



SkyFoundry Insider

Issue No. 40

February 2022

In This Issue

SkySpark’s GhG App Enhancements - Make It Easy to Calculate, Track and Report Greenhouse Gas Performance2

New Schedule App – Define Schedules and Supervise External Systems3

SkyFoundry Open Sources Core Software to Accelerate the IoT....4

The Next Step in SkySpark Machine Learning – Python Support.....5

Cloning Records Across Distributed SkySpark Systems – Advancing Edge-to-Cloud Compute6

SAML-based SSO (Single Sign On) and LDAP Enhancements...6

SkyPosium 2021 – A Quick Look Back.....7

The Technology Issue

Recent Additions to SkySpark and a Look at the “Why” Behind Them

The last year has seen major new features added to the SkySpark platform—new Apps for users, new data visualization features, and many other significant new releases.

In addition to bringing clients the most choice in the industry, one of the significant benefits to SkyFoundry’s worldwide network of partners is that we get input from around the world from partners that specialize in different market segments and apply SkySpark in a wide range of applications. They provide us with tremendous input into market needs, customer requirements, and product features that continue to make SkySpark the most deployed analytics solution in the built environment.

In this issue of the Insider, we look at some of the most notable features and focus on the “**why**”—why these made it to the top of our development priority list.

Let’s take a look! →

SkySpark's GhG App Enhancements - Make It Easy to Calculate, Track and Report Greenhouse Gas Performance



WHY – Corporate and government support for sustainability and “ESG” related policies (Environmental Social and Governance) is growing across all segments of the CRE industry. ESG issues are increasingly seen by shareholders as indicators of a company’s future success and are now a core component of most annual reports by public companies. Key elements of the “E” portion include energy efficiency initiatives, and Greenhouse Gas Emission tracking and reporting. Adding the capability to calculate, track and report GhG metrics in SkySpark® was a logical extension of our continually growing suite of energy focused applications.

THE LATEST UPDATES: In early 2021, SkyFoundry introduced the **GhG App** as a standard SkySpark feature. The GhG App integrates with online data resources (including the US Environmental Protection Agency or EPA) to make it easy to quickly configure location-specific emission factors to accurately convert energy consumption into GhG equivalents for use in visualizations, reports, KPI's and analytics. The App presents GhG performance in a wide range of engineering and “societal” metrics, i.e., miles driven, trees planted, smartphones charged, and many more.

We received quick and positive response to the GhG App and now, after input from energy and sustainability professionals, have captured that input in our recent GhG App release. **New features include:**

- GhG values calculated and visualized with the same resolution as meter data, whether that is 15-minute interval data or more frequent meter readings. This is especially important as more and more facilities adopt use of multi-source energy including renewable resources.
- GhG Emissions can now be visualized as the total for each emission source or meter in addition to summarizing each site.
- Users can select a baseline of a previous month or previous year to analyze emission performance.
- A new built-in view makes it easy to display equivalent emission values of societal metrics in SkySpark’s “card” format.
- New chart customization options including multi-axis charts (see image at right).
- A number of predefined emission factor models have also been added based on published data from the UK Government and British Columbia Provincial Government.



New Schedule App - Define Schedules and Supervise External Systems





WHY – Ah... schedules – simple stuff right? Not so fast. One of the surprising challenges in acquiring data from BMS and similar systems is the difficulty, or often inability, to get schedule details in a usable format. They are often stored very differently than sensor values. Given that schedules are directly associated with evaluating proper equipment operation there is a need to be able to represent schedules for use with analytic rules.

Another important use is in calculating true energy costs with the SkySpark Tariff Engine to make decisions on demand response actions or fuel source switching or other control actions. And, with SkySpark’s ability to issue commands, it can act to supervise schedules in connected systems.

SkySpark’s **Schedule App** enables users to define schedules to replace or add to the schedule information available from their control systems. The all-new Schedule App provides a quick easy way to define schedules in a graphical format. Additional features include:

- The ability to assign a color to each item and view color either by item or by value.
- A new schedule rule that allows you to easily match dates based on the "closest weekday". This rule is useful in cases where you need a holiday to be observed on the closest weekday if it falls on a weekend. Many tariff schedules specify holidays that follow this kind of rule.



SkyFoundry Open Sources Core Software to Accelerate the IoT

WHY – Acquiring data from building systems, sensors and other smart devices remains one of the key challenges in applying analytics and moving to data-driven facilities. Gateway development is not easy. Often the companies with the best hardware development capabilities do not have deep experience with communication connectors and in the necessary web-technologies to bring data to higher level applications securely. SkyFoundry has extensive experience with both hardware and software assets that have been proven in the field—many for as long as 10 years. This proven code is being open-sourced to help accelerate the development and reduce the cost of creating IoT devices for the built environment.

Haxall is an open-source software framework designed to streamline the process of building IoT data products—both hardware and software. It includes a core subset of proven SkyFoundry code that addresses critical functions needed to connect to devices and equipment systems, normalize their data, and make it available to other applications in open standard formats.

Haxall provides the software functions to communicate with external devices, capture semantic tagging at the edge, and communicate with external applications. It is powerful enough to run sophisticated applications **at the edge** or use it as a gateway to bring IoT data to cloud-based applications.

What Does Haxall Include?

Full Haystack 4 support. Haxall is bundled with SkySpark's full suite of Haystack APIs. This includes a rich set of APIs to model, encode, and query data using the Haystack 4 ontology.

Folio. SkySpark's revolutionary Folio database technology is at the heart of the Haxall platform. Folio is built 100% around the Project Haystack data modeling standard. Everything is stored as native Haystack "dicts" and queried using Haystack filters. Folio provides seamless integration with persistent and transient real-time tags making it a truly unique design to build IoT applications.

Axon. Haxall includes the Axon functional scripting engine. Create scripts to onboard, query, and transform your Haystack data using a rich library of hundreds of functions. The Axon function library has been leveraged for more than a decade to streamline working with Haystack data.

Haxall Daemon. A complete runtime called the "Haxall Daemon" is provided to quickly get your IoT applications up and running. The Haxall Daemon comes with "batteries included": tools to create new databases, built-in user management and authentication, a webserver with full support for the Haystack HTTP API, dynamic module management, and an Axon shell for remote management. In addition, Haxall provides runtime support for both the Java VM and JavaScript environments.

Connectors. Haxall includes a connector framework to acquire data from external devices and normalize diverse IoT protocols into the Haxall ecosystem. A suite of ready-to-use connectors is provided as open-source including **Haystack HTTP API, MQTT, Modbus, oBIX, SQL, and Sedona Sox**. An in-memory trending engine is included to historize connector data into time-series data. Haxall also includes an Arcbeam™ module that provides seamless connection of Haxall devices into SkyFoundry's SkySpark software.

How Haxall Helps Accelerate the BIoT. Haxall streamlines development and reduces the cost of building IoT data acquisition devices by providing proven software components. Haxall is available as open-source under the widely adopted Academic Free License ("AFL") v. 3.0 <https://opensource.org/licenses/AFL-3.0>.

The initial release of Haxall is available now on [GitHub](#)

The Next Step in SkySpark Machine Learning – Python Support



WHY – The Python programming language has emerged as the “go to” favorite for implementing machine learning (ML) and many data scientists in the industry have extensive experience implementing ML in Python. They look to use the toolset they are familiar with.

With the latest release SkyFoundry integrates **Python** into SkySpark to take your machine learning applications to the next level. SkySpark now allows you to run Python code in a sandboxed Docker container and return the results back to SkySpark.

To simplify use of the Python integration, SkyFoundry provides a default Docker image that has several common ML Python packages installed and ready to go (e.g. NumPy, Pandas, scikit-learn).

We have also written an application note to show users how to switch over to Python from the original SkySpark mExt which was based in the LBNL SMILE library.

Where will your imagination and data science skills take you? With SkySpark your options are wide open.

Switching from the mExt to Python for Machine Learning

[Overview](#) • [Setup](#) • [Porting](#)

Overview

This guide covers the high-level steps required to switch from the mExt to Python for doing machine learning in SkySpark. The old mExt made use of an open-source library called SMILE. For various technical reasons, it no longer makes sense to support that library in SkySpark. The mExt is considered deprecated, and will be removed at some point in the future.

Instead, beginning in SkySpark 3.1.2 we have made it possible to do machine learning using Python in a Docker container. We believe this approach to be more sustainable, robust, and will enable the use of all sorts of algorithms that simply were not available in the mExt.

Setup

You will need to enable the Python extension in SkySpark (which will automatically enable the Docker extension). After enabling those extensions, you should follow the instructions in the [Python docs](#) for getting your Docker environment setup for executing Python code from SkySpark. You should then read the full Python docs to get a picture of how the extension is used.

SkySpark - The Platform for Operational Data Analytics

Connect – Includes BACnet IP, Modbus TCP, Obix, Haystack, SNMP, Sedona, OPC UA, MQTT, SQL, Nest, Ecobee, CSV import, REST API, and a connector development toolkit

Collect - Store vast amounts of device data with SkySpark’s highly efficient Folio database

Transform - Normalize data and relationships across diverse data sources

Analyze - Automated analytics processes rules and algorithms on your data

Present - Apps visualize your data & analytic results

Deliver – Transmit data and issue commands to external applications via a variety of standard formats

Meet Real-World Application Needs – All buildings are different. SkySpark combines full user programmability with an extensive library of over 500 analytic functions, advanced math and now Python-based MML to enable you to meet diverse application needs

SkySpark brings together sensor and equipment data, web-services data, historical data, analytic results, KPI’s, energy cost and more in a single platform and makes all data available to external applications via our open API.

Cloning Records Across Distributed SkySpark Systems – Advancing Edge-to-Cloud Compute



WHY – The story starts with SkySpark’s unique distributed architecture – the ability to distribute SkySpark nodes from the edge to the cloud to create a true distributed data and compute platform. When you have a distributed computing platform you need technology to manage the software in the individual nodes and automatically share data records across nodes.

The management of software modules in multi-node SkySpark systems is handled by our comprehensive Provisioning tools. Cloning adds the ability to share data records across distributed nodes seamlessly and automatically.

Cloning enables simple, efficient replication of individual records between projects over an Arcbeam cluster. Once a “clone set” is defined those records will automatically be replicated into projects and kept synchronized. Cloning provides a very simple mechanism to share functions, apps, views, and templates across projects. Example, having a central standard library of analytic functions and sharing them across a multi-node system. All updates and changes are automatically shared through the Cloning feature.

SAML-based SSO (Single Sign On) and LDAP Enhancements



WHY – Integration with Enterprise level user management systems is an essential requirement for centralized management of user accounts and security permissions in large scale applications. Two widely used protocols for integration with enterprise identity providers are LDAP (Lightweight Directory Access Protocol) and SAML (Secure Assertion Markup Language). LDAP provides a Single Login experience so users can login to SkySpark using the same username and password they use for all enterprise applications. SAML provides Single Sign On (SSO) so once users are logged in to their identity provider, they don’t even have to type their credentials into SkySpark providing a more seamless experience.

In 2021 we fully released our SAML-based SSO feature for SkySpark. We also added numerous enhancements to SkySpark’s remote user management capabilities providing more flexibility for management of user records via LDAP or SAML. Now it’s possible to have the enterprise user management system provide authentication only, with authorization (role, access filters, etc.) controlled in SkySpark, OR the enterprise system can manage and update authorization dynamically. It’s also possible to use a combination of these techniques where some authorization parameters are configured in SkySpark while the enterprise system sets other authorization attributes. The enhanced flexibility simplifies enterprise integration and avoids changes to the existing centrally managed user database.

SkyPosium

2021 | The Worldwide SkyFoundry Community Event

A Live Event

October 13-14, 2021 | Denver, CO

The Event for the entire SkyFoundry User Community

SkyPosium 2021 – A Quick Look Back

One of the most exciting events of 2021 was the return to an **in-person SkyPosium** for the SkyFoundry community. SkyFoundry opened the event with a general session where we introduced many of the new advances highlighted in this newsletter, a preview of our roadmap, and tech sessions for developers.

As always, major portions of the program were delivered by the community who presented case studies and new applications. We recorded all of the sessions for those who could not join us in person and published all of the content on our support site. Here is a quick view of the program:

Day 1

- [Welcome and Opening Comments](#) John Petze, SkyFoundry
- [Haxall - the Open-Source Initiative to Accelerate the IoT - Strategy, Capabilities, Use Cases](#) Brian Frank, SkyFoundry
- [Machine Learning/Python/Docker in SkySpark - Next Generation ML, a Preview](#) Brian Frank, SkyFoundry
- [The Greenhouse Gas App - Live Demonstration](#) Ross Schwalm, SkyFoundry
- [The All-New Doc Structure and Search Capabilities](#) Brian Frank, SkyFoundry
- [Tour of New UI features - Bet You Don't Know Them All! - Overview and Demonstration](#) Brian Frank, SkyFoundry
- [Arc Workflow Tool Demonstration and Overview of Customization](#) Paul Bergquist, SkyFoundry
- [Product and Pricing Updates](#) John Petze, SkyFoundry

Day 2

- **Developer Sessions**
 - [Haxall - Developer Session](#) Brian Frank, SkyFoundry
 - [Single Sign-On, SAML, LDAP and Security Enhancements \(HTTPS Embedded Webserver\) - Technical Review](#) Ross Schwalm, SkyFoundry
 - [Greenhouse Gas App - Technical Review](#) Ross Schwalm, SkyFoundry
 - [Building Dashboards with Card Views - a Live Tutorial](#) Paul Bergquist, SkyFoundry
 - [Using the Notification Framework](#) Ross Schwalm, SkyFoundry
 - [Cluster Cloning - Record-Based Replication](#) Brian Frank, SkyFoundry
- **Community Presentations**
 - [Predictive Maintenance of Air Filters](#) Gia Nguyen, Altura Associates
 - [The Fastest Ice on Earth - Analytics at the Utah Olympic Park Slide Track](#) Greg Schlegel, BuildingFit
 - [BASGG - ASO - Executing Control Based on Analytics](#) Alper Üzmezler, BASGG
 - [Monitoring-Based Commissioning Program Using SkySpark: A Success Story at a K-12 School District](#) Brent Grimm, ETC Group
 - [Developing a Site Build Tool Easy Enough Even for Me to Use \(Successfully\)](#) Jeff Cheek, BuildingFit
 - [Group14 Chiller Optimization](#) Celeste Cizik, Group 14 Engineering
 - [Rule Tuning View](#) Chris Spurlock, BuildingFit
 - [SkySpark Rule Tester App](#) Alper Üzmezler, BASGG
- **Additional Community Presentations**
 - [Fantom Factory Training Programs](#) Emma Eynon, Fantom Factory
 - [Fantom Factory ViewBuilder Tools](#) Emma Eynon, Fantom Factory
 - [Energy Twin Machine Learning Extensions for Energy Data](#) Jan Siroký, Energocentrum
 - [StackHub Overview](#) Steve Eynon, StackHub
- [Roadmap, Q&A Discussion](#) SkyFoundry



The event also included a **vendor showcase** that provided attendees with the opportunity to meet with companies that offer complementary products and services to the SkyFoundry community.

Plans are already underway for SkyPosium 2022 – Stay Tuned!



The Leading IoT Data and Analytics Platform for the Built Environment

SkySpark® Analytics automatically analyzes data from building automation, metering systems and other smart devices to identify issues, faults and opportunities for savings. Learn why SkySpark has been deployed to over 1 Billion square feet of facilities around the world for energy management, optimization, monitoring-based commissioning and fault detection.



Find What Matters™ to Improve Equipment Performance and Reduce Operational Costs.

SkyFoundry

www.skyfoundry.com

**Learn More About SkySpark®
and How to Apply the Industry-
leading Data Analytics Solution
to Your Application**

*Join us for a comprehensive
demonstration webcast*

Find our calendar of upcoming sessions and
other events here:

<https://skyfoundry.com/calendar>

Or contact us at: info@skyfoundry.com