

# SkyFoundry Insider

The Data Center Issue - Improving Performance and Reliability in Data Center Applications Pg 1

Using SkySpark to Support a Fully Integrated Data Center Pg 2

SkySpark Analytics Helps the Data Center at Dolmabahçe Palace Stay Cool and Efficient Pg 4

Analytics – a Key Component of a Comprehensive Information System for Data Center Operators Pg 6

Connecting to Diverse Data Sources: SkySpark Adds OPC UA Connector Pg 7

Project Haystack Announces Two New Publications Pg 8

Update on SkySpark 3.0 – New Features Continue to be Released Pg 9

## The Data Center Issue *Using SkySpark® to Improve Performance and Reliability in Mission Critical Data Center Applications*

With application across over 10,000 buildings, SkySpark has been successfully applied in a wide range of applications in buildings of all types. In this issue we focus on the use of SkySpark in a very special market segment – Data Centers. Data Centers have unique operating needs, equipment systems, and of course, reliability requirements.

In this Insider, SkyFoundry partners provide a glimpse of how they use SkySpark to help data center operators understand and improve operations, reduce costs, and insure reliability.

The focus on data centers highlights another key point about SkySpark – its applicability to facilities and equipment systems of all types. SkySpark is used in colleges and universities, multi-site retail stores, the largest commercial office buildings, small footprint fast food restaurants, government facilities, military bases, hospitals, hotels, and industrial facilities. How can a software application fit such a wide range of applications? SkySpark has been designed to work with all types of data – real time and historical, as well as asset data and normalize the data in a way that makes it possible to perform analytics across these diverse data sets. Equally important to the success of SkySpark across diverse applications is the knowledge and service delivery capability of SkyFoundry's worldwide partner network. These organizations bring their specialized domain knowledge to applications across all types of facilities and applications. And by being fully programmable, SkySpark enables these experts to implement the analytic rules that fit the unique needs of the application.

## In This Issue

**SkySpark as Part of a Fully Integrated Data Center**



**Analytics Help the Data Center at Dolmabahçe Palace**



Pg 4

**SkySpark Integrates with DCIM for Optimized Operation**

Pg 6



**Announcing Haystack  
“Connections” Magazine and  
CABA White Paper**

Pg 7

# Data Centers: Demonstrating the Power of SkySpark Analytics to Improve Performance

## Using SkySpark as Part of a Fully Integrated Data Center Backbone

Controlco began its first implementation at a major data center, located in the Pacific Northwest four years ago. Recently, the integration and building automation services company was brought on to complete a second implementation at the facility, this time for one specific data center client. At the outset of this client's 8-megawatt project, which will go live later this year, Controlco was tasked with integrating the company's building management system into the IP backbone and providing a comprehensive operator information interface, while communicating with the data centers BMS systems.

Data centers require very controlled environments, and each client's data processing needs require unique sets of data collection and monitoring rules. For this particular client, their data center focus has been on transitioning from retail co-location capacity to wholesale, and utilizing up-and-coming management methodologies such as 400-volt distribution and cabinet level heat rejection. The company is also doubling cabinet densities from past data centers.

The client will use three rooms of the total 240,000 sq. ft. facility. In those rooms, Controlco has networked 352 pieces of equipment with 12,800 single project points. Each point generates data. Controlco is utilizing both Niagara and SkySpark for data collection and analytics, along with a client-facing DG Lux interface, creating a unique combination of information analysis, presentation and alarm management capabilities that reduces the need for the client to learn, manage and maintain multiple systems.

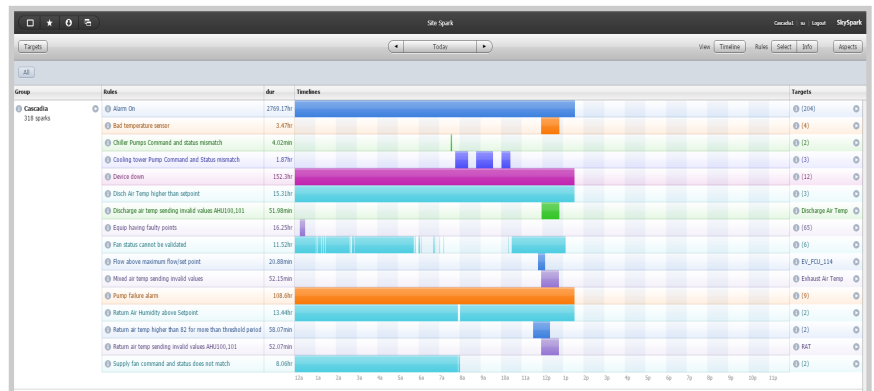


# Using SkySpark to Support a Fully Integrated Data Center Backbone

Con't from pg 2

To determine setpoints in air handling units (AHUs), for example, engineers have identified a target zone temperature for all points on that unit. SkySpark gathers the data values for those points and compares the data to the setpoint temperature throughout specific threshold periods of time. A Spark is generated if the zone temperature noted in the data is above or below the desired target set point and not adjusted by the AHU within the time threshold. This ability to monitor equipment function over time allows the project manager to know exactly where to look when a problem is identified, rather than simply having to manually check all AHUs in this example, or having to start by looking at a whole room full of multiple pieces of equipment.

Controlco engineers developed additional layers of analytics and alerting to customize the check times from up to one day to every 10 minutes, giving the client further knowledge and control of their sensitive data center environs. Controlco has also associated costs factors with Sparks, which is helpful for a number of applications.



The processing power of this combined solution allows for more data to be analyzed so that analysis can be more accurate. After the project goes online, the client can begin viewing Spark patterns, historic data charts, measure averages per year, and calculate actual cost savings against key performance indicators (KPIs).

## Controlco

Our thanks to Controlco for this case study:

<http://www.controlco.com/>

Among the goals for the client in moving its data center to the Oregon facility was to increase energy efficiencies by having greater access to renewable sources. Since many inefficiencies can also be spotted through automated data analytics they can be addressed more quickly through this enhanced data-driven solution, Controlco, Niagara and SkySpark have offered another way for the client to achieve its energy efficiency goals.

## SkySpark® Analytics: Combining a Comprehensive Library of Pre-defined Analytic Functions with Full Programmability

All buildings are different. With SkySpark you get the benefits of full programmability to define rules that fit the specific needs of your facility – your HVAC system design, your building usage and the scope of your project. But SkySpark offers more than just programming flexibility. It combines full programmability with an extensive library of standard analytic functions to streamline the process of turning your domain knowledge into SkySpark rules. The SkySpark library includes over 500 built-in analytic functions and we are always adding more. Many can be considered ready to go “end use” rules, and all can be combined with your own logic to implement rules that fit your application.

**With SkySpark its never one size fits all. Learn more about the SkySpark approach here:**

<http://www.skyfoundry.com/file/48/SkySpark---Combining-Full-Programmability-and-Built-in-Analytic-Functions.pdf>

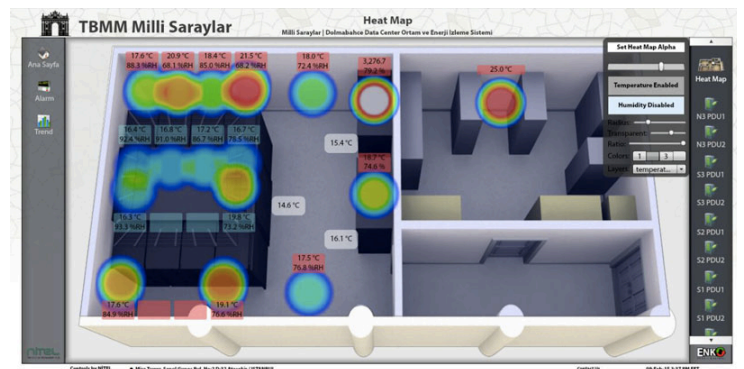
# DCIM Analytics Help the Data Center at Dolmabahçe Palace Stay Cool & Efficient



The Dolmabahçe Palace in Istanbul, Turkey, was built in the mid-19th Century and served as a main administrative center of the Ottoman Empire for many years. It is one of the most ornate and historically significant buildings in the region, and today houses a museum and a data center dedicated to tracking the many valuable artifacts found in the palace and its environs. Each historical element is barcoded and related information is stored within these critical servers. A new Data Center Infrastructure Management (DCIM) dashboard application was deployed for the data center by BASSG's partner in Turkey, En-Ko.

DCIM typically requires a multi-protocol gateway that can aggregate data from BACnet, Modbus and SNMP sources. A third-party device was used to integrate these data streams, and the Project-Haystack nHaystack Niagara module, was used to enable the addition of semantic tagging of the data. The data was then transmitted via the Haystack RESTful interface to SkySpark®. nHaystack unifies real-time and historical data structures in Niagara, reducing the number of steps needed to bring together real time and historical data.

This application combined SkySpark's high-speed historian with BASSG's Energy DVR™ application (image on right) to playback historical data on the visualization platform. This is an example of how SkySpark's support for open-source Project-Haystack connectivity greatly simplifies the integration process.



## CRAC Unit Performance Overview

Live and historical data from Computer-Room Air-Conditioning (CRAC) units are overlaid in this custom-designed home-screen so that data center staff can identify any problems at a glance. Data center operators rely on the help of the SkySpark analytics engine to find optimization strategies. Ventilation optimizations have already resulted in 15% run-time improvements. Here is an example of detecting a faulted Crac unit return air temperature sensor →





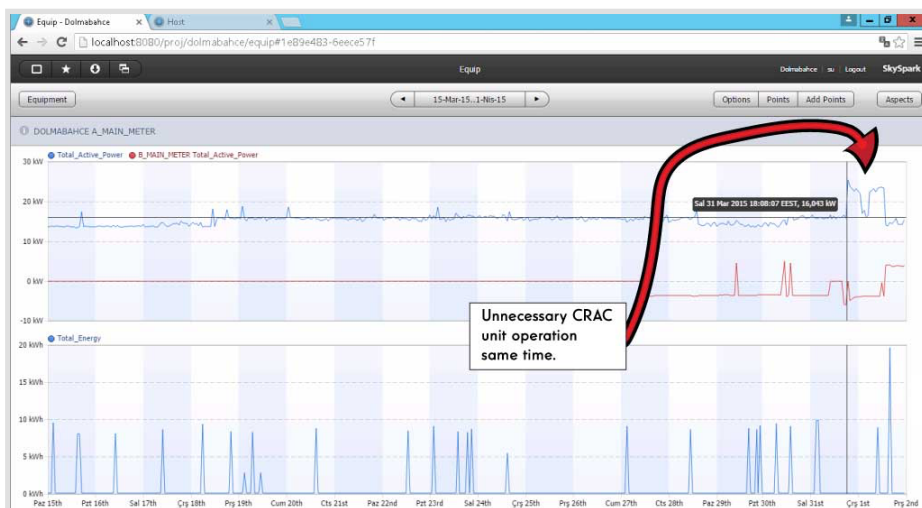
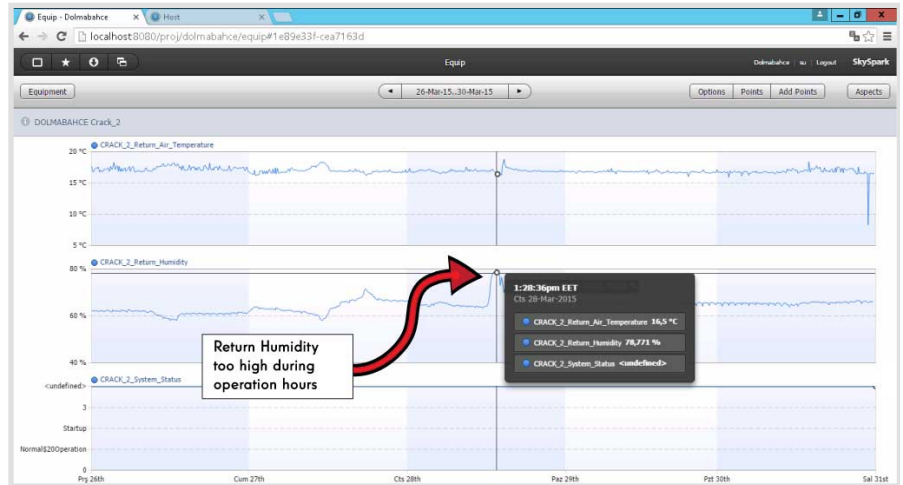
# Analytics at the Dolmabahçe Palace Data Center

(con't from page 4)

## CRAC Unit Cold/Hot-Aisle Reports

The power distribution unit holds a vast amount of data. SkySpark helped operators track peaks in energy usage. Any under utilized racks were reported to the owner.

With the high-speed historian, customers can playback temperature data at five-minute-interval resolution. Sparks and reports are generated for customers automatically. With SkySpark's analytics engine, optimization strategies are easy to identify and data center facilities staff are empowered to take faster action to avoid potential over-heating that would lead to server failures and downtime.



← Here SkySpark shows that both Crac units started at the same time for no reason. The control system sequence was reviewed and modified to correct both issues.

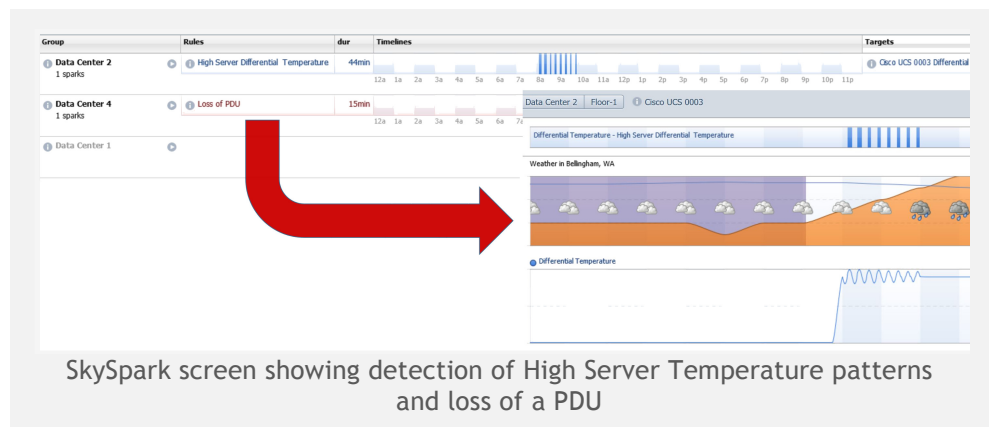


Our thanks to BASSG for this case study. You can find them at: <http://bassg.com/>

# SkySpark Analytics - a Key Component of a Comprehensive Information System for Data Center Operators

Data centers are full of, well, data. Beyond what's stored on the facility's servers from outside sources, data centers also produce their own data about the operation of servers, HVAC systems, and power systems that make up the backbone of the building itself. While buildings of all types benefit from a "single drivers seat" to enable monitoring and management of these disparate systems, data centers present an opportunity to take that approach a step further to solve long-standing issues unique to these types of facilities. At the forefront of this approach is an increased focus on analytics made possible when all systems communicate properly.

With its ability to communicate over Modbus, BACnet and SNMP, and support for other standard equipment protocols, SkySpark helps bridge the gap between traditionally segregated "silos" of systems found in data center facilities enabling all essential equipment systems data to be combined for analysis. In this project, SkySpark was combined with the data center product EntroCIM which provides the ability to communicate with the Intel chipsets inside the servers located throughout the data center, providing access to essential server and processor data that is not otherwise readily available. By integrating EntroCIM with SkySpark analytics operators are given a comprehensive understanding of the actual performance of their critical systems and insight into performance patterns, optimization strategies, failure forecasting, and other information crucial for day-to-day decision making.



SkySpark screen showing detection of High Server Temperature patterns and loss of a PDU

Another challenge for data center operators is "what-if"

analysis. If an operator has maintenance due and needs to take down a PDU (Power Distribution Unit) or router, how will the system respond? To solve this problem, an operator can enter a testing scenario where supporting components can be "failed" to determine how the remaining systems would respond. For instance, if a UPS battery is scheduled for replacement, the system can verify that the opposite bus is correctly setup to power the affected servers and infrastructure. SkySpark then analyzes the data generated during the test to identify and present results. This unique combination of technologies -- SkySpark analytics and EntroCIM, addresses the specialized challenges data center projects present, allowing for the integration of different system "silos" while providing operators with a clear focus on the metrics essential to Data Center Infrastructure Management systems. The addition of SkySpark analytics drives strategies and forecasting that allow data center operators to significantly improve the performance and reliability of their systems.



Our thanks to Hepta Systems for this case study. You can find them here: <http://www.heptasystems.com/>

# Connecting to Diverse Data Sources: SkySpark Adds OPC UA Connector

The Internet of Things has clearly become the buzzword of 2015. Hidden behind this term is the reality that there are literally thousands of different types of devices and dozens of protocols involved in the IoT field. And as much as we might wish for a narrowing in the number of protocols used by smart devices, its fair to say the IoT space has likely seen an increase in the last few years.

These protocols exist for many reasons: legacy systems that are used by large numbers of customers and equipment suppliers; the need to support communications on ultra low power devices (wireless for example); and new protocols that provide more modern streamlined communications in the latest generation of devices. The reality is that no one protocol can serve all of these needs. There is nothing wrong with this. I want my new car to have better features and perform better than my 10 year old car and accept that means changes to how the car works, and the skill sets needed to services it.



In our business as a provider of automated operational analytics software, we have to be able to connect to and manage data from a wide range of devices, applications and data sources used by our system integrator partners, which now number over 115 across the world. Fortunately, an increasing number of devices are coalescing around a reasonable number of accepted standards. Some are the device protocols we are all familiar with from the BAS industry like BACnet® IP, oBix™, and LONTalk®. From the IT domain we also see the SNMP growing in importance as data from IT assets is brought into integrated facilities management systems.

Newer protocols, like the Haystack protocol, are seeing rapid acceptance due to their ability to communicate semantic information that makes data self-describing. The addition of semantic information dramatically reduces project engineering costs which quickly justifies the learning curve for these new communication technologies.

Other important communications options come from the software application world. Here we can think of XML, connection to SQL databases, CSV file formats, (often the “common denominator” data format), and REST style communications. REST stands for representational state transfer. It is the primary style of communication of the World Wide Web, and is supported by the majority of modern, web-based software applications - SkySpark included. REST can reasonably be said to have revolutionized inter-application communications and data sharing. The Haystack communications protocol is REST interface. You can learn more about REST here: [https://en.wikipedia.org/wiki/Representational\\_state\\_transfer](https://en.wikipedia.org/wiki/Representational_state_transfer)

The latest addition to the SkySpark suite of native connectors is OPC-UA. OPC is most well known for its widespread use in industrial automation. Numerous products in the BAS market use it as well. The OPC Unified Architecture (UA), was released in 2008. It is a platform independent, service-oriented architecture that integrates all the functionality of the earlier “OPC Classic” specifications into one extensible framework. You can learn about OPC UA here: <https://opcfoundation.org/about/opc-technologies/opc-ua/>

With the addition of native support for OPC UA, another large population of devices, systems and applications can now be connected directly to SkySpark without the need for external gateway devices or software.

## Project Haystack Announces Two New Publications

Project Haystack recently announced two important new publications supporting the organization's work. The first, a "Project Haystack" white paper, has been published in cooperation with the **Continental Automated Building Association (CABA)**. The second is the inaugural issue of "**Haystack Connections**", the Project Haystack community's new magazine. Both publications provide insight into how the Project Haystack methodologies are helping organizations unlock the value of data by making that data self-describing through the use of metadata tagging.

A complimentary copy of the CABA white paper is available at:

<https://www.caba.org/CABA/DocumentLibrary/Public/Project-Haystack.aspx>

A complimentary copy of Haystack Connections is available at:

<http://project-haystack.org/download/file/Connections-Issue-01.pdf>



Since its formation in 2011, The Project Haystack organization (a 501c non-profit trade association) has grown tremendously, providing the industry with an open-source, collaborative environment where people work together to address the challenge of utilizing semantic modeling (also known as tagging) to make data self-describing and thereby streamlining the interchange of data among software applications. The community has developed a flexible, extensible data-modeling approach and standard models for common equipment systems. The standard includes detailed documentation describing the data-modeling techniques, significant libraries of consensus-approved equipment models, and software reference implementations to enable software applications to easily consume smart device data that is "marked-up" with Project Haystack data descriptions.

The work developed by the Project Haystack organization and community of supporters, is streamlining the process of managing, presenting and analyzing the vast amount of data produced by smart devices and equipment systems, and the techniques can be used with virtually any type of system data. The organization's work is not tied to any one vendor or communications protocol.

### Haystack in the News

Navigant Research published an article entitled "Overcoming the Building Big Data Challenge" by Alvin Chen — March 1, 2016. The article addresses the fact that while the cost of sensors has dropped, the amount of computational power and data storage has increased, and the amount of building data available has increased, the key challenge is how to manage and get useful information out of the data.

The article goes on to mention Project Haystack and the work the community is doing to address the challenge. Read the Navigant article here:

<http://www.navigantresearch.com/blog/overcoming-the-building-big-data-challenge>

To join, or for more information on the Project Haystack organization, visit: [www.project-haystack.org](http://www.project-haystack.org).






# SkySpark® - Analytics for a World of Smart Devices

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.*

## SkySpark 3.0 - Features Continue to be Released with Much More to Come !!!



In late 2015 we announced key parts of the SkySpark 3.0 roadmap. Since that time we have been busy delivering a range of features outlined in the Roadmap including:

The Tariff engine and Rate Modeler Functions were released as part of the 2.1.13 build in December 2015. The Tariff engine enables energy and spark costs to be calculated based on complex energy rates. You can find details in the Dec 2015 Insider: <http://www.skyfoundry.com/file/155/SkyFoundry-Insider-December-2015---Introducing-Rate-Modeler-and-Tariff-Engine.pdf>

The Historian features to address data quality, handling of missing data (known as not available or "na" data), and data provenance were added in the 2.1.14 release in January 2016. We also added Calculus functions to SkySpark's math capabilities.

A major new feature was added based on discussions at that occurred at our fall 2015 partner meetings – an OPC-UA Connector. See details on page 7 of this newsletter.

Next up is a major enhancement to the Folio database (being released in May) that will provide even greater speed and even more efficient data storage. Tests run on actual customer project databases show that for many projects the disk space required to store 3.0 data will be only 25% of that required for SkySpark 2.1!

Watch for more announcements on SkySpark 3.0 releases soon.