



"Find what matters"™

The OEM Software Platform for Smart Device Data Applications

A Comprehensive IoT Software Platform for
OEMs and SaaS Providers

*Shortens Time to Market and Reduces
Software Development and Support Costs*

Deploy in the Cloud or On-premise

CONTENTS

The OEM Software Platform for Smart Device Data Applications.....	3
SkySpark - A True Data-Focused Platform: Key Differentiators.....	4
A Wide range of Deployment Options to Fit Your Business Needs.....	5
The IoT is a Distributed Computing Challenge.....	6
SkySpark® Everywhere™ for Embedded Device Applications.....	7
SkySpark® – Software for a World of Smart Devices.....	8
Contact SkyFoundry.....	9



The OEM Software Platform for Smart Device Data Applications



SkySpark® by SkyFoundry is an extensible, software platform for smart device data that can be applied from the edge to the cloud. SkySpark provides data collection, data storage, data management, automated analytics and data visualization for Internet of Things applications in the built environment.

SkySpark® for OEMs and SaaS Providers – A Comprehensive Data Platform Shortens Your Time to Market and Reduces Your Software Development and Support Costs

If you are planning to offer your customers solutions to collect, manage and create value from the data produced by smart, connected devices you will invariably encounter the need to buy or build software to address the following needs:

- Connectors and APIs to access data from external devices, systems and applications
- Examples: Bacnet IP, Modbus TCP, Obix, Haystack, SNMP, Sedona, OPC UA, MQTT, SQL, CSV import (manual & automated), and REST API
- A Database optimized for efficiently managing huge volumes of “machine data” – time-series, multi-structured data with semantic modeling functionality
- Historian functions and services optimized to query and process large volumes of timestamp value pairs
- Analytics engine – to automatically process rules and algorithms against the data to detect issues, patterns, correlations, anomalies and trends in the data
- Visualization – Ready-to-go Apps that automatically generate rich displays of data and analytic results, and tools that allow you to create customized user Apps and reports
- A Distributed Computing architecture to enable deployment from the Edge to the Cloud

- Multi-tenancy support – a single server can host multiple customers, each with their own secure data realm
- Security: Certificate based authentication, TPM integration, HTTPS in the browser, LDAP support for centralized credential management
- Automatic notification of detected issues via email with links and report attachments

SkySpark provides all of these functions in a ready-to-go platform that has been proven across thousands of applications.

SkySpark - A True Data-Focused Platform: *Key Differentiators*

If you are working with IoT devices and their data, you likely know that machine data is different than conventional IT data. SkySpark has been designed from the ground up for the unique needs of sensor, machine and device data.

Device data is by nature “time-series” data – a sensor value with a corresponding time stamp. High resolution device data can yield an immense volume of timestamp/value pairs. For example, if we store minutely data for a single sensor (we call a sensor a “point”) we end up with 525,600 samples per year. 100,000 points with two years of minutely data results in 105 billion samples! Although that isn't actually an extremely large system for SkySpark, a traditional database with 100 billion rows would be considered a really big database. In order to provide high speed processing of data and a responsive user interface for viewing and analyzing data, efficient storage and processing of these volumes of data is a top priority for any IoT software platform and an area where SkySpark excels:

- SkySpark’s “Folio” database provides highly efficient data storage (>10 times more efficient than relational database technologies). A time-stamp value pair averages 12 bits of storage space in SkySpark. Most relational databases require 30-40 bytes to store the same data – some require much more. Beyond highly efficient storage of time series data, Folio adds the ability to model the meaning of data and define relationships between data items with semantic tagging. SkySpark uses the **Project-Haystack.org** meta data tagging methodology to represent the semantic meaning of data items.
- SkySpark integrates a high speed analytics processing engine with the Folio database. With SkySpark, analytic rules “run in the data” eliminating the delays and complexity found in software architectures where large data sets have to be pulled from the database and passed to a separate analytics process.

- Analytic rules and algorithms are fully programmable to meet the unique needs of almost any data type, equipment system and applications – just add your domain knowledge to SkySpark’s library of over 500 analytic and data transformation functions.
- System architecture is important too. SkySpark is scalable and distributable – it can run in everything from small “edge” devices to virtual servers in the cloud – on your favorite operating system
- SkySpark supports both cloud or on-premise deployment for server level applications to address the needs of applications that simply can’t allow huge volumes of data to be transmitted to the cloud
- This comprehensive platform provides lower Total Cost of Ownership than data technologies designed for business IT applications
- SkySpark is proven and ready to deploy – today SkySpark is used across of 11,000 facilities consisting of over 900M square feet of space (>80M square meters)

A Wide Range of Deployment Options to Fit Your Business Needs

SkySpark gives you the flexibility to deploy solutions that fit your customer’s needs and your business model – whether that is Software as a Service, on-premise deployment, or embedded products.

- Install SkySpark as an end-use application for facility managers, energy managers. Install on a laptop or desktop or server, or offer as a cloud service (SaaS)
- Use SkySpark as an in-house, tool for data acquisition and analytics - example for consultants involved in commissioning, energy analysis, monitoring and verification.
- “White box” model – Use SkySpark as a backend database and analytics engine and integrate with other data visualization applications for final presentation of data and results to users
- Embed SkySpark in small, low cost hardware devices (i.e., gateways, controllers) – collect data locally, process analytics at the edge, and serve rich visualizations directly from the edge device and/or replicate data to the cloud and serve visualizations from the cloud
- SkySpark’s fully distributed architecture supports deployment across large numbers of nodes that work together in “clusters” providing a near-seamless user environment that enables users to navigate their applications from the Edge to the Cloud

The IoT is a Distributed Computing Challenge

We hear lots of talk about the “cloud” as it relates to the IoT. 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 SkySpark® Everywhere™ Distributed Informatics Architecture does just that – giving you the flexibility to deploy the architecture needed for your unique application.

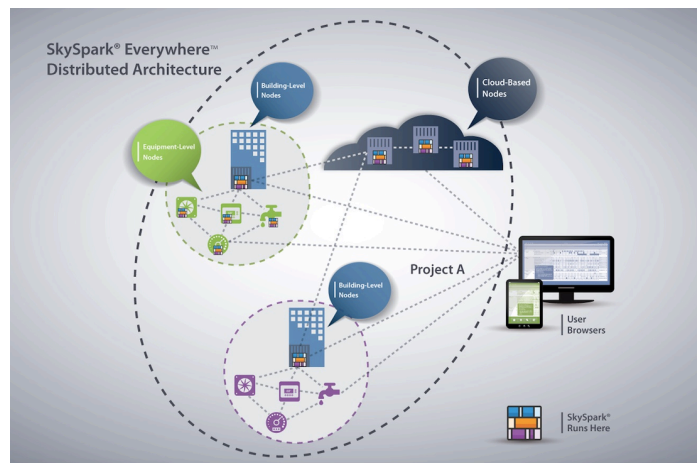
SkySpark Everywhere allows you to deploy the full SkySpark feature set across large numbers of computing devices – from servers to small edge devices – that work together to provide users with a near seamless interaction at the system level, while at the same time providing the full SkySpark feature set locally in edge devices.

SkySpark’s Distributed Informatics Architecture addresses key challenges for the next generation of IoT applications including:

- Seamlessly disseminating data and computation between edge devices and server/cloud environments
- “Sharding” both data and computation across a cluster of servers/nodes of all sizes
- Providing a seamless, responsive user experience across systems from 10’s to thousands of sites and millions of points
- Providing greater fault tolerance with minimal downtime
- Optimizing communications costs and efficiencies when systems are connected over cellular networks or intermittently connected networks

Learn more about the *SkySpark Distributed Informatics Architecture* here:

<https://skyfoundry.com/file/226/SkySpark-Everywhere-Distributed-Informatics-Architecture---White-Paper.pdf>



SkySpark® Everywhere™ for Embedded Device Applications

SkySpark has been designed from the ground up to be embeddable in small, low cost hardware and SkyFoundry offers an OEM program for hardware manufacturers. Our embedded OEM pricing model provides a very low “point of entry” for small devices starting at point counts of 10 points (sensors/control) per embedded device. Even at the smallest system sizes, all of the features of SkySpark are included:


- Communications connectors to major standard protocols including Bacnet IP, Modbus TCP, Obix, Haystack, SNMP, Sedona, OPC UA, MQTT, SQL, CSV import (manual batch or automated), and a REST API;
- The high speed Folio Database
- The SkySpark analytics engine
- All SkySpark Apps for viewing data and analytic results
- Our secure clustering technology that connects multiple nodes in a unified distributed architecture

The result is a comprehensive platform to get your IoT products and services to market quicker, at less cost, with the features your customers want. SkySpark has been successfully applied on off-the-shelf or custom designed embedded devices in the class of **Beaglebone Black™**, **Raspberry Pi** and other hardware with similar performance specifications. SkySpark is also qualified for use on the Dell Edge Gateway™ 3000 and 5000 and on a number of OEM hardware products.

Contact us today to discuss your application and learn more about our programs for OEMs and SaaS providers. We will be happy to provide you with all of the information you need to evaluate the SkySpark platform for your applications.

To learn more about SkySpark, including joining one of our live demonstrations, contact us today!


info@skyfoundry.com



SkyFoundry
PO Box 1749
Glen Allen, VA 23060

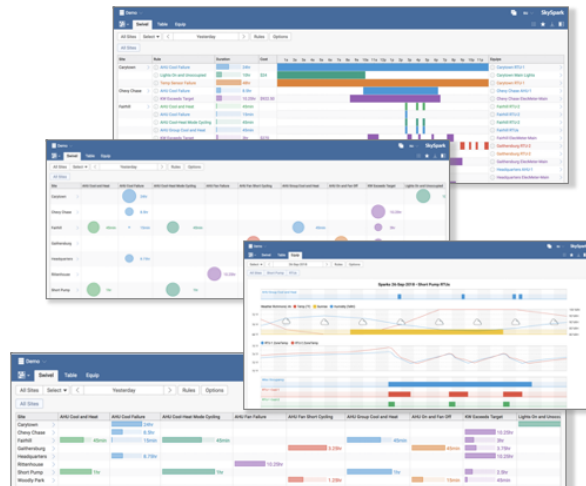
SkySpark® – Software for a World of Smart Devices

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



Access to this data opens up new opportunities for the creation of value-added services to help businesses reduce energy consumption and operational costs and to identify opportunities to enhance operations through improved control, and replacement or repair of capital equipment.

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. SkyFoundry's SkySpark® directly addresses this challenge.



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

ABOUT SKYFOUNDRY

SkyFoundry's mission is to provide software solutions for the age of "the Internet of things". Areas of focus include:

- Facility Automation and Management
- Remote device and equipment monitoring
- Energy management, utility data analytics
- Asset management

SkyFoundry products help customers derive value from the data in smart systems. Contact us to learn more.

info@skyfoundry.com

<https://skyfoundry.com>

SkyFoundry