server computing models. SkyFoundry's approach provides a single unified platform for data aggregation, management, search, discovery, analysis and prediction across diverse machine data types. The result is a truly new approach to data analytics that provides a complete picture of the myriad of interactions and states that machines evolve through including status, configuration changes and usage.

In today's world, information is not free (that's free as in "freedom," not free as in "free of charge"). In fact, thanks to present information architectures, it cannot access and merge device data with other information and cost effectively enable any kind of discovery or intelligence. What would truly liberated information be like? It might help to think of the atoms and molecules of the physical world. They have distinct identities, of course, but they are also capable of bonding with other atoms and molecules to create entirely different kinds of matter with unique and valuable properties. Often this bonding requires special circumstances, such as extreme heat or pressure, but not always.

In the world of information, such bonding is not all that easy. Today's software platforms focus on execution processes that generate one of three types of data - unstructured, transactional or time series. For each of these data types, a specific set of intelligence tools have evolved to provide "insight" but, in most cases, these tools limit the questions that can be answered to those known in advance. So for a user attempting to do something as simple as asking a multi-dimensional question that might involve such variables such as environmental conditions, energy use and equipment operation, creating new information from multiple data types that is an easily perceivable, manipulable, or map-able "model" of the answer to that question is a significant challenge.

The traditional approaches to data discovery and systems intelligence have two failings: they can't provide a holistic view of these diverse data types and, the types of intelligence tools available to users are, at best, arcane and typically not accessible to the people that have the application domain knowledge. The ability to detect patterns in data is the holy grail of smart systems and The Internet of Things because it allows a whole higher order of intelligence to emerge from large collections of ordinary data. The implications are obviously immense, but given the immature state of today's real-world systems, most people have trouble grasping the power and importance these capabilities can enable.

---

**05****Information Is  
Not Free**



## TODAY'S PLAYERS SHOW LITTLE APPETITE FOR RADICAL DEPARTURES FROM CURRENT PRACTICE

06

Cyber-Physical  
Computing Is  
Radically  
Different

When it comes to preparing for the global information economy, most people assume that “the IT and telco technologists are taking care of it.” They take it on faith that the best possible designs for the future of connected things, people, systems and information will emerge from large existing IT and software players. But those are big, unfounded assumptions. In fact, most of today’s entrenched players are showing little appetite for radical departures from current practice. Yet current practice will not serve the needs of a genuinely connected world.

What are the major obstacles that need to be overcome?

**Leveraging collective intelligence:** For all its sophistication, many of today’s connected product systems are a direct descendant of the traditional cellular telephony model or the remote services/M2M model where each device acts in a “hub and spoke” mode. The inability of these systems to interoperate and perform well with distributed heterogeneous devices and new, independently developed software applications is a significant obstacle. The many “nodes” of a network may not be very “smart” in themselves, but if they are networked in a way that allows them to connect effortlessly and make their data seamlessly available to each other and independently developed software applications, they begin to give rise to complex, system-wide behavior. This allows an entirely new order of intelligence to emerge from the system as a whole—an intelligence that could not have been predicted by looking at any of the nodes individually or their original functions alone. What’s required is to shift the focus from simple device monitoring to a model where device data is communicated easily to new data management tools and analytics applications to achieve true systems intelligence.

**Optimizing all assets - tangible and intangible:** New software technologies and applications need to help organizations address the key challenge of optimizing the value of their balance sheets, allowing them to move beyond just financial assets and liabilities to their physical assets and liabilities (like electric grids or hospitals) and then to their intangible assets and liabilities (like a skilled workforce). The task of optimizing the value of financial assets, physical assets and people assets requires new technologies that will integrate diverse information sources in unprecedented ways to solve more complex business problems.

**Flexible, scalable systems:** IT professionals rarely talk these days about the need for ever-evolving information services that can be made available anywhere, anytime, for any kind of information. Instead, they talk about web services, enterprise apps and now cloud computing. The Web stores information in one of two basic ways: utterly unstructured, or far too rigidly structured. The unstructured way gives us typical static Web pages, blog postings, etc., in which the basic unit of information

---

## THE IoT REALLY MEANS THE CONVERGENCE OF INFORMATION, CONTROL AND PHYSICAL SYSTEMS

is large, free-form, and lacking any fundamental identity. The overly structured way involves the use of relational database tables that impose rigid, pre-ordained schemas on stored information. These schemas, designed by database administrators in advance, are not at all agile or easily extensible. Making even trivial changes to these schemas is a cumbersome, expensive process that affects all the data inside them. Just as importantly, they make deep, inflexible assumptions about the meaning and context of the data they store. Both of these approaches to data-structure enforce severe limitations on the functions you want most in a global, pervasive-era information system: scalability, interoperability and seamless integration of real-time or event-driven data. The client-server model underlying the Web greatly compounds the problem.

**Automated development:** 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. While the devices may be ‘standard’ each new IoT implementation to connect them and work with their data requires too much customization and maintenance just to perform the same basic tasks.

Some things that should be kept simple are allowed to get unnecessarily complex. The drive to develop technology can inspire grandiose visions that make simple thinking seem somehow embarrassing or not worthwhile. But making device data work for us simply is the essential goal we need to achieve. This requires a re-think of the solution stack, not simply the introduction of additional layers.

The Internet of Things really means the convergence of information, control and the cyber elements of physical systems. It will require a remarkably agile network that could comfortably scale to billions or more nodes—some of them hardware, some software, some purely data, many of them coming into and out of existence or some in obscure physical locations or changing location constantly. Obviously, such a network cannot be “designed” in any ordinary sense. Certainly, it cannot be designed “top-down.”

And yet the Internet of Things must be designed in some sense. Some basic design principles must be put in place to guide the growth of a vast, distributed technological organism that must remain organized as it evolves according to a logic all its own. It demands that we design not only devices and networks but also information itself in ways not addressed by current IT.

---

## 07

### Designing The Future of IoT Data and Analytics

## THE TOOLS WE ARE WORKING WITH TODAY TO MAKE PRODUCTS “SMART” ON NETWORKS FALL SHORT

Artificial intelligence, machine learning and the Internet of Things are all in some way trying to break from today’s computing paradigms to enable real-world physical systems, but these promising technologies need to be viewed as applications that will work with this next generation IoT architecture that provides near seamless access to data of all types seamlessly to these applications. Said another way, it will not be feasible to re-invent the underlying infrastructure for each new application. Instead, these applications need to integrate easily with an appropriate networking and data management foundation.

We have seen that we can make a computer capable of beating the reigning genius of chess, but yet we can’t make a robot capable of walking across the street as well as any normal two-year-old child. The real world is not a strictly regulated, closed system like a chess game. Sensing a player’s moves on a wired chessboard and responding quickly and intelligently based on “knowable” algorithms is one thing. Sensing and responding to physical systems and states – is a fundamentally different challenge. To achieve the seamless intelligent solutions envisioned by the IoT requires an entirely new approach that can leverage common platforms and data models across diverse devices, data and domains.

### Next-Gen Architectural Approaches Required to Enable Real-World Physical Systems



## CUSTOMERS EXPECT EVOLVING SOFTWARE TOOLS TO BE FUNCTIONAL, UBIQUITOUS, EASY-TO-USE

The fact that a wide range of sensors, machines and equipment can transmit information about status, performance and usage, and can interact with people and other devices anywhere in real time points to the increasingly complex role of data in IoT systems. This only compounds when we consider the billions or more of networked devices that many observers are forecasting will be deployed and the scale of data they will produce.

The tools we are working with today to make products “smart” on networks were not designed to handle the scope of new functional capabilities, the diversity of devices and the massive volume of data-points generated from device interactions.

These challenges are diluting the ability of organizations to efficiently and effectively leverage the real value of connectivity. The rigid and fragmented nature of software offerings available today make it extremely difficult and expensive to use and re-use diverse device-generated data.

Customers expect evolving software tools to be functional, ubiquitous, easy-to-use and able to be added to their smart device systems even as they come from different vendors. What would be entailed to achieve these requirements?

- » An architecture that can align with and “form fit” to the physical world and its embedded and distributed computing schemes with software tools that are easy to use and address a broad range of real-time, and historical data analytics requirements – ideally a single unified framework to design and build solutions that can interoperate across diverse data environments and under widely differing usage scenarios.
- » A true distributed software architecture that can process and create value from device data locally at the edge - in remote sites, buildings, equipment rooms and telecom closets, while enabling higher level applications that generate value from enterprise or portfolio-wide data. We have been lead to believe that the only way to get value from IoT data is to bring all of it to the cloud. As stated earlier, it is not possible, cost effective or desirable to transmit every piece of data from IoT devices to the cloud in order to gain value from them. Numerous critical applications will require processing and analysis of data locally to support user needs at the edge and to automatically filter and semantically enrich that data before delivering it to higher levels of the architecture for other application needs. To enable the promise of the IoT we need to forget about pushing gargantuan amounts of data to warehouses and repositories and the only architectural approach.

---

**08**
**What's Really  
Required?**

## DATA MANAGEMENT, MODELING AND ANALYTIC TOOLS ARE THE CORE ENABLERS OF IoT VALUES

- » Software that can enable easy integration of any combination of inputs and data types - message - feed - stream - data - in real-time with stream processing that provides complete independence from traditional database technologies and rigid schemas; software designed without the bias and dependencies of a single product or service application that will cause integration pain when attempted to be used beyond its original scope.
- » Truly scaleable software architecture, data models, tools and functions with no constraints on where it gets deployed – whether on a chip, device, server, cloud, or hybrid system; yet able to provide extensibility with common features and functions at each level of the architecture.
- » Software and tools that allow users to easily and quickly build their own applications and automate the generation of visualizations that present critical insight and value to the user.

We are reaching a critical juncture in market development where organizations will soon be crying out for a completely new approach - one where the effort invested to develop new data-centric IoT applications can be quickly and easily re-used again and again across an ever broader spectrum of devices, integrations and end use applications.

The bit, the byte, and later the packet made possible the entire enterprise of digital computing and global networking. Until the world agreed upon these basic concepts, it was not possible to move forward. The next great step in the convergence of IT and OT—completely fluid information and fully interoperating devices, data, people and systems—requires an equally simple, flexible, and universal abstraction that will make information itself truly portable in both physical and information space, and among any conceivable information devices.

SkyFoundry's SkySpark 3.0 platform takes on the toughest challenges of interoperability, database dependency, semantic modeling, and information delivery, and enables adaptable real-time intelligence to empower users and developers to exploit the vast potential of the Internet of Things. SkySpark achieves these goals via a range of key technologies:

- » Give meaning and context to time series data streams by using “tags” Instead of relying on conventional relational database schemas to convey definitions and associations. Tags provide the hooks that the analytics engine uses to correlate and analyze the data that informs temporal, spatial and state-based contextual processing.

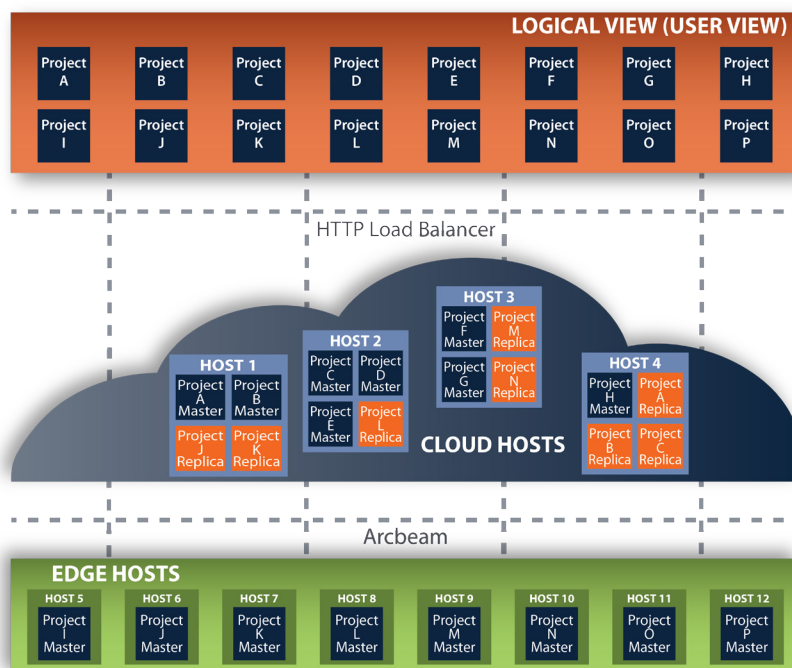
# 09

**SkySpark:  
Next Generation  
Tools To Make Data  
Meaningful and  
Truly Useful**

## SKYFOUNDRY'S PLATFORM ENABLES TRUE REAL TIME INTELLIGENCE

- » Blend and manage diverse datasets – any type of information from any device, whether historical batch data or streaming source. SkySpark's Folio database technology combines a number of non-relational techniques in order to accept, store, normalize and analyze diverse data sets. Folio combines the techniques of document orientated, graph oriented and time series database design to provide a solution suited to performing analytics on diverse sets of sensor and machine data.
- » Integrate multiple data acquisition protocols, to enable simultaneous management and analysis of data from a wide variety of sources via a diverse assortment of communication connectors.
- » Distributed architecture for edge processing that embraces the non-hierarchical nature of the IoT. Designed to run in everything from very small footprint embedded devices to server clusters in the cloud with a single scalable software architecture that brings the analytics processing “to the data” enabling high speed continuous analysis of data.
- » Integral services that provide automatic generation of information visualizations dramatically reducing the application development costs typically associated with delivering data and insights to the user.

### SkySpark 3.0 Platform Logical Architecture



## THE SKYSPARK PLATFORM FUNDAMENTALLY CHANGES CONVENTIONAL PARADIGMS

This last point a key and often overlooked challenge in IoT applications. Finding patterns and transforming data is the core value of an analytics platform, yet if the process required to create presentations of those results to the user entails significant costs, complexity, long lead times and lack of re-use, the potential benefits will not be realized. To make data analytic results useful for operators, analytic findings need to be automatically converted into visualizations that are easy for operators to understand and act on. Automatic conversion of analytic results into rich, intuitive visualizations that show trends, correlations and data relationships without requiring custom development is a significant advance in the state of the art.

Recognizing the needs for highly specialized information presentations cannot be overlooked, however, so SkySpark combines the power of automatically generated visualizations with the ability to combine view elements into custom applications a concept SkyFoundry calls compose-ability. This is not simply assembling information and charts into reports, but rather combining different visualizations elements (or widgets) into unique visualizations that can be viewed as a interactive application on a PC or mobile device, or delivered as a static report.

- » Combining pre-packaged algorithms for machine data analysis along with the ability for domain experts to implement their own algorithms to meet unique application needs. Lab-developed algorithms are useful, but the real world rarely matches the lab.

Data management, modeling and analytic tools are the core enablers of new Internet of Things values. However, numerous hurdles have constrained growth in machine data analytics. We believe that overcoming these hurdles will drive a focus on developers, ecosystems and the synergy of relationships required to drive value. One of the major hurdles is data semantics – being able to convey not just data values but the meaning of that data.

SkyFoundry's founders are a part of an open source community that is working on a unique semantic modeling approach to address the challenge of making device data self-describing. This approach allows users to create data representations of a physical devices and systems elements that can be shared along with the data produced by the devices. Project Haystack is an open source initiative to develop tag naming conventions and taxonomies for modeling of equipment systems and their operational data – consider it a “mark-up language” for data.

---

# 10

**New Relationships  
and Ecosystems To  
Extend Innovation**



## DEVELOPING ANALYTICS APPLICATIONS THAT INTEGRATE DIVERSE DATA SOURCES AND FOCUS ON END USER VALUE WILL SEE BIG ADOPTION

In much the same way the HTML allows different browsers to accept and properly display a web document based on a known set of “markup rules,” Project Haystack is developing a standardized methodology and consensus-approved data models and tag libraries for diverse devices such as a temperature sensor or an electric meter. The essence of an “device” is completely abstracted from its real-world embodiment and is mutually interchangeable. The structure Project Haystack is utilizing is based on a simple model that includes entities such as:

- » **Site:** single building with its own street address
- » **Equip:** physical or logical piece of equipment within a site
- » **Point:** sensor, actuator or setpoint value for an equip (the final element that data is associated with)

Each of these entities can then be tagged to capture essential information about its characteristics, purpose and relationships to other entities. It’s important to note that Haystack is cleverly combining a data modeling methodology with a community-developed tagging libraries (taxonomies) and the required communication protocols and APIs to exchange Haystack modeled data between applications and software tools, without them needing to be pre-aware of each other. This is a model we believe will be applicable in other industry sectors such as industrial systems. The open source nature of the project allows experts from any domain to build on the foundational work of the community to develop tagging models for systems in their own sphere.

SkyFoundry is leveraging this new semantic modeling approach to avoid the confinements and limitations of the today’s differing data types and relational database-oriented tools. It allows data to maintain their fundamental identity while bonding freely with other data. Facilitating discovery, based on data and information accessibility and cumulative systems intelligence, is one of the fundamental purposes of SkyFoundry’s platform. They are designing a system for a genuinely connected world in which there are no artificial barriers between pieces of information.

The SkySpark platform fundamentally changes the conventional paradigm, treating data from sensors, devices, systems and the physical world as “neutral” representations. In other words, treating diverse data types equally. This enables processes connecting diverse data in any combination to be rapidly built and deployed.

### PROJECT HAYSTACK

Today’s automation systems, sensors, equipment, metering systems, and smart IoT devices produce tremendous amounts of data that can be very hard to organize and use across different applications because it is stored in many different formats, has inconsistent naming conventions and very limited data descriptors.

In essence it lacks information to describe the meaning of the data and without meaning a time-consuming manual effort is required to utilize the data for discovery and analytics.

Taking cues from other data intensive applications such as Facebook, Twitter, Google and others, the open source Project Haystack community addresses this challenge by defining a methodology to describe the meaning of data from smart devices with tags and standardized semantic data models.

This enables software applications to automatically consume, analyze and present data from devices and equipment systems. The Project Haystack community has streamlined the interchange of data and the techniques for managing the vast amount of data generated by today’s smart devices.



## THE SKYFOUNDRY TEAM UNDERSTANDS THE CRITICAL NEED TO PROVIDE A FULL END-TO-END DATA MANAGEMENT AND ANALYTICS SOLUTION

11

**New Business  
Mode Drives  
New Values**

To date, remote services and M2M systems have largely been focused on single purpose, vendor-centric services - in large part because of technical complexities, and business model challenges. Existing technology has proven cumbersome and costly to apply across conflicting protocols, different data types and incomplete component-based solutions. The challenges of developing data-centric analytics applications that integrate diverse device data sources and focus on end user value have seen big adoption hurdles.

Simple applications largely focus on the product manufacturer's own service delivery chain with simple "hub and spoke" remote support. While there is value in these models, there are significant untapped opportunities for providing new value for the users and customers. Return from simple applications, while valuable, is limited to the manufacturer's service delivery efficiency. Contrary to what current market offerings depict, however, the value of connectivity does not have to end with just simple applications focused on a single class of device or machine.

As technologies mature, applications based on deeper, peer-to-peer interactions between devices, systems and people will drive more compound and dynamic value streams which will in turn open up new collaborative business model opportunities that have the potential to drive much greater value for the all members of the value chain - from customers to service providers and equipment manufacturers.

Moving from "Simple" to "Compound" applications involves multiple collaborating systems with significant interactions between and among devices, systems and people. No longer is the focus solely on the product supplier's ability to deliver support for their product efficiently. Rather, value is brought to the customer through business process automation and optimization.

The nature of compound and complex systems applications is just beginning to be understood. What the SkyFoundry team have come to understand is the critical need to provide a full end-to-end data management and analytics solution that is consistent and uniform at all levels of the architecture.

Combined with alliances and a new breed of IoT analytics solution partners and developers, SkyFoundry enables partners to add unobtrusive machine data analytics to a broad range of equipment and systems. It treats user concerns—from product usage, performance, and efficiency to health and support—as a unified challenge that can be addressed by a new generation of data analytics provided to customers under a range of service models. In taking this perspective, we believe SkyFoundry is pioneering a new meta-market opportunity with vast potential.

### ABOUT HARBOR RESEARCH

Founded in 1984, Harbor Research Inc. has more than thirty 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.