FALL 2019

Project Haystack

ISSUE 06

connections

Journal of the Haystack Community - Solutions for Interoperable Device Data

Collaboration

Solving the Data Challenge as a Community

• Data Tagging at Ford Motor Company • Introduction to Haystack 4
• ARC Review • PitchFest • Working Group Updates

www.project-haystack.org

Project Haystack

Founding Members

ConserveIt

intel

J2 INNOVATIONS
A Siemens Company

legrand®

LYNXSPRING®

SIEMENS
Ingenuity for life

SkyFoundry

Associate Members

Accu Temp
Building System Integrators

ALTURA

BAS SERVICES GRAPHICS

bueno
Built Environment Optimisation

buildingfit
Making SkySpark® Work For You.

CABA
Continental Automated Buildings Association

EMA
ENERGY MANAGEMENT ASSOCIATION

Intellastar

INTELLIGENT BUILDINGS™

IOT

KMC CONTROLS

KNX®

KODARO

sensorFact®

SMARTGREEN

TRIDIUM

WideSky

Yorkland Controls
ENVIRONMENTAL SOLUTIONS

CONTENTS

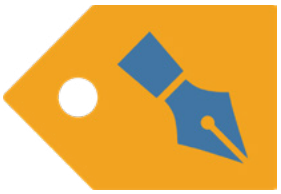

EDITORIALS

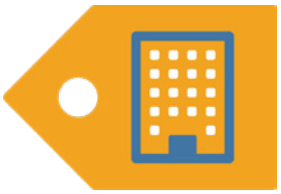

PROPERTY MANAGER'S VIEW

CONTRIBUTED ARTICLES

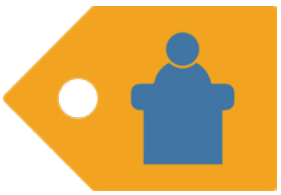

ANALYST'S VIEW

HAYSTACK
FROM THE GROUND UP.

Enjoy native Haystack tagging that can be used in a massively parallel manner.

75F comes pre-tagged, allows you to simultaneously use machine learning algorithms across multiple pieces of equipment, and will automatically update algorithms as ASHRAE GPC 36 (High Performance Sequences of Operation for HVAC Systems) evolves. Stop settling and harness the power of cloud with 75F's software-defined hardware and Building Intelligence Solutions.

Download our latest case study & step up to Haystack at cloud scale.

LEARN HOW WE SAVED HOM FURNITURE $45k+ IN ENERGY COSTS

75f.io/hom-case-study-ph

- **Reduce Energy Costs by 30-50%**
- **Enjoy an Average Payback in Less Than 2 Years**
- **Manage Multiple Buildings from a Single Dashboard**
- **Installed in Half the Time of the Competition**
- **Enjoy Wireless Communication and Bluetooth Pairing**
- **Control Several Equipment Types on a Single Platform**

www.75F.io | @75F_io | (888) 590-8995

CONTENTS *cont'd*

RECAPS OF PITCHFEST

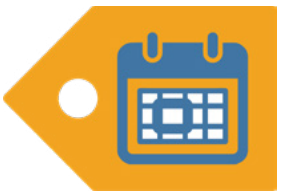

EVENTS

WORKING GROUP UPDATES

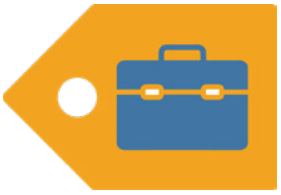

TOOLS FOR DEVELOPERS & INTEGRATORS

SOCIAL MEDIA CURATION

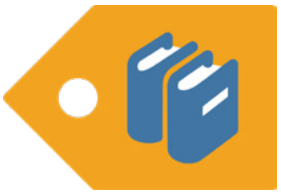

DIRECTORIES

World's first hardware agnostic DDC Project Sandstar.

Join Project Sandstar

https://bas.co/projectsandstar

Better Data that Builds Cars, Smarter Communications and Smart Cities

by Robin Bestel, Managing Editor, Project Haystack Connections Magazine

Welcome to the Fall 2019 issue of the Project Haystack Connections Magazine. This issue is a true testament to the collaboration that the Haystack Community is committed to and the accomplishments that have been made this past year and the ones anticipated as we head into 2020.

In March, Project Haystack exhibited at ISH in Frankfurt, Germany and in May, we held our biggest collaborative event, the **2019 Haystack Connect Conference**, which brought this community together again for three days in San Diego. We held a **PitchFest** during the conference and have included a few of them in this issue. The Haystack Connect 2019 Speaker Presentations are available at **www.haystackconnect.org**. In September, a one day, first-ever **Project Haystack event** was held in London and sponsored by J2 Innovations. A recap and videos of the presentations are available at "**Project Haystack Went to Europe**".

Looking ahead, automatedbuildings.com is collaborating with us on an educational session "**Haystack 4 - The Continued Evolution of Semantic Tagging – What it Is and Why it Matters**" during AHR Expo 2020. And, to show our appreciation for their support, I want to mention that ControlTrends has opened their ballots for the 2019 **ControlTrends Awards**, being held February 2, 2020 at B.B. King's Orlando.

This issue holds the most contributed articles, **11 in total**, by Project Haystack members and supporters that include Ford Land, ARC Advisory Group, Lynxspring, Tridium, Conserve It, Open Source Strategies, J2 Innovations, Intelligent Buildings, BASSG and Bueno Systems. Brian Frank provides us with a full update on Haystack 4, as well. I thank all of you for "collaborating" on another, very informative Connections Magazine.

One of our presenters at Haystack Connect 2019 has contributed to our **Property Manager's View**. Tim Scott is a building management system subject matter expert at Ford Land, a subsidiary of Ford Motor Company, that manages over 5 million square feet of commercial office space. He tells how they are already implementing Haystack tagging as part of their **10-Year Campus Reconstruction Project**.

For our **Analyst's View**, Larry O'Brien of the ARC Advisory Group, has allowed us to republish his recent article "***Project Haystack Turns Data Into Useful Information For Smart Cities***".

Project Haystack **Working Groups** contributed updates to their work. Working Group **Cybersecurity**, championed by Fred Gordy of Intelligent Buildings and **ATES Systems**, championed by Jaap Balvers of BAM Energy Systems, part of Royal BAM Group nv, are just getting started and looking for "Collaborators" to join their group.

As always, we have a section dedicated to **Tools for Developers and Integrators** and **How To Get Involved**, a curation of social media about **Projects, Practices and Products**, and our **Members Directory**. We also list all of our **Advertisers**, for whom we thank for their support and sponsorship of Connections Magazine.

Just in case no one noticed yet, Project Haystack has a new identity!

With the announcement of Haystack 4, it was decided that it was time to update our "Tag Logo".

Project Haystack

We also updated our Project Haystack Founding Member, Associate Member and Supporter logos. Here, in this issue, they are officially announced and are being used.

Project Haystack
Founding Member

Project Haystack
Associate Member

Project Haystack
Supporter

We created a special one that we will use to promote Haystack Version 4.

Haystack 4

Along with this, Project Haystack launched a new **Marketing website** to compliment the **Developer website**. Now, the Developer website is focused solely on the continued development of the Haystack methodology and assignment of tags, the activities of the Working Groups and the developer Forum conversations. The Marketing site is designed to inform the community on all the who is where and when, **Project Haystack Logo files and new Marketing Kit**, press releases, our members list and how to become a member, and of course, the home of our **Connections Magazine**.

It has been my pleasure to work with everyone to publish this Fall 2019 issue of Project Haystack Connections Magazine. The Project Haystack community truly represents "What Collaboration Can Accomplish"!

PlantPRO® performance reliability optimisation

Conserveit

BECOME A RESELLER OF AN AWARD-WINNING CHILLER PLANT OPTIMISATION SOLUTION

Call for PlantPRO VARs, MSIs and
Intelligent Building Services Providers Worldwide

WHY PLANTPRO®?

- Advanced Self Learning Algorithms based on ASHRAE Study
- Chiller Plant Performance Monitoring
- Chiller Plant Analytics and FDD
- Advanced Chiller Plant Optimization
- Chiller Plant Measurement and Verification
- Chiller Plant Control and Automation
- Chiller Plant Continuous Commisioning and Tuning

WHY PLANTPRO® CORE?

- Low Cost Chiller Plant Control System
- Chiller Plant Performance Monitoring
- Chiller Plant Alarming
- Standard Chiller Plant Optimisation
- Chiller Plant Charting
- Chiller Plant Control and Automation
- Chiller Plant Continuous Commisioning and Tuning

FEATURES AND BENEFITS OF PLANTPRO®

- Smart Machine Learning
- Multiple Chiller Manufacturers
- Up to 40% Energy Savings
- Predictive & Preventative Maintenance Diagnostics & Analytics
- Up to 0.35kW/RT or 10 COP Optimized Plant Efficiency
- Easy BMS/CMS Integration

powered by

niagara framework®

Project Haystack
Founding Member

We're looking for PlantPRO VARs, MSIs and Intelligent Building Services Providers Worldwide

conserveitiot.com

plantpro@conserveitiot.com

FIN Framework 5

Open Fast Robust Extensible

Deployable on:

> **Cloud**

> **Server**

> **Controller**

Open | Fast | Robust | Extensible

Whether you want to easily engineer building automation, be faster to market or make better use of data to run buildings more efficiently and productively, FIN is the next generation Building Automation System for you.

FIN is a software technology that uses the Haystack standard, combining the core functionality of a Building Automation System for connecting and controlling devices with the added benefits of a Building Operating System to manage and leverage data.

J2 INNOVATIONS
A Siemens Company

To find out more about how FIN can help your business, visit **www.j2inn.com**

Message from the Board

The Value of Collaboration

This issue of Haystack Connections echoes the central theme of the Haystack Connect 2019 Conference – **Collaboration**. Bringing a true, open-source development effort to the buildings industry which has historically been dominated by prescribed standards and proprietary technology, has been both rewarding and challenging.

The central focus of Haystack is to enable software developers, system integrators, building owners and operators to easily use the data produced by the equipment, devices and sensors that make up the built environment. So, at its heart, Haystack is a software effort. When we look to the world of software, its undeniable that open source movements have changed the landscape, solving problems no single manufacturer or supplier could have addressed on their own. There are some things in life, science and technology that "take a village" to address and Project Haystack is one of those. As an open-source effort, the results are directly correlated to the contributions and participation of the community.

These are exciting times for Project Haystack. We recently released, the 4th generation of Haystack, Haystack 4, and launched a new Marketing Website.

marketing.project-haystack.org

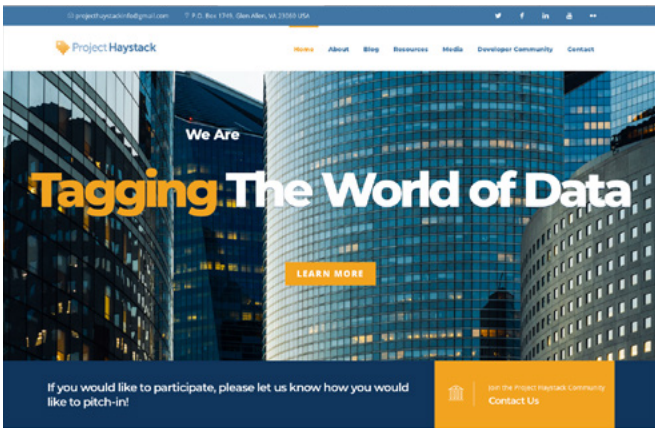

In May, we held the fourth Haystack Connect Conference and last month, due to the efforts of Chris Irwin, Executive Envoy for Project Haystack, we conducted the first-ever European Project Haystack event. Also, in September, Project Haystack was the subject of an article by Larry O'Brien of ARC Advisory Group. And, at the upcoming AHR Expo 2020 in February in Orlando, Florida we will be leading an education session on *The Continued Evolution of Semantic Tagging – What it Is and Why it Matters*.

Adoption of Project Haystack is at an all-time high. It is being utilized in a wide variety of applications and use cases, and deployed in thousands of buildings incorporating millions of square feet globally. That said, we need to continually grow the community and supporters. If you and your company are not supporters yet, consider it. It could not be a better time.

The Board would like to take this opportunity to thank everyone for ***collaborating*** to create another excellent issue of Connections Magazine for Fall 2019.

John Petze
Executive Director
Project Haystack

Marc Petock
Executive Secretary
Project Haystack

Data Tagging at Ford

FORD LAND

Tim Scott of Ford Land, a subsidiary of Ford Motor Company that owns and operates approximately 5 million square feet of commercial office space in Dearborn, Detroit and Allen Park, Michigan, shares how effective semantic tagging is helping to ensure data arrives clean, structured and ready to be processed by Ford Land's selection of applications and tools to better manage their properties.

Ford Land, a subsidiary of Ford Motor Company, owns and operates approximately 5 million square feet of commercial office space in Dearborn, Detroit and Allen Park, Michigan. From multi-story office buildings and business parks with research and development space, to individual offices and move-in-ready suites, they provide a variety of properties designed to meet all types of business needs. Since its inception in 1970, Ford Land has honored its original intent to be an active guardian, manager and developer of what has become a thriving real estate community.

Ford Land is the real estate division of Ford that drives project management, engineering, space planning, dealership design and facility maintenance for all Ford facilities globally. In its **10-Year Campus Reconstruction Project**, they are renovating 20th Century buildings, adding two new data centers, improving manufacturing facilities through energy efficiency programs, creating state of the art building technologies in Corktown, a Detroit neighborhood that has a rich history with Ford Motor Company. We are adding mobility technology centers and offices needed to support our growth towards becoming the worlds most trusted mobility company.

As Ford reimagines its business for a new century, it has turned to Detroit's oldest neighborhood as a launchpad for its plans. In Corktown, we are creating an innovative hub where Ford and its partners will work on autonomous and electric vehicle businesses, and design urban mobility services and solutions. Michigan Central Station and a number of Ford-owned properties nearby will anchor a 1.2 million-square-foot campus that will be home to approximately 2,500 Ford employees and an additional 2,500 people from strategic partners.

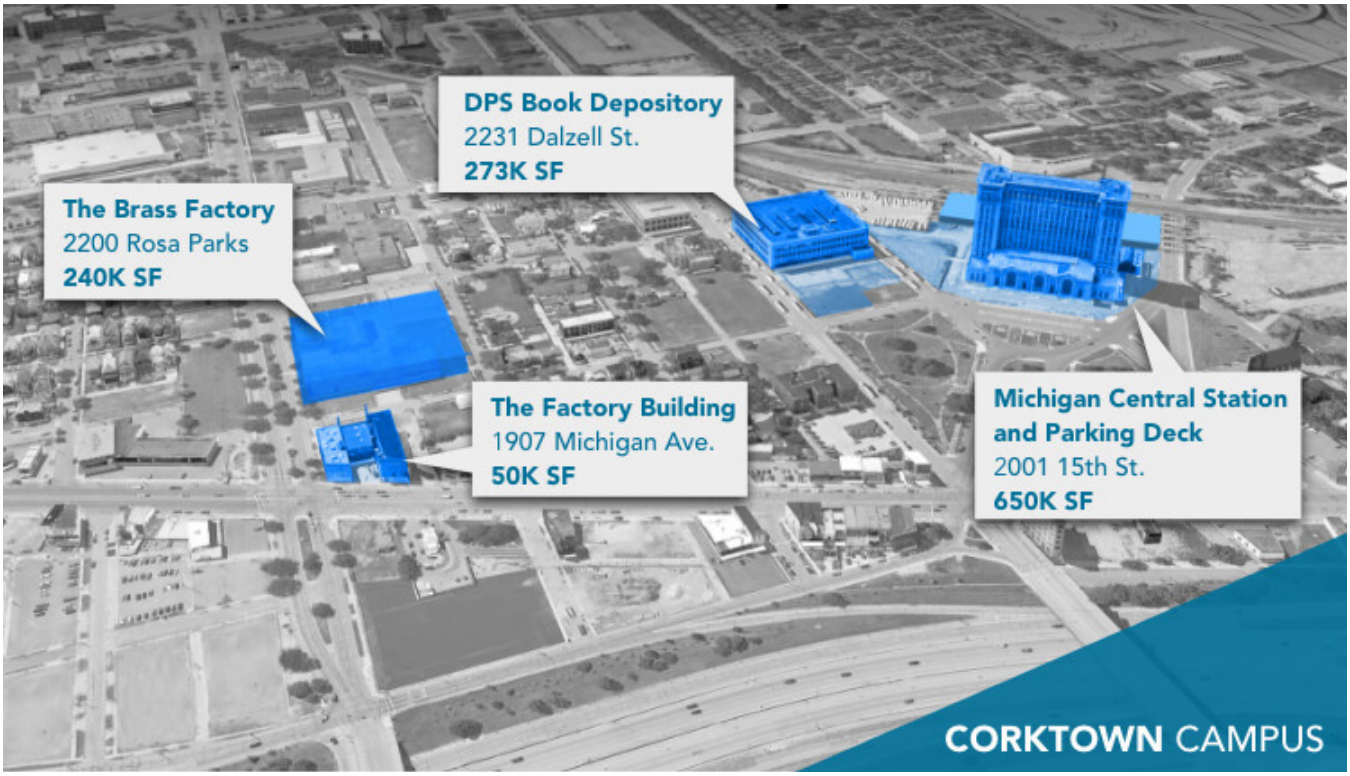

The centerpiece of Ford's Corktown vision is Michigan Central Station, a beloved Detroit landmark that was acquired in June 2018. Ford plans to return the grand hall of the station to its original grandeur and attract local shops and restaurants in support of a vibrant, inclusive public space for all, while the tower will be transformed into modern office space.

With this vision in mind, there are a few guidelines that further direct how Ford Land is deploying Smart Building technology across their diverse portfolio:

1. A joint IT/OT team is involved in all controls-related decision-making, this ensures a completely integrated approach from project inception. A reference topology makes clear to anybody that does system integration or control work on campus what is authorized when connecting into Ford's newly upgraded, secure wireless and IP infrastructure.
2. Establishing a Master Systems Integrator (MSI) role that will be involved in every new construction and controls retrofit project to enforce standards for all building engineers and controls contractors doing work throughout all properties in our portfolio.
3. When adding control software for HVAC, lighting, occupancy measurement, physical security and other applications, sufficient attention is paid to how data is modeled and managed, so that Ford management is sure of its integrity and ownership. Ford wants to ensure successful deployment of data analytics to help with the deployment of autonomous driving which will benefit us going into the future of transportation systems.

"More recently we have been adopting the Project Haystack tagging system. The job of developing mapping tools and methods for translating certain data sets into Haystack tags is already underway. "

Given the scope of what Ford is working to achieve, the core engineering team has established a strong standard for tagging data. Ford-specific naming conventions and tagging dictionaries have been introduced over the last decade with consistency enforced. More recently we have been adopting the Project Haystack tagging system. The job of developing mapping tools and methods for translating certain data sets into Haystack tags is already underway. The construction and operation of data centers and manufacturing plants presents a need for different types of ontologies and tags than the more standard commercial building HVAC areas where Haystack is so strong. Ford Land BMS leaders are grappling with building-out tagging dictionaries.

Ford Land has mapped out a number of approved connection topologies for wired and wireless infrastructure and assets, which are applied depending upon the business function of the facility in question.

Any asset that connects to the Ford Office Automation Network — be it equipment delivering a building service like lighting, HVAC, or physical access control, or be it a metering device for electricity, water, air, etc. — is never allowed direct access to the extranet. All connections are through secure gateways and devices that have been reviewed, tested and approved to meet cybersecurity standards by Ford's IT TCS Network Engineering team. Programmable Logic Controllers used extensively on the manufacturing floor are required to be connected through an access control list provided by Ford IT, in order to communicate over the network.

A private virtual LAN, running at the Ford Enterprise Data Center, hosts Ford Land's Global Niagara database along with control applications, energy and performance analytics, and reporting tools. MSIs and controls contractors design workflows such that any valuable data can travel seamlessly and securely from edge devices through secure JACE boxes, protocol converters and other approved gateways; to the Ford Office Automation Network or, for some manufacturing facilities, to the Ford Manufacturing Plant Network; and up into this Ford Land private virtual LAN.

Effective semantic tagging helps to ensure that this data arrives clean, structured and ready to be processed by Ford Land's selection of applications and tools that will turn it from raw bytes into insightful information that will help us better manage our properties.

Tim Scott is a building management system subject matter expert at Ford Land. He oversees the facility operation of the Tridium Niagara infrastructure at Ford Motor Company globally. He graduated with a B.S. in HVACR Engineering Technology from Ferris State University..

What Project Haystack Means to Building Owners and Operators

LYNXSPRING®

"How well and how quickly building owners and operators can put data to work for themselves is dependent on how well their data is organized."

We have all seen and experienced the dramatic shift in building management systems over the last 10 years. This rapid evolution has set the stage on how buildings and facility are managed and how services are now delivered today.

With this shift has come the importance of data to the operation and performance of buildings and creating occupant experiences. Data accessibility and making use of it is no longer a nice to have; but rather, a must. Data impacts every facility in many ways, regardless of type or size. Data has opened up a whole new world of possibilities, and continues to change how we do business, inside and out.

The use of data is an incredible opportunity for building owners and operators. How well and how quickly your business can put data to work for you is dependent on the capabilities of how your data is organized. We can access and retrieve data; we can normalize it; we can tag it and make it easier to use across applications of all types through the standardization of semantic data models and web services (thank you Project Haystack); we can process it at the Edge, in the Fog, at the Enterprise and in the Cloud.

Haystack creates and supports many operational value points and contributes to a number of business outcomes.

Here are some:

- Contributes to a foundation for a good data management plan
- Provides a common "linkage" for data to be tied together, exchanged and managed consistently
- Distills down the type of data that is needed and required
- Ensures all users are working with the same versions of the organization's data
- Enables faster time to experience results from data (Reduces amount of non-productive data preparation work)
- Addresses the challenge of data interoperability and the use with multiple software tools and applications
- Enables effective assessment of data quality, consistency and accuracy
- Creates better understanding of information generated by all building equipment
- Provides data uniformity and supports a common schema between traditional tools and formats for data exchange
- Provides for an open, industry-standard methodology-works with any system
- Serves tagging in accordance with agreed-upon industry definitions and names found in all building asset classes

- Streamlines process of delivering the right data to the right people at the right time
- Reduces overall cost of data and ongoing "data maintenance"
- Improves the cost-effectiveness of performing analytics and supports deriving greater value from operational data
- Proven, deployed and validated in millions of square feet and buildings across different parts of the world

Data in the built environment is evolving faster than ever before. It is now an irreplaceable asset within the built environment. Building operational performance outcomes will only be realized if organizations can understand and leverage the data they produce. Project Haystack speeds up time to building operational insights and aids better decision making and drives outcomes. Project Haystack is not only changing the way buildings operate, but also how stakeholders co-operate.

"Data in the built environment is evolving faster than ever before. It is now an irreplaceable asset within the built environment."

Marc Petock is Executive Secretary on the Board of Project Haystack and Chief Marketing & Communications Officer at Lynxspring, Inc., a Founding Member of Project Haystack and leading developer and manufacturer of smart building technologies and solutions.

An Introduction to Haystack 4

Haystack4

The understanding of the need for semantic modeling of device and equipment data has matured significantly in the last decade and the requirements and techniques for applying semantic modeling to equipment data are advancing rapidly. As we have learned, semantic modeling is critical for humans to work with and understand the ever-increasing amount of data coming from their systems, but the process of manually applying that semantic model is not scalable. We need our tools to simplify and automate how the semantic model is applied.

Haystack 4 builds on the 8 years of experience in applying Haystack across thousands of buildings worldwide, the input from practitioners in the community throughout that time, as well the collaborators that have participated in the activities of Haystack Working Group 551 over the past year. The way you configure tags today using Haystack will not change, but as you will see, the way the tags get modeled within Haystack systems will enable the tools you use to become smarter so you spend less time manually configuring tags and more time getting value out of the raw data coming from your IoT devices.

Starting at the Beginning – A Standard Vocabulary to Describe Facts About Our Data

Haystack 1 pioneered the concept of applying semantic modeling to equipment and devices data. It employed a simple approach of applying tags to items to define what they "meant". Tags described things like units of measure, as well as facts and characteristics about data. For example, the tags:

```
discharge, air, temp, sensor, point,
unit:"°F"
```

tells us that a number represents a numerical value of discharge air temperature expressed in degrees F produced by a sensor. Depending on the system, this number could be named AO_ 21, nvoTemp, or Register_43015. Without the simple tags mentioned above you couldn't do much with the sensor data.

Haystack 1 therefore provided us with a standardized vocabulary to markup "things" and the data they produced. Starting in 2011 that provided the industry with its first widely adopted solution for standardized, open, data modeling for device and equipment data, which allowed us to agree on the terms to use to help define what things are. In the world of semantics that's called a vocabulary.

Haystack 2 introduced a REST API, in 2013, to provide a standard way to query a system that applied the Haystack semantic model to its data. As the demand for open protocols and open systems in the built environment continued to rise, offering an open API was important to ensure customers had a standard way to easily access the data in their systems.

Haystack 3, released in 2016, added several new data types to help machines better understand and process the different types of data formats for the IoT. The importance of data types for machines can be thought of by using a simple example. Imagine getting added to an email chain, where some of the older messages were in a language you didn't understand. You could copy and paste the text into an online translation tool and may be able to figure out what was said, but not as quickly or as easily, if you natively spoke that language. Every machine uses the concept of a String. What is stored in that String may or may not make sense without extra processing. Having to account for these different scenarios adds complexity to systems that can cause implementation problems as networks grow. Therefore, having standard data types reduces the number of scenarios our systems need to support.

Haystack 4 – Bringing a Formal Taxonomy and Ontology Model

With Haystack 4 we undertook addressing the next level of sophistication in semantic modeling – developing a taxonomy and an ontology to support the ability to represent machine-readable relationships of things and their data.

By Taxonomy we refer to a way of defining the relationships of things. For example, we say that water is a subtype of liquid because it is a specific type of liquid. The converse is that liquid is a super-type of water. Haystack 4 utilizes the concept of subtypes to organize all terms into a tree-based taxonomy. This provides us with defined and agreed upon relationships of things. We will touch on the concept of "types" more in a moment.

By Ontology we refer to the way a semantic model captures relationships between things, such as which AHU feeds air to a VAV. We need a structured taxonomy to achieve the benefits of a rich ontology of devices and equipment systems. A powerful use case for analyzing data from the IoT, is tracking the flow of energy across systems. The energy could be used to convey heated or cooled gas through a duct or liquid through a pipe, but without a standard way of representing the flow of energy, or any relationships between things, we can't drive the industry forward by making our tools more capable of automatically analyzing these relationships. Haystack 4 extends the standard to support the implementation of both a taxonomy and the resulting ontologies.

What Happens to Pre-Haystack 4 Systems?

It is worth repeating, the way you configure tags today does not change. Also, the tags you are using today are not changing. The important difference and significant benefit, is how a Haystack system models those tags internally.

Today you may have an AHU with at least the following tags:

```
ahu, equip, hvac, siteRef:…
```

That AHU likely has many points including a Discharge Temperature Sensor with at least the following tags:

```
discharge, air, temp, sensor, point,
unit:"°F", equipRef:…, siteRef:…
```

The system you used to configure those tags probably didn't help by automatically adding a discharge temp sensor point (as well as several other points) after you added a piece of equipment with the ahu tag. You either manually created a template for yourself, or if you were lucky, the system you used had a proprietary template that helped. The terms we use to represent tags are not changing in Haystack 4, but the way the tags are internally represented gives us a standard way to identify required relationships. Keep reading for an overview of how the new **def** system makes this possible.

Examples of Working with Haystack 4

This next generation of Haystack moves us closer to transforming device data into knowledge along with these important benefits:

- Enables us to infer relationships between items and apply more powerful filters and queries
- Enables the development of more advanced tools for automating semantic tagging
- Validation of tagging through certification testing
- New standard model to describe the spaces, floors, rooms, zones, etc. within a building

Defining Types of Equipment and Other Things

A key feature of Haystack 4 that enables more comprehensive modeling of taxonomies and ontologies is a simple, flexible and elegant approach to defining types of equipment and devices. Prior to Haystack 4, there was a flat list of terms that represented tags. These terms were used the same way as a hashtag you see on any social media, such as #iot or #data, as a way to quickly find content related to those topics. The challenge, which you probably have noticed, is that the number of tags seems to continue to grow as subtopics of a main topic become popular. Now you see a message with like 30 hashtags at the bottom to try and guarantee that more people will find the message. As more hashtags get added, less relevant content comes up in searches.

The new **def** model in Haystack 4 allows a term to be represented with significantly more context rather than just a term from a flat list. That additional context simplifies the number of tags needed, but more importantly allows a system to automatically determine relationships. Prior to Haystack 4, the only guaranteed way for you to know a relationship between two tags was to go and read the tag list on www.project-haystack.org. Now, a Haystack 4 system can use the metadata included in a **def** for the concept, generally referred to as "subtyping", to automatically organize virtually any "thing" or entity described with Haystack terms into a taxonomy tree. Here's a simple example defining that water is a type of liquid:

```
def: ^water
doc:Water in its liquid form
is:[^liquid]
```

Now, when you are modeling a water system, even if only the water tag is applied, an analytical algorithm that generally applies to all forms of liquid can be used because a Haystack 4 system can infer from the water **def** that it is also liquid.

Now let's examine a few parts of the new ahu **def** below (see the full **def** on the Project Haystack website):

```
def:^ahu
coolingProcess:^coolingProcessType
cools:^air
dehumidifies:^air
doc:Air handling unit - mixes outside air and
return air
ductConfig:^ductConfigType
heatingProcess:^heatingProcessType
heats:^air
humidifies:^air
is:[^airHandlingEquip]
ventilates:^air
wikipedia:`https://en.wikipedia.org/wiki/
Air_handler`
---
def:^heatingProcess
doc:Process used to heat a substance
is:[^choice]
of:^heatingProcessType
---
def: ^gasHeating
is: ^heatingProcessType
processUses: ^naturalGas
doc: "Heating by the combustion of natural
gas"
```

The ahu **def** specifies several features of AHUs including that they are a subtype of equip which cool, heat, humidify, dehumidify, and ventilate air. We also see that AHUs have a heating and cooling process. Heating process is a choice which is selected for AHU instances. The ontology defines the valid choices for heatingProcessType which includes gasHeating, hotWaterHeating, elecHeating and steamHeating. In the example above we show the definition for gasHeating, which in turn tells us that this heating processes uses naturalGas as a fuel.

To further explain our example earlier where Haystack 4 systems can now automatically infer that an ahu needs a discharge air temperature sensor, let's trace the discharge-duct **def**. Notice that the discharge-duct **def** is contained by airHandlingEquip and an ahu is a subtype airHandlingEquip (as you can see above). Therefore, we know that an ahu contains a discharge-duct. Because a discharge-duct conveys air, and we know air is a subtype of substance (air is a gas, which is a fluid, which is a substance) as defined by the Haystack 4 taxonomy, Temperature is a quantity of a substance, as you can see in the temp **def** so we can infer that anything with a discharge duct, which conveys air, needs a sensor to measure the airs temperature. Hopefully this helps you start to understand the power that Haystack 4 will bring to the IoT.

```
def:^discharge-duct
containedBy:[^airHandlingEquip,^airTerminalU
nit]
conveys:^air
doc:Supply air discharged from equipment
ductDeck:^ductDeckType
ductSection:^discharge
is:[^duct]
wikipedia:`https://en.wikipedia.org/wiki/
Duct_(flow)`
---
def:^air
doc:The mixture of gases which surrounds the
earth
is:[^gas]
wikipedia:`https://en.wikipedia.org/wiki/
Atmosphere_of_Earth`
---
def:^temp
doc:Temperature - measure of hot and cold
is:[^quantity]
prefUnit:["°C","°F"]
quantityOf:^substance
```

The Subtyping concept is a powerful knowledge modeling tool that helps us apply semantics to devices and equipment systems encountered in the real world. While it may initially appear complicated, please remember that ultimately humans are not supposed to be manually traversing the Taxonomy and Ontologies defined in Haystack 4. Systems will be doing the heavy lifting automatically allowing us to focus on the things humans are better at. One of the most important benefits is that it adds important new capabilities to define relationships that model how spaces, equip, points, and processes are related to each other. Almost all Haystack implementations will model the containment of physical spaces and equipment. It is also typical to model the flows of energy and substances such as electricity, air, and water. Subtyping and Relationship modeling provide the ability to implement advanced "filters" for querying relationships among items.

Where Are These Things in My Building?

Something people have been asking for within the Haystack Community is a model for spaces, floors, rooms, zones, etc., and Haystack 4 makes that a reality. The physical location of equipment is beneficial when automatically generating support tickets for maintenance based on the results of automated analytics. This standard model will also be a powerful data point when considering comfort within buildings. Now, you can quickly determine which equipment impacts the largest amount of space in a building or identify specific spaces, such as a data center, which should be exempt from automated demand response actions. As you are reading this, I am sure you can think of additional ways this new space model will bring value to the Haystack Community.

Take a look at an example of how the new model works below. We start with a space def, which has several subtypes, including floor, room, zone-space and others. The area def has been added as a tag on the space def so each subtype of space inherits the area tag. The full space def model can be reviewed here: https://project-haystack.dev/doc/lib-phIoT/space

```
def:^space
containedBy:[^site,^space]
doc:Space is a three-dimensional volume in the built environment
is:[^entity]
---
def:^area
doc:Area of a shape or floor space
is:[^number]
prefUnit:[“ft²”,”m²”]
tagOn:^space
---
def:^room
containedBy:[^site,^space]
doc:Enclosed room of a building
is:[^space]
```

Support for RDF/Linked Data

Another major new feature of Haystack 4 is support for RDF. RDF (Resource Description Framework) is a family of World Wide Web Consortium (W3C) specifications originally designed as a metadata data model. It has come to be used as a general method for conceptual description or modeling of information that is implemented in web resources, using a variety of syntax notations and data serialization formats. [Wikipedia]. It provides ability to accomplish semantic modelling but is typically the domain of software engineers and is not very accessible by industry professionals dealing with the real world of equipment systems.

When originally creating the Haystack approach, we felt it was important that modeling of equipment systems not require users to understand advanced software modeling concepts as a starting point. Rather, Haystack allowed users to focus on the more tangible "facts" or descriptors about data and equipment that they readily understood. Haystack took a much simpler and more accessible approach to enable industry practitioners to describe the characteristics of the equipment systems and devices that they encountered by adding simple descriptive tags. This made Haystack very accessible to industry professionals and has been a key reason for its widespread adoption and success.

As the understanding of semantic modeling of device and equipment data has matured, users are seeing interest in taking advantage the techniques and capabilities available with RDF. Haystack 4 has been designed to provide high-fidelity RDF expression of Haystack models. This allows software developers to utilize Haystack with the RDF techniques and semantic modeling tools they may be familiar with, without losing compatibility with the tag-oriented approach typically used by industry practitioners and tools designed for use by the people in the field.

Brian Frank is President and Co-Founder of SkyFoundry, a software company specializing in storage, analysis, and visualization of data from the IoT. He also serves as the technical lead for Project Haystack, working with the Project Haystack community to curate domain models and technical specifications.

Rules, Machine Learning, and Processes for Haystack Tagging

opensource STRATEGIES

"Rules-based tagging requires programmers to set up and maintain rules. Simple rules could be implemented by anybody with programming experience. Deductive rules engines would require someone with experience in logical programming. So it's natural to wonder -- can computers learn to tag data points automatically?"

We all agree that tagging building data is a good idea. Once data points from your Building Automation System (BAS) are properly tagged, you'll be able to visualize them on beautiful dashboards, run automated fault detection, and optimize building energy use in real time.

But getting there might mean sifting through thousands of cryptic data points set up over the last 10, even 20 or 30 years, all by different people. This could take weeks for just one building. With all the new technologies out there, are there better ways to do this? Let's take a look at the different ways of automating data points tagging and the processes required to put them into place.

Can Machines Follow Rules?

The first and most obvious way to teach a computer how to tag data points would be to set up simple rules like this:

```
If point name contains "Set" or "SP"
Then tag point as his, temp, set, point
```

But what about "Settings" in the BAS? We could just add to the rule to avoid confusion:

```
If point name contains "Set" or "SP"
and point name does not contain "Setting"
Then tag point as his, temp, set, point
```

But what if somebody misspelled "Setting" as "Seting", or just used "Set." as a shorthand for Setting? A better way to do it would be to use the fact that temperature should have values between 60 to 80:

```
If point name contains "Set" or "SP"
and point value is greater than or equal to 60
and point value is less than or equal to 80
Then tag point as his, temp, set, point
```

Looks nice, right? Programmers among us, though, would immediately notice that this could create a lot of repetitive code, which would be hard to maintain. For example, if there's a building where the temperature is set above 80 or below 60, we'd have to change every rule for temperature which would be a lot of work given how many temperature values come out of a BAS.

So wouldn't it be better if we had two rules?

One rule to determine if a point is a temperature:

```
If point value is greater than or equal to 60
and point value is less than or equal to 80
Then tag point as his, temp, point
```

Then a second rule to determine if a temperature is a set point:

```
If point has tag temp
and point name contains "Set" or "SP"
Then tag point as set
```

These rules introduce a subtle but potentially serious problem of their own. The sequence of when each rule is run, now matters. For example, if the second rule is run before the first rule, then the data point would not be tagged as a set point. In general, more complex rules require deductive rules engines, which would continuously run all the rules until no new tags are added to the data points.

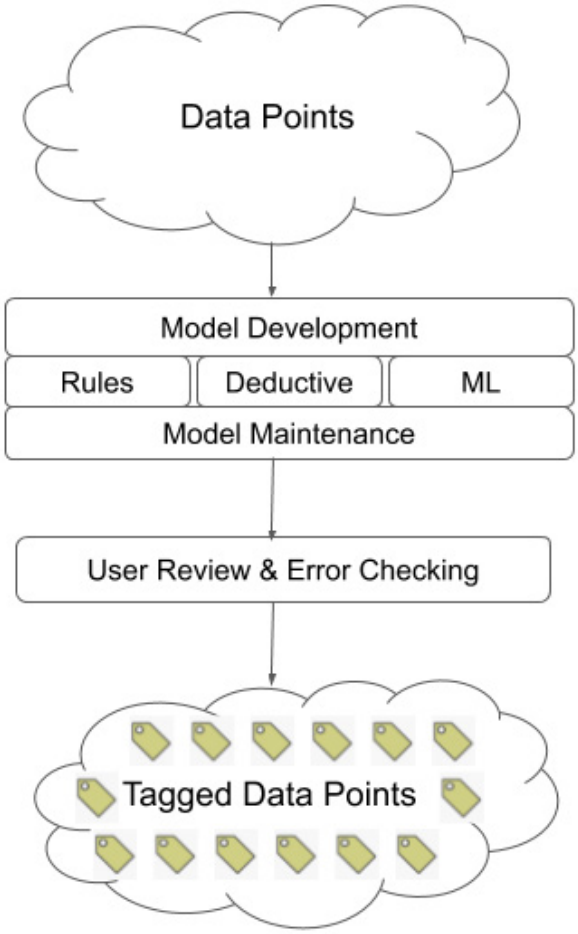

Can Machines Learn?

Rules-based tagging requires programmers to set up and maintain rules. Simple rules could be implemented by anybody with programming experience. Deductive rules engines would require someone with experience in logical programming. So it's natural to wonder -- can computers learn to tag data points automatically?

Indeed they can, with Machine Learning (ML). The process looks like this. Start with a set of data that has already been tagged correctly. Subdivide this data into a training data set and one or more testing data sets. Then feed the training data set into an ML algorithm to get a classification model. Finally run the classification model on the testing data sets you've set aside to see how good your training results are.

A recent article shows that if you start with 100 to 150 correctly tagged data points, a model could usually be trained to classify the rest of the building's points with high accuracy. The problem though, is that models trained by one building or operator, may not do well when used to tag data points for other buildings or operators. Furthermore, building ML models requires specialized skills, especially with selecting the right training data and analyzing the models' robustness. In other words, machines are really like people. It's hard to teach them and they don't know what they don't know, when they hit the real world.

Can We Trust Machines?

If you go to a website, search for cats and once in a while see a puppy, you might not be too upset. But if you mis-classify a set point as a "setting" and turn the thermostat up to 100 or down to zero, the consequences could be serious. So, however you tag your data points, manually with rules or ML models, processes will be needed to check the results and correct the inevitable errors.

Such a process should allow users to preview the tags before they are applied to data points. It should also allow a user to export data points with all their tags, so the user could check them, and then upload them again to fix any errors. Since most users are used to working with spreadsheets, exporting and importing CSV files with tags would be a good way to support this.

Do Crowds Scale?

Another feature users often ask for is to be able to export and import rules, so they could use them at different sites. The same feature would allow different users to share tagging rules, so a natural question, in this age of social media, is: Can we crowdsource rules? What if different users shared or open-sourced their tagging rules, and then rated them?

Or, maybe you don't believe in all that happy community stuff and would prefer to buy tagging rules from a company that specializes in developing them. For example, such a company could make different rules or models by OEM vendor and offer periodic updates as its rules are used at more buildings.

Both crowdsourcing and commercial development of tagging rules could work if they could scale the same rules across a large number of buildings. At the same time, they face the same fundamental problem as ML: Is there enough similarity in data points across buildings and operators? Or would there be so much differences between each building or each operator that you could never successfully train an ML model, crowdsource rules, or profitably support a commercial business?

Metrics for Evaluating Tagging Processes

So in the end, which way is the best? That depends on how the technologies would be integrated into a full process. These metrics should be used to evaluate any process being considered:

Accuracy - There are several standards of accuracy that must be met by any tagging process. First, all the points we need must be identified and tagged. Second, critical points should be individually checked to make sure they are tagged correctly. Finally, a random selection of points should be checked to determine the overall level of accuracy.

Total Cost - Just like an energy efficiency project, rules and ML models save time but also require time to develop and maintain. And just as energy prices vary, the value of time saved vs time spent may be different because of differences in skills involved. Therefore, it is important to compare different processes in terms of both the amount of time and value of the time, based on the skills involved.

Perceived Value - Users can usually accept their own mistakes if they create a few rules, but they would blame large rule sets or ML models if there are errors in them. Therefore, more complex technologies must result in significantly greater efficiencies and must offer users ways to understand and manage the results if they're to be successful with real world users.

Conclusion

Rules. Machine Learning. Crowdsourcing. They are all ways to make our lives better by eliminating the repetitive work of tagging data. In the end though, they're all the same. Trying to scale a set of rules, whether simple, deductive or neural, across a large number of buildings, OEM manufacturers and human operators. Will it work? That ultimately depends on the quality of all the data points naming conventions and data out there.

Meanwhile, no matter what process we follow, people will have to keep a watchful eye on what machines do. Our software must also support the people in their roles if they're to be successful in the long run.

To learn more, visit: www.opentaps.org.

Si Chen is the Product Manager of the open source opentaps Smart Energy Application Suite (SEAS), which helps building owners analyze energy data, identify efficiency opportunities, and finance energy projects. This article is based on his work for the buildings of a major East Coast municipal government.

Overhauling the Chiller Model in Project Haystack

ConserveIt

"Many of the principles of Project Haystack offer significant improvements to how we model machines, yet we found the current model to be somewhat limited and also missing some critical details that open up a variety of data and applications."

At Conserve It, our daily development work includes the software modeling of chillers and chiller plants for application across HVAC systems. During the 10-plus years that we've been doing this within our commercial product PlantPRO, we have continuously learned a few key things.

Recently, we had cause to review the way we model machines in plants and in doing so, we took the time to compare our models with those that have been accepted in Project Haystack.

Many of the principles of Project Haystack offer significant improvements to how we model machines, yet we found the current model to be somewhat limited and also missing some critical details that open up a variety of data and applications.

Starting From the Basics

If we look at a basic chiller, it is essentially a number of pieces of equipment that are related to each other. Typically, a simple chiller will have two heat exchangers, a refrigerant circuit and a compressor. Whilst it is possible to build a specific model for this easily, it would be better to create a more generic model that can capture more complicated scenarios, such as multi-circuit and multi-compressor machines. Any new model that captures more complicated chillers should also take into account heat pump and dual production machines (producing chilled and hot water simultaneously), which is something we see on a daily basis.

If we look at a simple chiller example then the machine could be tagged, without worrying about data point entities yet, as follows:

```
// chiller or plant machine entity
id:@a, hvac, equip, chiller

// first heat exchanger, also referred to the evaporator
id:@b, exchanger, function:cooling, water, equip, equipRef:@a

// second exchanger, also referred to as the condenser
id:@c, exchanger, function:sink, water, equip, equipRef:@a

// refrigerant circuit
id:@d, equip, circuit, equipRef:@a, refrig, gas

// a single compressor
id:@e, equip, comp:1, screw, equipRef:@a, circuitRef:@d
```

Let's examine this model more closely.

A few new tags have been introduced, however what has been created is a more detailed model of a chiller. One that captures the relationships of all the "sub-equips" in a chiller. Creating higher level abstractions in the model allows us to capture more complex scenarios and even add value in these simpler cases.

The first new tag is *exchanger*. This identifies an equip as a typical heat exchanger vessel, or perhaps energy exchanger would be a better term, that would be found in a chiller or plant machine.

In this model we are using the *equipRef* tag as the means of relating all the sub-equips together to the parent *equip*.

The *function* tag is a key tag that provides for flexibility in other scenarios. This tag is an enumeration and could have the following values:

- *cooling* - the exchanger produces cooling water
- *heating* - the exchanger produces hot water
- *sink* - the exchanger rejects energy, typically heat, should be used in conjunction with either an air or water tag to denote the rejection medium
- *source* - the exchanger imports energy from either water or air, should be used in conjunction with an air or water tag to denote the energy source
- *source_sink* - the exchanger can either be importing energy or rejecting energy to a medium, it cannot do both simultaneously, should be used in conjunction with an air or water tag to denote the energy source and rejection medium

The circuit tag marks a sub-equip in the plant machine as a loop that conveys a fluid or gas of some description. When used with the *refrig* and *gas* tags it denotes this sub-equip as a refrigerant circuit in the chiller.

The *comp* tag creates another *equip* entity within our chiller that represents a compressor. The compressor has a tag denoting what compressor type it is, in this case screw, and this is done on the compressor as it is possible for plant machines to use different compressor types on the one machine.

Lastly, I have created the tag *circuitRef*, which is a most important tag especially when looking at more complicated machines. This tag allows us to know the circuit a compressor is connected to in a chiller. This information is absolutely crucial in multi-circuit/multi-compressor machines as it allows us to do much more in-depth analysis of the operation of a machine.

So far so good. we have constructed a new model for a machine. Now, what about points?

Exchanger Points

As the model for a plant machine has now been generalised a lot more, it is also possible to generalise the data point entities for an exchanger and other sub-equips. For an exchanger we would alter the existing point definitions for the chiller model to the following:

```
entering, water, temp, sensor
leaving, water, temp, sensor
leaving, water, flow, sp
water, flow, sensor
water, delta, pressure, sensor
water, valve, isolation, cmd
```

The existing chiller model defines many more points which in practice are superfluous. For example, in the hundreds of plants where we have deployed our commercial PlantPRO solution on, we have never seen an entering and leaving flow meter or sensor on the evaporator of a chiller. The cost of doing such an installation would be beyond most facilities and given that the flow through a chiller is generally the same at the entering and leaving points to the exchanger, then it need only be measured once.

Circuit Points

For the refrigerant circuit we need to introduce some new point definitions. Luckily there are not so many to do, but there are still some complicating factors that need attention. For a refrigerant circuit we would look to have entities such as:

```
discharge, pressure, refrig, gas, sensor
discharge, temp, refrig, gas, sensor
suction, pressure, refrig, gas, sensor
suction, temp, refrig, gas, sensor
```

They look all good and they capture some vital information. The refrigerant temperature sensors may not always be present so they are optional, and it is possible to convert pressure to temperature via a calculation as well.

The complication on a refrigerant circuit comes in to play when there are multiple compressors connected to a single refrigerant circuit. Each compressor may have its own suction and discharge pressure sensors rather than there being 1 suction and 1 discharge pressure sensor for the whole circuit, but in reality, most applications only care about the suction and discharge readings for a circuit. This can be handled by knowing the relationships of compressors to circuits, which of course are now captured in this new model.

Kicking the Complexity Up a Notch

To this point we have created a simple model for a single refrigerant circuit and single compressor chiller. How would this look if we had a 2 circuit 4 compressor machine?

```
// chiller or plant machine entity
id:@a, hvac, equip, chiller

// first heat exchanger, also referred to the evaporator
id:@b, exchanger, function:cooling, water, equip, equipRef:@a

// second exchanger, also referred to as the condenser
id:@c, exchanger, function:sink, water, equip, equipRef:@a

// refrigerant circuits
id:@d, equip, circuit, equipRef:@a, refrig, gas
id:@e, equip, circuit, equipRef:@a, refrig, gas

// compressors connected to circuit @d
id:@f, equip, comp:1, screw, equipRef:@a, circuitRef:@d
id:@g, equip, comp:2, screw, equipRef:@a, circuitRef:@d

// compressors connected to circuit @e
id:@h, equip, comp:3, screw, equipRef:@a, circuitRef:@e
id:@j, equip, comp:4, screw, equipRef:@a, circuitRef:@e
```

By constructing the model in this way, it is now possible to do some really in-depth analysis of the machine, particularly on the refrigerant circuits. By creating a relationship between compressors and circuits, it is now possible to determine which circuit is active - a vital piece of information that we use in PlantPRO when conducting performance analytics on a chiller.

Ready to Go Even Further?

In our work, we see a lot of plants with chillers and heat pumps, or chiller and heat recovery machines of various kinds. How could we model these? Pretty easily now with our much more generic and abstracted model. See below for a single compressor and single refrigeration circuit air source heat pump:

```
// heat pump or plant machine entity
id:@a, hvac, equip, heatpump

// first heat exchanger
id:@b, exchanger, function:heating, water, equip, equipRef:@a

// second exchanger, also referred to as the condenser
id:@c, exchanger, function:source, air, equip, equipRef:@a

// refrigerant circuit
id:@d, equip, circuit, equipRef:@a, refrig, gas

// a single compressor
id:@e, equip, comp:1, screw, equipRef:@a, circuitRef:@d
```

That is a simple heat pump example. There are many more complex heat pumps as many manufacturers offer either water reversible or refrigerant reversible heat pumps which means the function of an exchanger is a dynamic setting and changes periodically depending on ambient temperature or time of year. This introduces further complexity into the model, but they are not insurmountable now that we can label an exchanger as having a function.

Heat recovery chillers are machines that can only operate when they can reject all of the heat they pull out of the cooling water into a heating loop in a building. This model is now easy to construct using the new ideas presented here.

```
// heat recovery plant machine entity
id:@a, hvac, equip, recovery, chiller

// first heat exchanger, also referred to the evaporator
id:@b, exchanger, function:cooling, water, equip, equipRef:@a

// second exchanger, also referred to as the
```

```
condenser
id:@c, exchanger, function:heating, water,
equip, equipRef:@a

// refrigerant circuit
id:@d, equip, circuit, equipRef:@a, refrig,
gas

// a single compressor
id:@e, equip, comp:1, screw, equipRef:@a,
circuitRef:@d
```

Simply by changing the function of an exchanger a new type of machine can be created.

What we have been able to illustrate in this discussion is that there are alternative ways to model plant machines than what is currently in the Project Haystack standard. Furthermore, the standard can be enhanced to incorporate the ability to model different types of plant machines. With the "electrification" of our building being driven by utilities around the globe, and the drive to move away from fossil fuels growing ever larger, our community will need to know how to work with heat recovery machines of various kinds so they can provide meaningful applications to stakeholders and customers.

In practice, through our work with PlantPRO, we actually have the requirements to model much more complicated machines from a number of different manufacturers. Machines with three exchangers whose operating mode can vary dynamically, modular machines where each module has multiple compressors and circuits and the function of an exchanger can also change dynamically, but for different reasons.

At Conserve It, we have had to find new ways such as these, to model such highly complex machines in order to be able to understand the data they generate.

If anyone has any interest in discussing this further or even taking it into a working group, I am happy to continue to work on this with the community. At the end of the day we need to find a way to model chillers and chiller plants more effectively and we are sure that others in our community also have this need.

As the Chief Software Architect at Conserve It, Richard McElhinney manages and drives the development of their industry-leading chiller plant optimization technology and associated products. He also serves as Vice President on the Project Haystack Board of Directors. .

Callum Rosel is an experienced Software Development Manager at Conserve It, who enjoys helping others expand their knowledge and understanding of the IoT world.

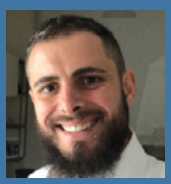

Filippo Bernardello is a Plant Software Engineer and the EMEA Product Manager at Conserve It SRL the brand-new European branch of Conserve It. With diverse project deployment experience, Filippo is a highly skilled engineer with in-depth knowledge of smart buildings, building automation, energy efficiency and IoT.

Machine Learning to Apply Haystack Tagging at Scale

bueno
Built Environment Optimisation

How using a Machine Learning tagging bot allows Deployment teams to significantly lower the time spent tagging and increase the quality of algorithms.

At BUENO we bring in data points from a variety of building systems, apply meaning to this data using Haystack Tagging models and use these to deliver value through energy savings, smarter comfort control and data driven maintenance tasking. To make any use of data we extract, we need to understand what that point represents, and to apply our intelligence effectively, we need to understand the context behind each data point. Haystack tagging models are used to give our technology this context, allowing us to apply a common set of algorithms to many different building types.

To ensure these outputs of our technology are reliable we need to apply tagging models at a high degree of accuracy. In the not so distant past, the process of tagging a site was extremely manual and time consuming for our Deployments team, usually resulting with some human error. This meant our product was expensive to deploy to a site and our analytics were not as accurate as they could be due to incorrect tags. As our company has grown, we have been creating more complex analytics and bringing on buildings at a faster rate, meaning our previous methods of tagging were becoming unsustainable.

I recently participated in the Haystack Connect 2019 Conference which involved a roll-out of the new Haystack 4 Tagging standard, as well as technical sessions on various applications of Haystack. One pain-point multiple presenters spoke of was the labor intensity of applying tags on building systems. During the conference, I was lucky enough to present BUENO's solution to speeding up the tagging process.

Introducing Unicron

Our solution to supporting the company's growth, without exponentially scaling the deployments team, was a machine learning tagging bot called Unicron. We figured we could learn from the numerous buildings we had already poured our blood, sweat and tears into tagging as a training dataset to infer what tags to place on new deployments. The end goal being a user interface that our Deployments team can use to simply check suggestions the machine has pushed forward and accept or reject them.

To most people, Machine Learning is a buzzword that somehow "magically" spits out all the correct results. In this article, I hope to lift the hood and provide a basic explanation of how our system acquires results. In short, Unicron is constantly examining how tags are applied on our current sites and builds a model based on its learning. This model is then used to make inferences about new points that come in and provide tagging suggestions to a Deployment Engineer.

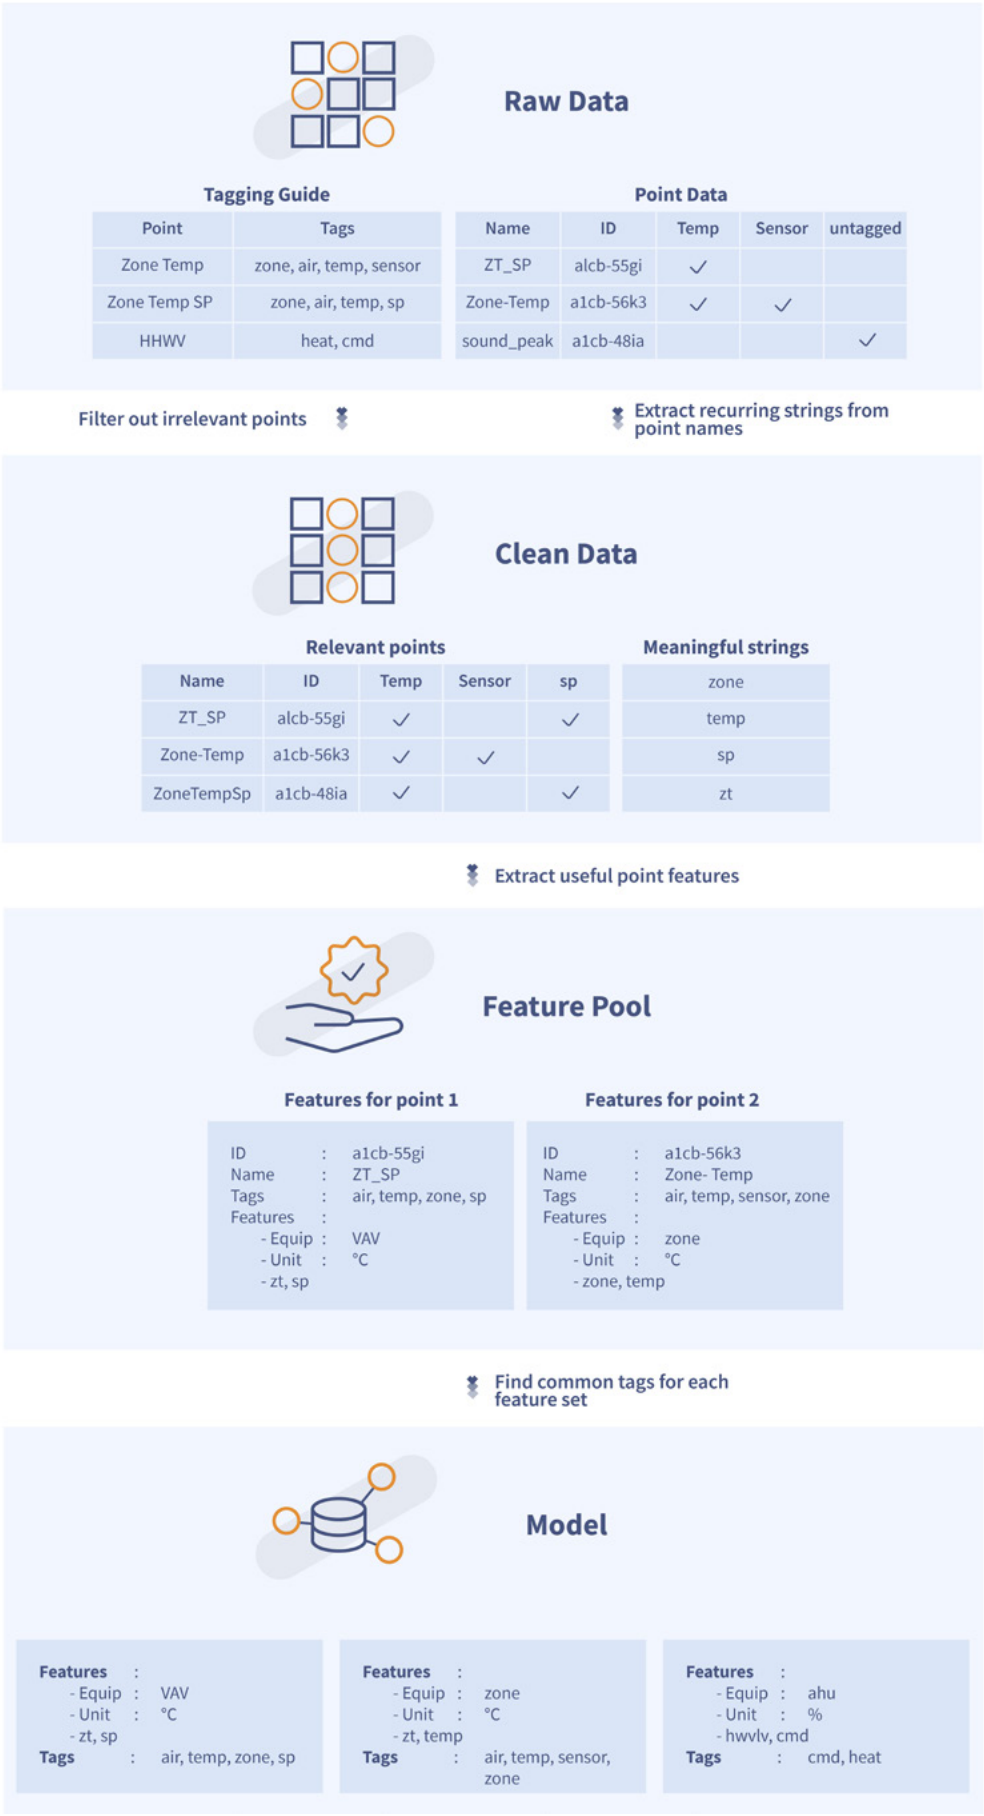

Figure 1. Basic information flow for Unicron learning cycle.

As ***Figure 1*** shows, the model is built by reading the BUENO tagging guide and all tagged data points we have in our system. The tagging guide is basically a big dictionary that tells our Deployments team what tags need to be placed on what points, it is used in Unicron to inform the algorithm what tags are relevant. An example of the raw point data is shown in the figure, these tagged points are what the algorithm uses as a training set for the model.

The raw data is cleansed, meaning all irrelevant data points and tags are filtered out until only the useful information is left, this will include point names, the units, a reference to the equipment it sits under and its tags. The clean data also includes a long list of meaningful strings. The meaningful strings will be any string combinations that appear frequently in the names of the raw data points.

Using the relevant points and meaningful strings a feature pool is built. The feature pool will have an entry for each point, consisting of useful point features and some basic information. We have defined the useful features as:

- The type of equipment the point sits under
- The unit it reads in
- The kind of data it reads in (i.e. number, boolean)
- The meaningful strings that have been matched within the name of the point

As can be referenced in Figure XX, there is also an entropy value, this is calculated based on how much of the point name has been matched to meaningful strings, we use this as a proxy for the confidence that we have completely understood what the point does.

The model building algorithm will look at the feature pool and extract all recurring pairings of feature sets and tags. The most likely pairings of features to tags will be saved to a database. This database will be used as a knowledge base to assume tags for new points.

All of the learning tasks involved here are run constantly, meaning that when updates are made to the points we currently have or a new building is deployed, the feature pool will start to update immediately. Our model will slowly start learning from the new feature sets it finds and update itself accordingly. This way, we can add new tagging models or update our tagging model to align with the new Haystack Tagging standard, without updating our tagging software.

"The unique flexibility of the Haystack methodology and tagging allows you to add uniform data descriptors to your data, no matter what their original names may have been."

How Tagging Suggestions Are Created

The flow Unicron uses to create suggestions can be seen in ***Figure 2***. We will read in all points on a site as raw data and use some of the data we have extracted during the learning stage to find the appropriate tags.

The point features Unicron uses includes the point's name, the type of equipment it lives under, the unit it reads in and the meaningful strings included in its' name. The type of equipment that the point lives under will have been tagged using scripts and hence will already be known. The point's features are extracted using exactly the same method as in the learning process.

Tagging suggestions are built by matching these features to the feature sets in the model. The confidence of a match is calculated as a mixture of: proportion of the point name that we have matched to meaningful strings (similar to the entropy mentioned previously), and the distance between the feature sets matched.

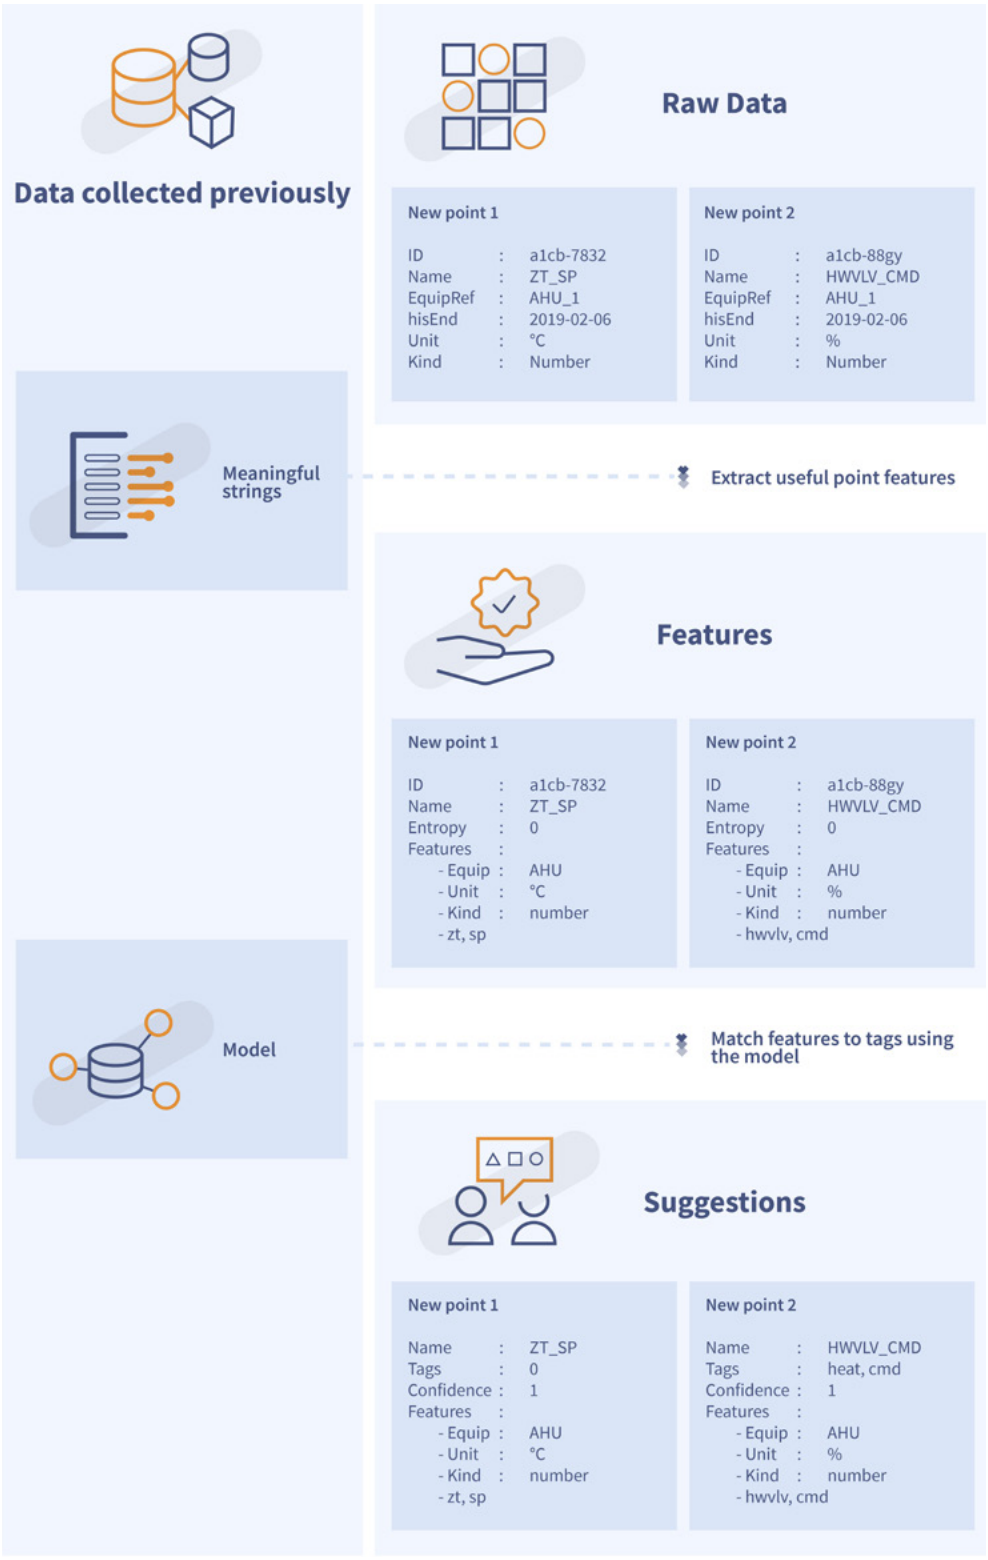

Figure 2. How suggestions for new points are created.

How We Use It To Make Tagging Faster

Accept?	Confidence	Features	Number of Points	Point Name	Suggested Tags
TRUE	1	Equip: floor Unit: unk Kind: Bool **tenant, zone, ah**	108	Zone 1_TenantAH_3	afterHours, cmd, occupied
TRUE	1	Equip: mvcd Unit: °C Kind: Number **temp, spd**	45	TempSp_D	sp, temp, zone, zonePair:D
TRUE	1	Equip: elecmeter Unit: pf Kind: Number **powerfactor**	30	Powerfactor	pf, sensor
TRUE	1	Equip: weatherSation Unit: °C Kind: Number **dewpoint**	4	dew_point	air, dew, outside, sensor
FALSE	0.85	Equip: weatherSation Unit: km/h Kind: Number **wind spd**	4	wind_spd_kmh	sensor, speed, wind

Figure 3. Example tagging suggestions Unicron creates.

When we bring on a new site, our Deployments team will connect all of a sites data points to our servers and start to read in histories. Next, they will access our bot and tell it which site they want to tag, and the bot delivers suggestions in a user interface similar to ***Figure 3***. As you can see the 'Accept?' field will default to true when there is 100% confidence, but to false when the confidence level is lower.

The Deployments team will check the suggestions and change the 'Accept?' field to 'TRUE' where necessary. The bot will use this feedback to make the appropriate tagging updates, leaving the Deployments team to only manually tag those points that the bot had incorrectly matched. On the next iteration of the learning process, the bot will use these recently tagged points on this new site as feedback and update the model with this new knowledge.

Results

On average Unicron will match 94% of a sites' points to a tag set with 100% confidence. We find that most new sites that we bring in have very similar naming patterns to those that currently exist in our system, hence direct matches to the model can often be made. The percentage of points matched can decrease significantly when the points are unusually named or the site contains equipment we don't currently have tagging models for. We also found that points matched with 100% accuracy are matched to the correct tags more than 98.1% of the time.

Clearly the algorithm is not perfect and in need of continuous improvement. In the future, we plan to improve the calculations we use for our confidence metric, add further detail to the features we extract from each point, update it to deal with different languages and apply the tagging bot at an equip level. We have already experienced some extremely valuable improvements to the time we spend tagging and the quality of the tags applied to our current stack.

The Deployments team estimates that using the bot reduces the time they spend tagging by 50%. This decrease can be attributed both to the improved User Interface that allows the Deployments team to view points with identical features in a single row, as well as the convenience of having suggested tags automatically available. This means they can spend more of their time working on tagging the complex relationships between equipment and use the bot to do the grunt work of tagging points. The ultimate result is a reduced price of entry to our platform for our clients, and greater job satisfaction for the Deployment team.

In addition to using the bot to tag all new buildings we bring onto our platform, we have used the bot to revise and update the tagging on our current buildings. After this exercise, we saw a 60% decrease in analytics bugs raised that were found to be tagging issues. Hence, Unicron has greatly increased the accuracy of our algorithms, leading to an immeasurable reduction of engineering time spent looking at faulty analytics, and best of all, higher confidence in our algorithms.

Summary

At BUENO we use Haystack tagging models to make sense of the data we ingest from buildings, but the process involved in tagging a site has historically been very manual. Using a Machine Learning tagging bot has allowed us to significantly lower the time spent tagging and increase the quality of our algorithms. I hope this article helps to make sense of the machine learning process and one way it can be applied to make tagging easier!

Lucy Kidd, a Data Scientist at Bueno Systems, is passionate about leveraging her knowledge and studies to make the world a better place. She is a developer of analytics and a tagging specialist, working towards the goal of making our buildings better!

Data Abstraction - The Coming Revolution

J2 INNOVATIONS
A Siemens Company

"None of the open-standard protocols provide sufficient metadata, and although a handful (like BACnet), provide some structure, others provide no structure at all (e.g. Modbus). This is why Haystack matters so much."

If you're reading this, you're probably already familiar with Project Haystack, and concept of "tagging" (the use of metadata to provide context and meaning to data). Recently, Project Haystack expanded the standard in its release of Haystack 4. Haystack 4 expands the methodology for data-modeling, defined by the technical words "ontologies" and "taxonomies." Ontologies are "a set of concepts and categories in a subject area or domain that shows their properties and the relations between them." A taxonomy is a way to organize things from general to specific.

What you may be less aware of, is that the evolution of Haystack is a manifestation of a broader revolution in data management that has been happening in other sectors (like Life Sciences) for many years. Only recently, has the building services sector begun to see the value and adopt this method of data management.

Why The Late Adoption?

There are several reasons why our industry has been slow to adopt:

- Building automation markets have been dominated by large system manufacturers who have little interest in providing open technologies for data exchange due to the commercial gains they've realized from the "closed" systems approach.
- Building systems have typically operated in functional "silos" without any means to easily exchange data except by conversion of one protocol to another.
- Building services engineers and consultants involved in system design are typically not academics, so tend to approach system design from a practical perspective rather than an academic one, resulting in each project being handled differently.
- The buildings systems markets are fragmented with many players, making it difficult to successfully champion a standardised approach.

Why This Is Happening Now

There are several changes in the industry that have caused this evolution. First, the quantity of data now available from building systems has multiplied, creating new challenges in the way it is analysed and reported. Building operators are realizing the value leveraging this data can bring to building assets, however, doing so is challenging due to the lack of standardized data sets.

This is compounded by the fact that there is a need for data to be more easily processed without human intervention, which would feed sophisticated software that would automatically configure and analyse system performance.

Another reason for this increasing rate of adoption is that Haystack community members like Altura, Bueno, Conserve IT, IOT Warez, J2 Innovations, Lynxspring, SkyFoundry, and Siemens are now leveraging the Haystack standard, making it easier for their products to take advantage of standardised data sets.

Other sectors, such as the life sciences, have talked in terms of data models, ontologies and taxonomies for years, as they have had to handle huge data sets and structure them in useful ways. These terms are still unfamiliar to most building services engineers, facilities managers, and others who work with the real-time data generated by modern building automation and management systems. Instead, our industry has been pre-occupied with the basic task of collecting data itself in standard ways, as the system manufacturers have slowly migrated their network protocols from proprietary ones to BACnet, Modbus, KNX and similar open standard protocols.

As open standards emerged in the 90s, many had hoped for one protocol standard to become universally accepted. Instead, multiple standards have been adopted and most people now accept that the future will include the use of several protocols. On a typical project, each building function has opted for different protocol standards (e.g. BACnet for HVAC, Modbus for electrical systems, KNX and DALI for lighting). This situation requires management software to be able to "speak" several protocol "languages" at once, which poses the question: What will be the common data abstraction to enable ALL the data to be processed in a common way?

None of the open-standard protocols provide sufficient metadata, and although a handful (like BACnet), provide some structure, others provide no structure at all (e.g. Modbus). This is why Haystack matters so much. Not only do the Haystack tags provide the missing metadata in a standard way, but the Haystack 4 now provides a better data structure. Haystack began as a dictionary of tag definitions and has since evolved to become an open standard data abstraction protocol (used over REST and MQTT). The tagging and structuring of data in a standardised way overcomes the problem of ambiguity when trying to interpret the meaning of data.

How Data Modeling Helps in Life Sciences

To understand better how data-modeling helps structure data and enable easier access to compare data sets, let's look at how it helps in life sciences. The use of ontologies and taxonomies for information retrieval by medical researchers goes back many decades. Breakthroughs and development of new technologies catalysed the rapid growth of new areas of research, much akin to the increasing prevalence of smart sensors and real-time data output now occurring in the buildings world.

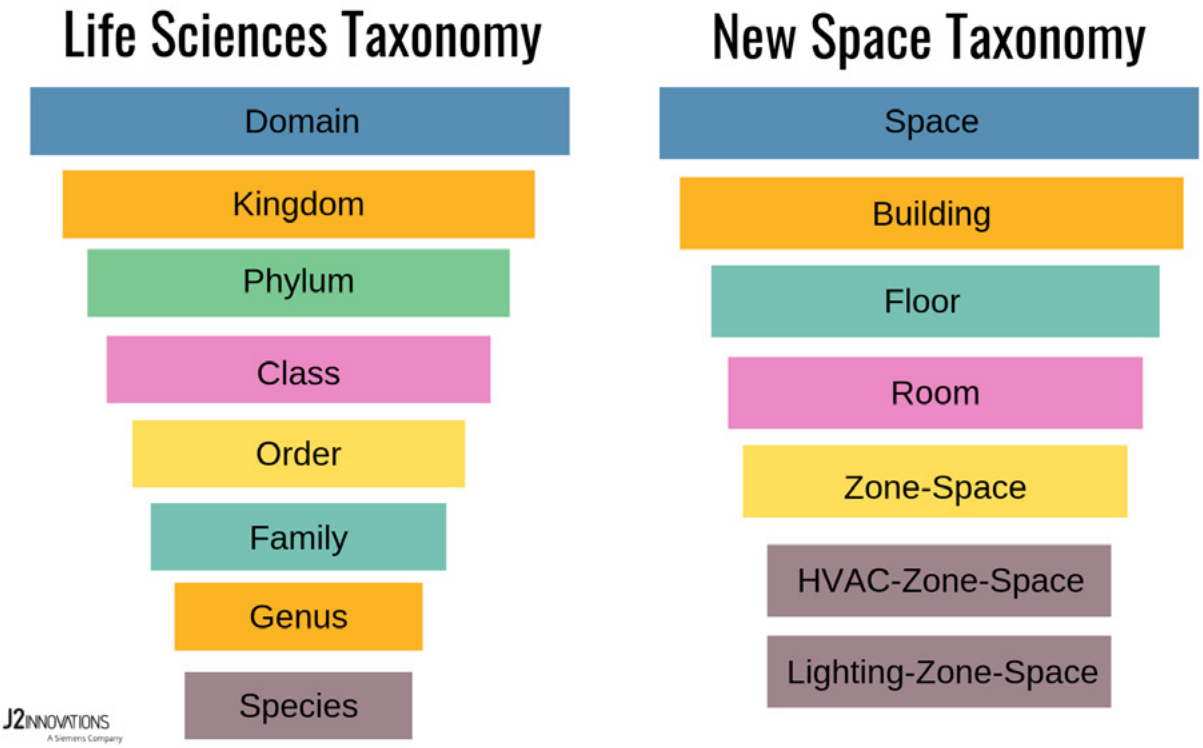

Modern medical research is increasingly heterogeneous and autonomous. With multi-disciplinary research on the rise, terminological, syntactic, and semantic differences across disciplines creates barriers for efficient information modeling and retrieval. Other contributing factors include existing differences in naming conventions, identifier codes, and schema. Applications for standardised data-modeling include information retrieval, data annotation, and data integration.

An example that illustrates the problem of data ambiguity is when searching for references to "gene". The concept of a gene is a "DNA fragment that transcribed and translated into a protein" but another definition of a gene is a "DNA region of biological interest with a name that carries a genetic trait for a phenotype". While both definitions are correct, additional factors are needed to precisely identify what is meant. The answer is to deconstruct the definition to achieve greater precision. Coding and non-coding DNA are separate notions, so further sub-classification is necessary. This is what is meant by a taxonomy – a classification scheme.

As early as the 1960's, researchers foresaw these issues and started development of Medical Subject Headings (MeSH) and later the Unified Medical Language System (UMLS). Domain experts from across disciplines worked together and reached consensus on the organization of medical concepts and the relationships between these concepts. This formal vocabulary contains different term types, including Descriptors, Qualifiers, Publication types, Geographics, and Entry terms.

MeSH descriptors, for example, are represented in a tree-like structure (a taxonomy) in 16 categories with increasing specificity up to 13 hierarchical levels. Top-level categories including anatomic terms, organisms, diseases, etc.

To facilitate a faster and more efficient user experience when searching the library database, the US National Library of Medicine employs trained indexers to read and tag thousands of articles added each month. By matching search terms to MeSH terms, search results are far more relevant because results are according to the tags and not the text.

In the building automation world, the challenge has been the creation of a single database to hold ALL the data relating to a building project. The creation of the Building Information Modeling (BIM) standard has been a step in the right direction in achieving a comprehensive 3D model of a building's structure and services. However, the multiplicity of data sets from various construction and other disciplines, has not been integrated into a single database.

Discussing this topic recently with a company involved in providing such database integration, they explained they have had to deal with nearly 100 different data models so far. As we saw in the life sciences example, this illustrates how important it is for the construction and building automation industries to reach a consensus in-order to achieve a common, standardised approach for each aspect of building projects.

Project Haystack has already made major strides in this direction, and is gaining significant traction, having already been used in over 30,000 buildings. We need to co-opt suppliers and specifiers across more of the building systems to engage in the standardisation process so that the scope of Haystack tags can be extended across more disciplines. Like the Life Sciences field did decades ago, we also need to map the Haystack definitions with other existing data-models in the building space.

Chris Erwin recently joined J2 Innovations, the developers of the FIN open framework, as V.P. of Sales for EMEA and Asia.

NIAGARA SUMMIT

APRIL 19-21

SAN DIEGO

MANCHESTER HYATT

Tridium invites you to its biennial gathering of the Niagara Community next April in San Diego, California. We are planning an agenda around the technology and business topics that are driving our Connected World. Registration opens soon. Contact us at niagarasummit@tridium.com to get on our mailing list for the event or if you are interested in exhibiting.

TRIDIUM
www.tridium.com

Contributed Article

Project Haystack Cybersecurity Working Group

INTELLIGENT BUILDINGS™

Project Haystack has been the driving force for the standardization of data models and web services. This standardization has unlocked the power of gathering a vast amount of data. The goal of the Cybersecurity Working Group (CySecWG) is to harness and capture information and configuration data to enable cybersecurity features that analyze and monitor systems.

Smart Building Cybersecurity awareness is growing. However, building control systems standards are varied, as many different types of integrators are installing the systems. Project Haystack is the leader in standardizing the building control space and cybersecurity. In early 2019, I started the Project Haystack Cybersecurity Working Group (CySecWG). As the year progressed, I was learning how to navigate the information available, with the help of the Project Haystack community.

Foundation Standards

Early on, the decision was made to build tagging based on an established methodology. That methodology is the National Institute of Standards and Technology (NIST). Formerly known as the National Bureau of Standards, NIST was founded in 1901 and is a non-regulatory federal agency within the U. S. Department of Commerce. Its mission is to promote U.S. innovation and industrial competitiveness by advancing measurement science, standards, and technology that enhance economic security and improve our quality of life. NIST is a voluntary practice for a mandatory standard within Federal agencies.

The NIST Cybersecurity Framework is a set of best practices, standards, and recommendations that help an organization improve its cybersecurity measures. The NIST Cybersecurity Framework seeks to address the lack of standards when it comes to security. There are currently major differences in the way companies are using technologies, languages and rules to fight hackers, data pirates and ransomware.

IDENTIFY	The Identify Function assists in developing an organizational understanding to managing cybersecurity risk to systems, people, assets, data and capabilities.
PROTECT	The Protect Function outlines appropriate safeguards to ensure delivery of critical infrastructure services. The Protect Function supports the ability to limit or contain the impact of a potential cybersecurity event.
DETECT	The Detect Function defines the appropriate activities to identify the occurrence of a cybersecurity event. The Detect Function enables timely discovery of cybersecurity events.
RESPOND	The Respond Function includes appropriate activities to take action regarding a detected cybersecurity incident. The Respond Function supports the ability to contain the impact of a potential cybersecurity incident.
RECOVER	The Recover Function identifies appropriate activities to maintain plans for resilience and to restore any capabilities or services that were impaired due to a cybersecurity incident. The Recover Function supports timely recovery to normal operations to reduce the impact from a cybersecurity incident.

The Working Group – CySecWG

With Project Haystack and the NIST Cybersecurity Framework, the building blocks for creating standard identification and classification of control system components are already available; they just need to be assembled. The purpose of CySecWG is to establish a tagging schema that facilitates the control system for enabling key data gathering that applies to the Functions and Categories of the NIST Cybersecurity Framework. In doing so, the data will be available to monitor and identify threats, provide data for forensics, monitor and report configuration changes, and allow for a more automated approach to measuring compliance.

CySecWG currently consists of 20 members and is still in the formative stages. Members need not be cybersecurity experts; all are welcome who have an interest in cybersecurity for control systems. Input from those who have not been involved in cybersecurity in the past will aid in the development and eventual socialization of the tag library by establishing an easily understandable product.

Objectives

Cybersecurity for control systems is still in the early stages and is best approached with a "crawl, walk, run" strategy. Even though we are in the "crawl" stage, there are many gains to be won. During this formative time, the objective is to outline basic tagging for the purpose of identifying device attributes for monitoring anomalies and configuration changes. Another tag assignment would be identifying criticality levels. For example, an attribute could be a criticality level depending on a device's objective. Two similar devices could have vastly different roles in the system. For example, one device that controls an air handling unit that supplies a common area would be considered low criticality, whereas another device of the same type that supplies a blood storage area would be considered high criticality.

There may be an opportunity to incorporate NIST privacy and security controls by tagging devices and/or objects with associated controls. The NIST privacy and security controls are listed in **Figure 1**. Subcategories listed below each control could be used to identify additional attributes.

ID	FAMILY	ID	FAMILY
AC	Access Control	MP	Media Protection
AT	Awareness and Training	PA	Privacy Authorization
AU	Audit and Accountability	PE	Physical and Environmental Protection
CA	Assessment, Authorization and Monitoring	PL	Planning
CM	Configuration Management	PM	Program Management
CP	Contingency Planning	PS	Personnel Security
IA	Identification and Authentication	RA	Risk Assessment
IP	Individual Participation	SA	System and Service Acquisition
IR	Incident Response	SC	System and Communication Protection
MA	Maintenance	SI	System and Information Integrity

Figure 1. NIST 800-53 Privacy and Security Controls

A tag could also be used for the purpose of assigning impact. 800-53 has Security Control Baselines for determining the security controls for low-impact, moderate-impact, and high-impact information systems. The three security control baselines, shown in **Figure 2**, are hierarchical in nature with regard to the security controls employed in those baselines. Impact attributes could be included in Haystack as a means for identifying the operational impact of a device.

PRIORITY CODE	SEQUENCING	ACTION
Priority Code 1 (PQ)	FIRST	Implement P1 security controls first.
Prioity Code 2 (P2)	NEXT	Implement P2 security controls after implementation of P1 controls.
Prioity Code 3 (P3)	LAST	Implement P3 security controls after implementation of P1 and P2 controls.
Unspecified Proioity Code (P0)	NONE	Security control not selected in any baseline.

Figure 2. Security Control Prioritization Codes

Summary

Work is scheduled to begin in early 2020. This work will be focused on laying a solid foundation for established guidelines. There are many more guidelines that NIST and other agencies offer for cybersecurity. CySecWG will determine which are directly applicable to the industry and begin to build a cybersecurity tagging structure to facilitate the identifying and monitoring control systems, and to lay foundational compliance attributes.

https://project-haystack.org/forum/topic/667

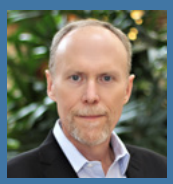

Fred Gordy, Director of Cybersecurity at Intelligent Buildings, is a Smart Building industry expert and thought leader with 20 years of experience in secure control system development and implementation for Fortune 500 companies throughout the U.S. and abroad. He is one of the first in the Smart Building industry to address the inherent risks posed by control system technology.

Contributed Article

The Haystack Byte Journey Continues

BASSC

"In the previous issue of Connections Magazine, I mentioned in my article a product called Mobilytik. Using the Haystack server and tags, we can offer two new features which are Offline Metering and Personal Feedback."

Our Haystack journey continues from where we left off. Project Haystack's primary capability is that it can act as a server and client. This feature enables creators to utilize it as a converter technology of data types from various sources such as images, personal feedback, sensors, input-output from controllers, and others.

My previous article in the Winter 2019 issue of Connections Magazine, I mentioned Project Sandstar, which is a hardware-agnostic DDC. Project Sandstar merges the Sedona framework with Project Haystack into a seamless whole.

In that article, I also mentioned a product called Mobilytik. We have improved this application, offering two new features which are Offline Metering and Personal Feedback.

A Haystack byte can be hiding in different locations and different forms. The byte could be lurking in an image or a video feed or within voice form. We can derive this byte and encapsulate it so that we can pass the data to the Haystack server.

Offline Metering

We have generated a new feature called Offline Metering. There are cases within facilities where budget constraints do not allow managers to install a digital meter that has BACnet or Modbus protocol capability.

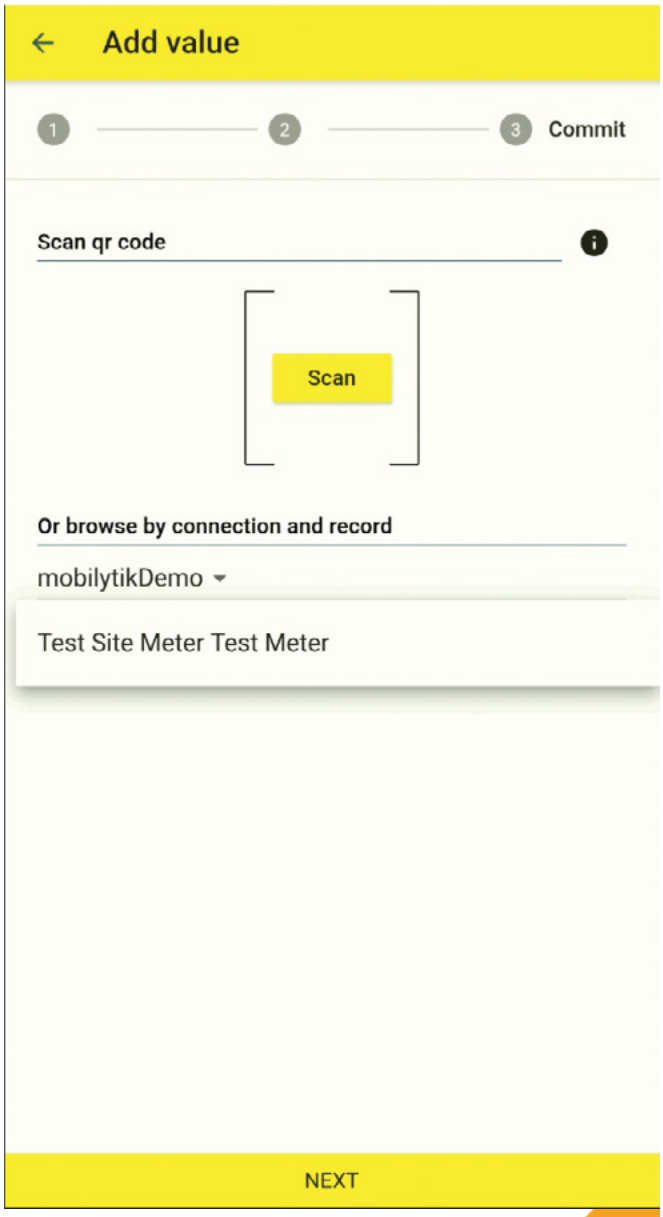

Within Mobilytik, we have generated a feature where the facility technician can derive the current meter value and upload it to the Haystack server.

The facility manager can easily find the virtual meter points based on the barcode. Mobilytik generates the barcode from Haystack's id tag. The engineer adds a simple Haystack tag called offlineMeter to the Haystack record within the Haystack server.

When the end-user logs into the Haystack server for the first time, offlineMeter points are cached. The end-user can either select the meter manually or can scan the barcode that is attached to the meter.

At this stage, we can take a picture of the analog meter. Google's tensorflow machine learning algorithm helps us to parse the image and digitize the text.

As our next step, we generate cached data to be committed to the Haystack server. The Power meter could be in a location where there is no internet connectivity. There could also be more than one meter within the basement. The system allows committing all historical data changes in one click.

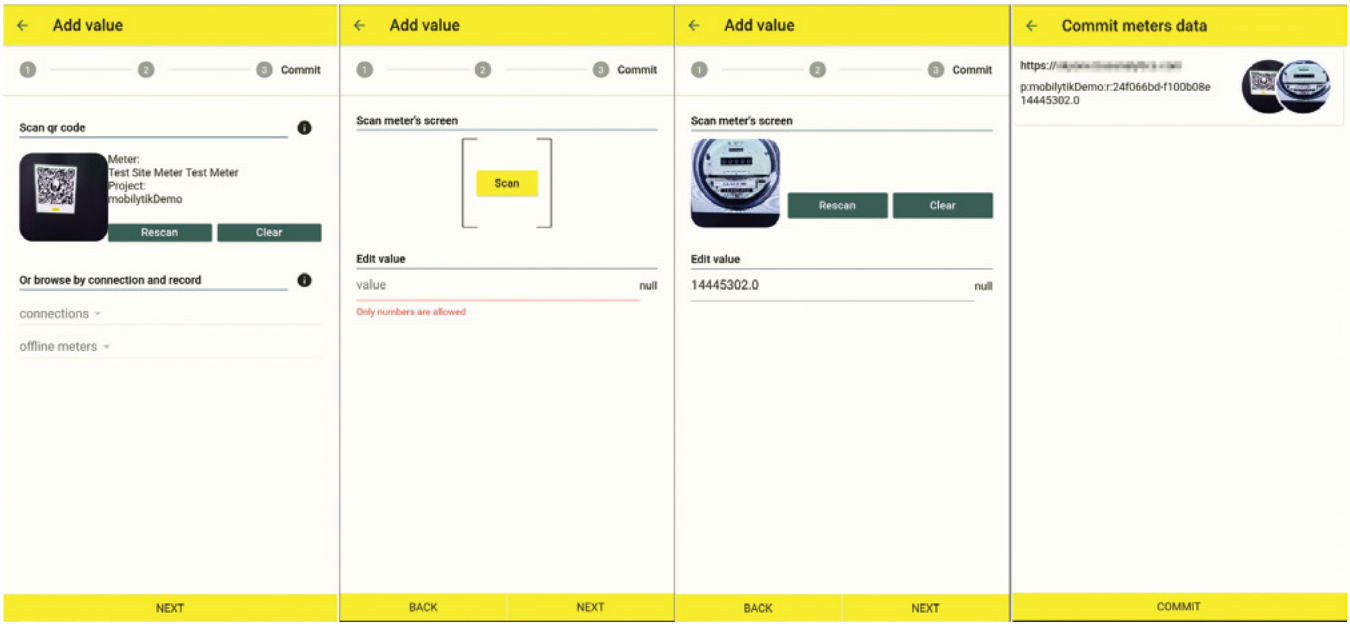

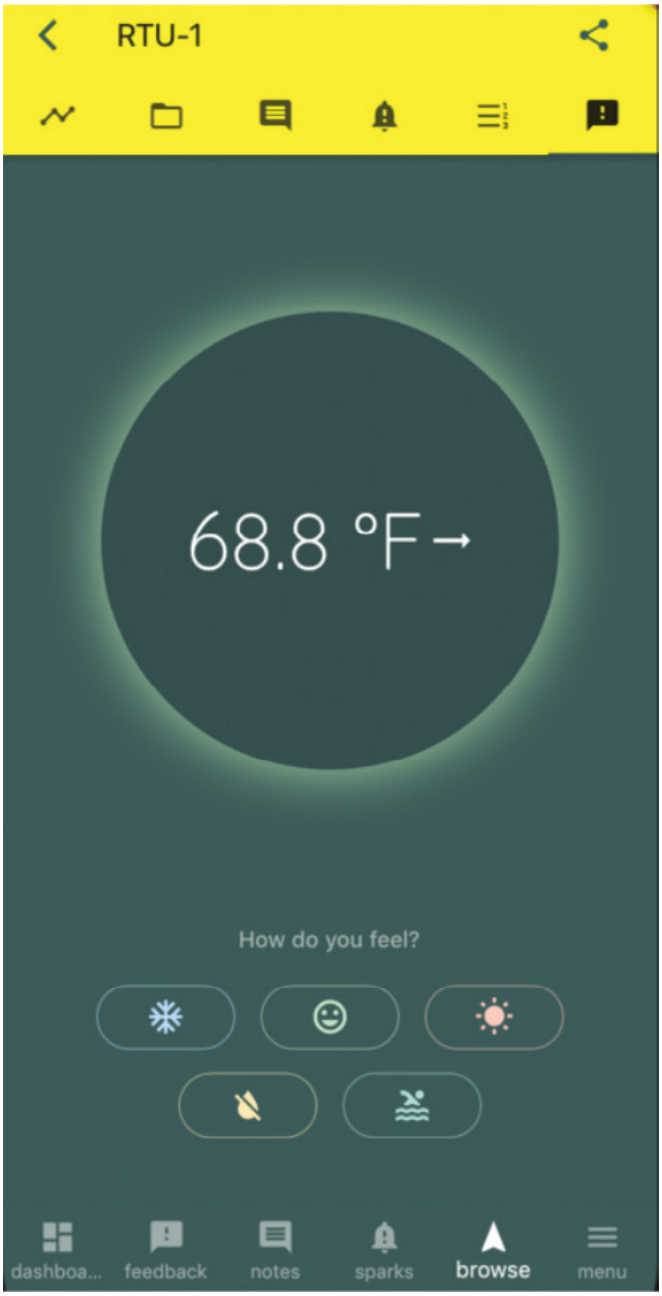

Personal Feedback

The goal of this feature is to automate tenant feedback and apply corrective behavior within the controllers automatically. Equipment of the end-user is defined as a tag on the user record. This tag defaults the GUI to the correct equipment.

The GUI contains buttons such as cold, warm, hot and humid. When the end-user clicks on the button widget, the application generates a record along with the historical data for the end-users feedback status.

An engineer can utilize this historical data to automate optimizing the equipment's run-time. With Project Sandstar-based controllers, automated changes to temperature setpoints will help ensure tenant comfort.

Engineers can also create portfolio-level feedback reports. Feedback KPI reports will help facility managers optimize tenant comfort and happiness as well.

All on a Haystack byte's journey.

Alper Üzmezler is a Managing Partner of BAS Services & Graphics, LLC, an innovator in Building Automation Technology and BAS Analytics delivery that reduces implementation and facility management energy costs.

nHaystack Version 3 Supports Additional Namespaces and Ad Hoc Tags

TRIDIUM

Version 3 of nHaystack has been completed by Tridium's engineering team under the direction of Richard McElhinney.

For those not familiar with nHaystack, it is the open-source Niagara module that enables Niagara stations to connect to external applications and transport tagged data using the Haystack protocol. Learn more about nHaystack and find official builds here. In Version 3 of this module, which is in beta, Niagara4 tags are delivered over the Haystack protocol. These tags can be added and edited with Niagara facilities or with the familiar nHaystack GUI tool. System integrators that have made an investment in Haystack tagging can more easily take advantage of Niagara4's tagging features, including Search, Hierarchies, and System Database.

Additional enhancements have been made to Version 3 as requested by Niagara users that are pushing the state-of-the-art in semantic tagging. First, the Niagara4 tags delivered over the Haystack protocol can come from multiple Niagara tag namespaces instead of only "hs". This allows tags applied using custom dictionaries to be combined with tags applied using the Haystack tag dictionary. Second, non-standard tags, i.e. ad hoc tags, are now handled consistently whether using Niagara tagging facilities or the nHaystack GUI tool - they will always become direct tags.

> ***"The Niagara community is engaged and asking for more support."***

Further development of nHaystack is underway. The Niagara community is engaged and asking for more support to accelerate effective data tagging toward the goal of semantic data interoperability of all the devices in a smart system.

Watch this space. Please contact Richard McElhinney if you are interested in participating in nHaystack Version 3 beta testing.

Eric Anderson, a Software Engineer, has been working on tagging and tag hierarchies since he joined Tridium in 2015. Tridium created and continues to enhance the Niagara Framework®, an open platform that facilitates system integration and control.

Project Haystack Turns Data into Useful Information for Smart Cities

ARC Advisory Group

Larry O'Brien has been part of both the cybersecurity and smart city practices at ARC and has authored a number of research papers, market reports, and articles including publishing this one following a Podcast with Project Haystack Executive Director, John Petze.

If the age of digitization has provided us with anything, it is data. Millions of data points from millions of sensors connected to an increasingly wide range of systems and applications. Turning all this data into useful information, however, can be a challenge. Data comes in widely varying formats, with different vendors handling and expressing data in different ways. Different protocols also have different ways of expressing data, and the lack of a single standard protocol in the world of smart cities and building automation compounds the issue.

The Cost of Point to Point Integration

Lack of common context for data also comes with a cost. In today's smart city implementation projects, much of the integration between systems and applications must be done in a point to point fashion, with the aid of systems integrator or engineering firm. This is a great cost to end uses and owner-operators. When looking at smart city software and system implementation projects, more than half of the cost of an overall project can be taken up by custom integration efforts. This offers a significant business opportunity and potential for both reduced project and operational costs, since these custom integrated environments must be maintained and tended to throughout their lifecycle.

Project Haystack

The consolidation of data into a common environment or context is necessary if you want to turn that data into useful information that can allow you to make informed decisions about how to run your cities, communities, buildings, and more. This is where Project Haystack comes in, which is an open industry initiative that is focused on providing a common metadata methodology for building automation, smart city, and other applications.

According to Project Haystack Executive Director, John Petze, "The challenge in our domain (building systems) is that device data are stored in many different formats, communicated via numerous protocols, have inconsistent non-standard naming conventions, and have very limited descriptors to enable us to understand meaning without direct human knowledge of the device producing the data. Ideally, we want data to be self-describing. Without that, a time-consuming manual effort is required before data can be used effectively to generate value."

Structure and Purpose

Project Haystack was formed in 2014 as 501(c) tax-exempt non-stock corporation to function as "a trade association with the purpose of fostering the common

association and interests of software and technology companies focused on developing semantic modeling solutions for data related to smart devices including: building equipment systems, automation and control devices, sensors and sensing devices, promotion and education with respect to the semantic data modeling industry for building automation systems, and to engage in educational activities directed towards the improvement of business conditions of the semantic data modeling industry for smart device data… All work developed by the Project Haystack community is provided for use as open source software under the Academic Free License 3.0."

More Than Just Data Tagging

Project Haystack's vision is to streamline the use of IoT data by creating a standardized approach to defining "data semantics", related services, and APIs to consume and share the data and its semantic descriptors. Project Haystack aims to make data "self-describing". It does this by relying heavily on data modeling methodology – and its tagging approach can be used in media from Excel spreadsheets and CSV text files, to data tables in embedded devices, XML representations, Web services and others.

Haystack also includes consensus-developed tagging libraries (taxonomies) published and made available for download and use (at no cost). Haystack also encompasses the REST communication protocol designed to exchange Haystack tags between applications. The group is also developing reference implementations and complementary applications are also being developed by various community members and companies. These include:

- Haystack Java Toolkit: lightweight J2ME compliant client and server implementation
- NHaystack: Niagara module to add Haystack tagging and the Haystack REST API
- Haystack CPP: C++ Haystack client and server implementation
- Haystack Dart: client library for Dart programming language
- NodeHaystack: node.js client/server implementation

Perhaps most important, however, is the community that has formed to address the challenges of data modeling for building systems and IoT devices. The Project Haystack community continues to grow and expand the equipment and device models (taxonomies) and extend the range of applications served by Project Haystack.

Open Source

Project Haystack is operated as an open source project, which makes it easy for anyone to get involved. Anyone can easily take advantage of the work of Project Haystack and contribute to it. Anyone can contribute on the forum by signing up on the website. Domain experts in each space such as chillers, data centers, or refrigeration can join or start a discussion. Equipment manufacturers who would like to see specific tag models for their products are also a great source of input. All the work done by Project Haystack is easily available to the industry community. It can be downloaded without even registering an account on the Web site. There is no cost or obligation associated with using Project Haystack techniques, tagging libraries and open source reference implementations.

Member Companies

Founding member companies for Project Haystack include Siemens, Intel, Conserveit, J2 innovations, SkyFoundry, Legrand, and Lynxspring. Associate members include almost 20 additional suppliers and service providers of all types. Membership is open to anyone. Project Haystack encompasses the entire value chain of building systems and related intelligent devices. For membership, contact projecthaystackinfo@gmail.com.

Larry O'Brien is Vice President, Research and part of the cybersecurity and smart cities and infrastructure teams at ARC, with a 20-year background in process control, process safety, and field devices/field networks.

For the first time, during Haystack Connect, companies were given an opportunity to "Pitch" their wares that support the Haystack Community. Here are a few of those "Pitches".

"Evidence for Building Retrofits that Improve Organizational Productivity" Collaborative Research Project

Continental Automated Buildings Association (CABA) members are poised to undertake a structured initiative to demonstrate that smart, integrated technologies produce organizational productivity gains, in addition to energy and environmental savings. The result will transform the high-performance buildings industry by establishing the built environment as a strategic tool for organizations to achieve their goals.

Greg Walker

Research Director

Continental Automated Buildings Association

walker@caba.org | 613.686.1814 x227 | 888.798.2222 | www.caba.org

Onyxx® BH311 BACnet® to Haystack Data Pump

Today, there is an increasing number and variety of equipment, sensors, devices and building automation systems that are available to connect within the environment and the amount of data that is available from them.

With all this data, the challenge lies in how to manage and get useful information out of the data especially as these data sets often come with various formats, naming conventions, and syntaxes.

Streamlining the interchange of data from BACnet® devices and building systems, the Onyxx® BH311 BACnet® to Haystack Data Pump provides network communication and data exchange.

The Helixx® Framework is embedded in the Onyxx® BH311 Data Pump and handles the BACnet® to Haystack protocol translation between BACnet® points to manageable Haystack points. Acting as a BACnet® client device, the Onyxx® Data Pump manages all BACnet/IP, BACnet/Ethernet or BACnet MS/TP devices connected to it.

Robert Hirsch

Chief Technical Officer

Lynxspring, Inc.

robert.hirsch@lynxspring.com | +1 (816) 347-3500 | LEARN MORE

New PADI Platform

The pitch proposes a new platform “Padi” to allow all types of users of digital assets to gather, distribute, manage and collaborate around them to increase productivity. With Padi, all of your digital assets, together with collaboration notes and related materials will be in a single secure, cloud-based repository, shared with those who need to work with them. While the elusive a single pane of glass has been promised for years, the advancement of cloud platforms and Internet-based standards is finally making this a reality for those working around buildings and facilities.

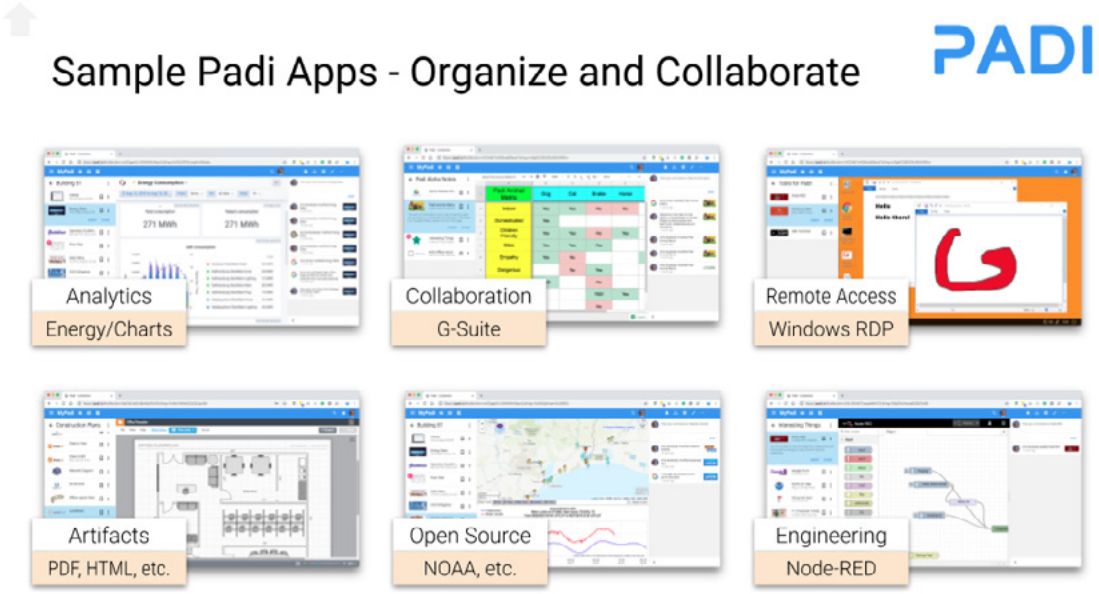

Anto Budiardjo

CEO

PADI

anto@padi.io | +1 (817) 946-3162 | LEARN MORE

Past & Upcoming Events

Demand for Project Haystack participation is up at events around the world and community members have risen to the call.

ISH 2019

Project Haystack Exhibiting Internationally

Project Haystack exhibited at the ISH 2019 Conference in March at Messe Frankfurt in Frankfurt, Germany.

Over 2,500 exhibitors, including all market leaders, launch their latest products, technologies and solutions onto the world market at ISH. ISH has gained a leading role worldwide as the occasion per se when the sector

comes together – 64 percent of the exhibitors and 40 percent of the approximate 200,000 visitors come from outside Germany.

ISH and Light Building in Frankfurt are international presentation platforms for the products of manufacturers of market-ready solutions, systems and products for house and building automation. Held in alternate years, the two leading trade fairs offer companies working in the field of house and building automation the best possible opportunity to present their latest products and innovations – at Light Building where the focus is on electrical engineering and at ISH which spotlights the heating, air-conditioning, ventilation and sanitation sector.

Haystack Europe 2019

London 26th September 2019

Project Haystack had a good turnout for the first ever event in London in September.

Presentations explained why data modeling and tagging are essential for the future of smart buildings and how

Project Haystack created an open-source standard for building services related data. Presentations also covered the implications for the specification and design of BMS and other systems in smart buildings, how tagging can enable intelligent alarm management, the automation of control systems configuration and commissioning, analytics and the integration with BIM.

Haystack Connect 2019 Recap

The 4th biennial Haystack Connect conference, was another great success demonstrating the continued growth of the community and widespread adoption of the Haystack standard by manufacturers, developers and systems integrators large and small. Attended by Project Haystack community members and representatives from the smart building and IoT industry, 24 sponsors and exhibitors from around the world, the community forged ahead once again, addressing the evolving challenges of making data easier to use and interoperable across the built and IoT environments.

The event featured over 50 speakers from systems integrators, technology providers, OEMs, and building owners and operators to presentations focused on the technical side of Project Haystack and the Working Groups, and the public review of the new Haystack 4!

SOME HIGHLIGHTS

Haystack 4 is the culmination of over a year's worth of work resulting in new vocabulary, taxonomy, ontology and inference, that can better describe real-world data through tagging and semantic modeling.

The first "Pitchfest". Think Shark Tank for Project Haystack, allowing companies 10 minutes to pitch their products and use of tagging and data modeling.

And, always a highlight of every Haystack Connect is the performance of the legendary Interoperable Blues Band (IBB). Not just any band, but the integration of talented attendees playing with the featured band, "Band in the Round".

All Haystack Connect 2019 Speaker Presentations are available at: www.haystackconnect.org/schedule.

You can find a detailed recap of the event in the June 2019 issue of www.automatedbuildings.com.

Documentation describing the Haystack 4 design is available at: https://project-haystack.dev.

Past & Upcoming Events

In 2020

AHR EXPO 2020

Haystack 4 - The Continued Evolution of Semantic Tagging – What it Is and Why it Matters

Monday, February 3
10:30 am - 11:30 am | W311C

Understanding the need for semantic modeling of device and equipment data has matured significantly in the last decade, and the requirements and techniques for applying semantic modeling to equipment data are advancing rapidly.

PRESENTED BY: AutomatedBuildings.com.

CABA Smart Buildings Summit 2020y

May 3 - 5, 2020

PRESENTED BY:
Continental Automated Buildings Association.

WG Working Group Updates

Tagging initiatives are made official by launching a Working Group with a defined proposal and good visibility. Join a WG now!

#551

Haystack Kind Reference

Champion:
Brian Frank,
SkyFoundry

After eighteen months, WG 551 has transitioned to public review. We have given WG 551 the umbrella term "Haystack 4". Public review is being run using a new website found at https://project-haystack.dev (there is a link on the primary website too).

Haystack 4 includes the following key new features:

- new design for definition of tags and tag sets
- organization of tags into a taxonomy tree structure
- symbol relationships between tags for a richer ontology
- enhancements to the filter query language
- RDF support

https://project-haystack.org/forum/topic/551

#595

Project Sandstar

Champion:
Alper Uzmezler,
BASSG

Since the last Haystack Connect meeting, the Project Sandstar WG slack room has grown to 96. We have completed Sixfab IoT integration, indoor air quality integration to Sandstar. Currently, we are working on integrating project Sandstar to human counter with stereo camera with machine learning identification technology.

Please join our slack room and promote the world's first hardware-agnostic DDC that is opened source. As a reminder, project Sandstar is an open source initiative that combines Project Haystack with the Sedona Framework.

https://project-haystack.org/forum/topic/595

#609 AHU Standing

Champion:
Jay Herron,
BuildingFit

Here are the result of the changes from the AHU working group, which was closed in May: The following tag definitions were changed to:

- ahu: Air Handler Unit: An enclosure with a fan that delivers air to a space via ductwork and
- performs one or more of the functions of cleaning, heating, cooling, humidifying, dehumidifying,
- ventilating or circulating the air.
- mau: Makeup Air Unit: An AHU that brings in 100% outside air and no recirculated air.
- rtu: Roof Top Unit: A unitary AHU that is weatherized for outdoor use.
- fcu: Fan Coil Unit: A unitary, factory-made enclosure with a fan and coil that recirculates air in a
- space and performs one or more of the functions of cleaning, heating, cooling, humidifying,
- dehumidifying, or ventilating the air. Unit heaters are examples of FCUs.
- heatPump: Heat pump: An air handling equipment which employs a vapor compression cycle
- with a reversing valve to provide heating or cooling to a space.
- ductDeck: Cold, hot, or neutral deck. Can be applied to 'ahu' equip if the AHU is part of a
- 'dualDuct' or 'tripleDuct' system but only serves a single deck type.

The following definition was added:

- crac: Computer Room Air Conditioner: An FCU designed to condition a computer room.

https://project-haystack.org/forum/topic/609

#667

Cybersecurity

Champion: Fred Gourdy, Intelligent Buildings

The Cybersecurity Working Group (CySecWG) began in 2019 however it has not fully kicked off. The plan is to begin 2020 with monthly calls to work on cybersecurity use cases. The use cases will identify what tags should be used for the purpose of monitoring, policy adherence, and classification. The foundation for these tags will revolve around the NIST Cybersecurity Framework.

https://project-haystack.org/forum/topic/667

#734

ATES Systems

Champion: Jaap Balvers, BAM Energy Systems, part of Royal BAM Group nv

The ATES Systems WG is developing a standard for tagging aquifer thermal energy storage systems. This type of system is quite common in The Netherlands and surrounding countries. Typically, these are “open” systems with at least one “hot” and one “cold” groundwater source. Heat or cold from the groundwater source is converted into energy-efficient hot or chilled water using heatpumps.

We are open for more members to join our effort, and plan to share the first draft proposal on the Project Haystack forum soon.

https://project-haystack.org/forum/topic/734

The Project Haystack community develops and freely offers a range of reference implementations to enable product manufacturers and application developers to quickly implement Haystack tagging and communications in their products.

Wiki

Haystack Wiki: Source for docs, and tag definitions

Java

Haystack Java Toolkit: Light weight J2ME compliant client and server implementation

Niagara

nHaystack: New Updated Niagara module to add Haystack tagging and Niagara REST API for AX and N4

C++

Haystack CPP: C++ Haystack client and server implementation

DART

Haystack DART: Client library for Dart programming language

Node

Node Haystack: node.js client/server implementation

Python

pyHaystack: Python client implementation

Check out these documents and audio resources to quickly come up to speed on Project Haystack tagging benefits and the methodology.

Download Detailed Reference Implementation Document. "Implementing Project Haystack: Applying Haystack Tagging for a Sample Building."

Download Harbor Research whitepaper with technical overview. Defines the concept of tags, breaking down and explaining the essential data elements.

Listen Audio Stream of "Making Internet of Things Device Data Just Work!" a Memoori webinar featuring John Petze and Marc Petock on Project Haystack.

Open REST API Description. Explains simple mechanism to exchange tagged data over web services

Download Haystack Guide Specification. Now available in English, French and Chinese.

Download CABA whitepaper that outlines how to use Haystack tagging in applications related to buildings, energy, and facility management.

Want to get involved in the Project Haystack open-source community? There are a number of ways and levels of involvement.

Contribute your expertise: Participate in the Project Haystack open forum discussions.

Join a Working Group: Project Haystack has members working together on developing tag sets and resolving other challenges related to particular topics. See the list of active Working Groups that you could join today here.

Become a Member: Project Haystack Corporate Associate Memberhip has many advantages. Email us to learn more at projecthaystackinfo@gmail.com.

All the latest Marketing information about Project Haystack has a home now at marketing.project-haystack.org.

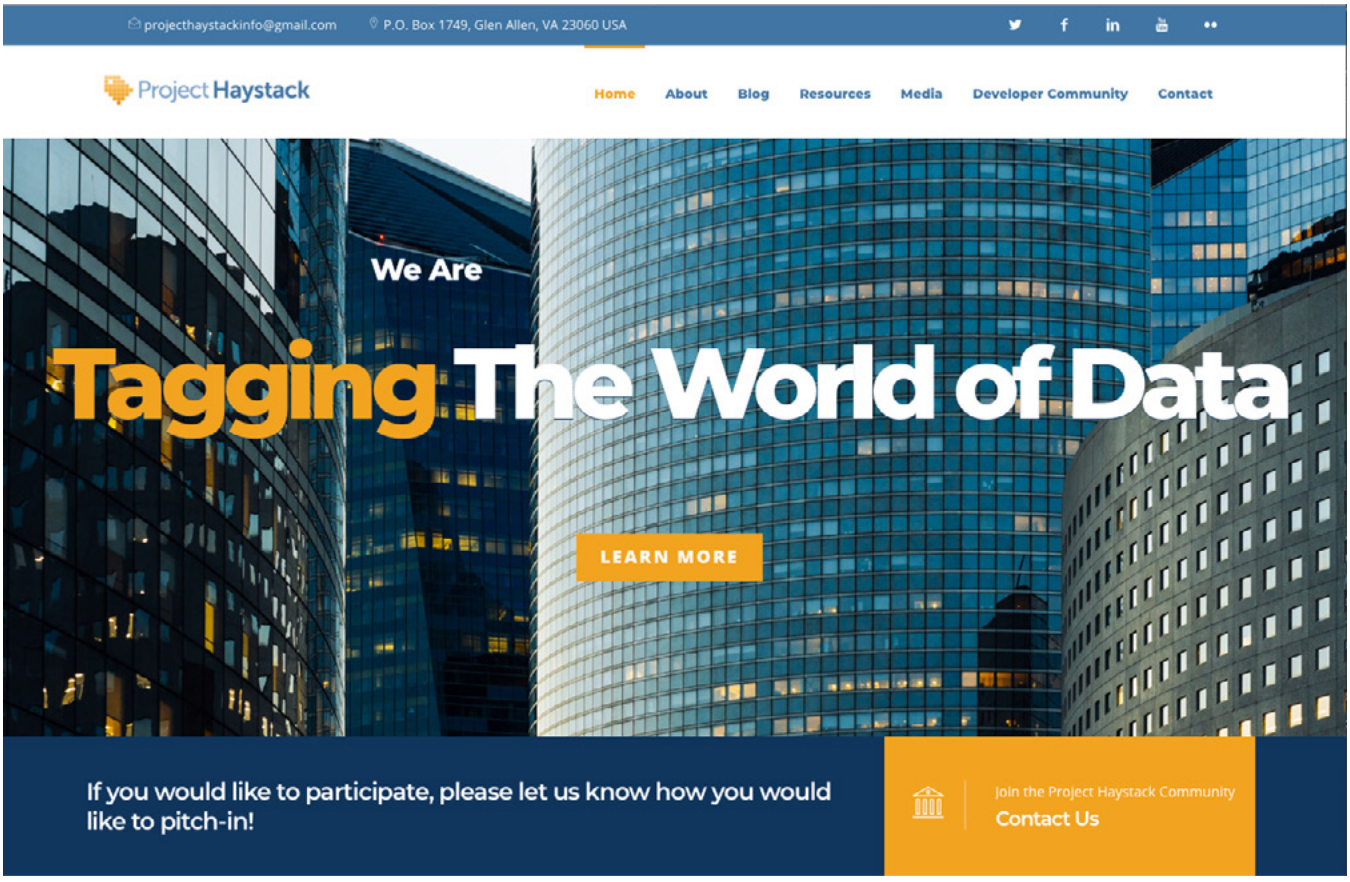

Social Media Curation

Here is some of the information shared by Project Haystack members on Twitter and LinkedIn. Follow their tweets and shares to learn about Haystack-enabled recent projects, new products and practices.

Projects

Discover Legrand's educational building reference projects.

How BuildingFit deployed SkySpark across a 6.5M sq. ft. for Banner Health.

How BuildingFit helped Macau Casinos with customized solutions to monitor and maintain their 30M square feet.

Energy Management Association announces collaboration with California Energy Alliance.

Practices

Eyeing 5G Edge Computing, Intel aquiring Smart Edge.

Home automation interest rises in line with IoT: REPORT.

Legrand signs the new French Climate Pledge.

Eight ways to secure your data on IoT devices.

"We need to explore how to make the building user experience seamless and amazing..."

Products

FIN Framework5 more robust than before.

Conserve It Edge IoT 534 combines a fully programmable controller and web server duties in a single device.

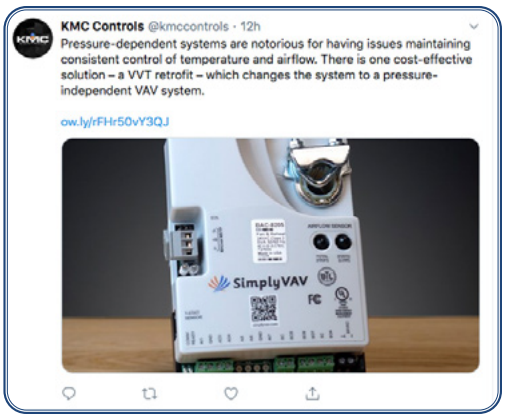

KMC Controls VVT retrofit - changes systems to a pressure-independent VAV system.

Siemens solutions enabling lab space to become more #energyefficient.

Conserve It distributing Sensors and Metering by Veris Industries and Inovonics.

ControlTrends Podcast Episode "Let's Talk Niagara 4.8".

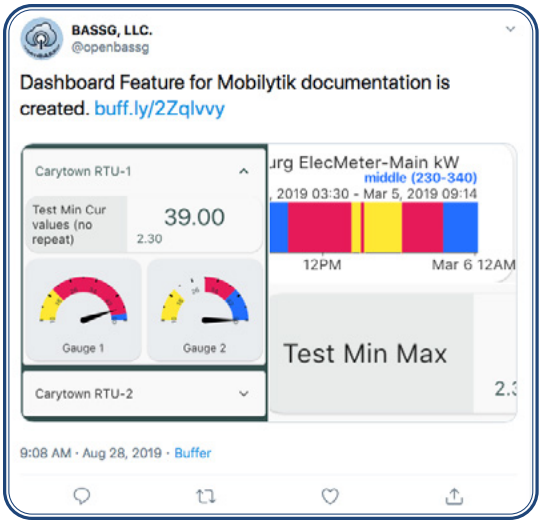

Dashboard feature for Mobilytik documentation created.

Tridium and Conserve It announce general availablilty of Niagara Framework® 4.8.

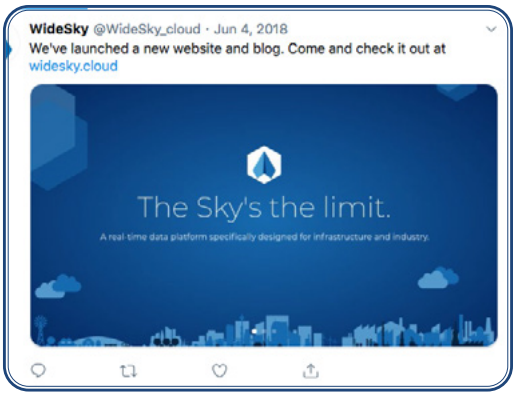

WideSky launches new website.

Project Haystack

connections

Become an Advertiser

The Haystack Connections Magazine advertising program is a cost-effective way for companies that provide complementary products and services to reach the growing and dynamic Project Haystack Community. This community is at the very forefront of intelligent buildings and the IoT. Haystack Connections is a premier advertising vehicle to reach this prime audience. With 8,000+ known readers, it is an incredibly cost-effective advertising opportunity. For rate info, email robin@haystackconnect.org.

Members

Founding Members

Conserve It was founded in 2007 with a focus on centrifugal chiller efficiency systems. Over time it has diversified into complete HVAC&R plant management including monitoring, reporting and controls, energy performance contracting, energy management consulting and distribution of industrial and building automtion products and sensors from leading international suppliers worldwide. Conserve It provides a range of unique products and services in this area.

Intel has been leading the pursuit of Moore's Law for its entire existence. We have continuously advanced silicon technology and moved the capabilities of the industry forward. Today, the unmatched scope and scale of our investments in R&D and manufacturing ensure Intel continues to maintain industry leadership and drive innovation to provide our customers and consumers with leading-edge products in high volume.

J2 Innovations brings powerful engineering tools, visualization and software technology to those involved in BAS installations. J2 is the developer of FIN Stack, a software technology that combines the core functionality of a Building Automation System (BAS) for connecting and controlling devices with the added benefits of a Building Operating System (BOS) to manage and leverage data. The technology uses Project Haystack tagging and data modeling to provide unprecedented capabilities and functionally.

As a leader in electrical and digital infrastructure solutions for all types of buildings, Legrand helps enhance everyday life for its customers. Legrand's Eliot program (Electricity and IoT) is speeding the deployment of Legrand's connected devices and accelerating the evolution of connected buildings. Eliot is powering development of new Legrand products for the benefit of private and professional users alike.

Embracing open software and hardware platforms, Lynxspring develops and manufactures innovative edge-to-enterprise solutions. We enable better building automation, better energy management systems, better control systems and specialty machine-to-machine and IoT applications. Deployed in billions of square feet of commercial buildings across North America and beyond, Lynxspring's smart solutions simplify integration and interoperability, and help connect your smart building's data.

Siemens Building Technologies consists of three Business Units: Building Automation (BAU): Control Products and Systems (CPS); Fire Safety and Security (FSS). These business units combine offerings for building security, life safety and building automation within one company as a service and system provider, and as a manufacturer of respective products. By virtue of the unique combination of these business sectors, the company occupies a leading position worldwide.

SkyFoundry's mission is to provide software solutions for the age of the "Internet of Things". Areas of focus include building automation and facility management, energy management, utility data analytics, remote device and equipment monitoring, and asset management. SkyFoundry products help customers derive value from their investments in smart systems.

Associate Members

Accu-Temp Systems is committed to delivering safe, comfortable environments for its customers. It leverages tools like secure mobile devices, cloud computing and advanced analytics. It offers systems integration services that help building owners protect their investment in existing direct digital controls, extending their useful lifetime while enjoying next-generation access and control.

Altura Associates is a professional services firm that goes beyond the traditional consulting model. Our team works closely with our client organizations to develop programs that offer immediate and lasting impacts, build capacity, and drive long-term value. The team combines expertise in mechanical/electrical engineering, energy management, environmental science, and financial analysis.

BASSG is an innovator in building automation technology and BAS analytics delivery. Its BASSG branded in-house developed easy-to-deploy, multi-system software tools reduce BAS implementation and facility management energy costs. BASSG also has multiple distributorships and can be a one-stop provider for everything-BAS at unbeatable value.

BUENO Systems is the Australian leader in data and information driven operational property services. BUENO delivers superior data related and technology driven services based on fault detection, optimization and business intelligence that simplify their clients operations and enhance their effectiveness across all building sectors and building information systems.

BuildingFit creates unique solutions for clients to ensure a proper fit between SkySpark® and their team. We do this through site construction, analytics, custom programming, SkySpark® Apps, reports, training, SkySpark® Licensing. BuildingFit is a SkyFoundry endorsed SkySpark Essentials provider.

The Continental Automated Buildings Association is an international not-for-profit industry association dedicated to the advancement of integrated technologies for homes and buildings. The organization supported by an international membership of over 300 organizations involved in the design, manufacture, installation and retailing of products relating to home and building automation.

EMA is a new and innovative association that is dedicated to advancing the quality of energy management products and services for the benefit of the building owner. The founding members are certified Energy Management Professionals (EMP), a program that was developed by ACG, the world's leading association of certified commissioning authorities. Management of the program has been transferred to EMA.

Intellastar Technology is at the Intersection of Smart Buildings and Smart Grid. The InferStack Software Platform is deployed in Servers and T-Star Field Devices, communicates over Intellastar Connect Cellular Data Service, to provide a complete technology to deliver Smart Buildings and Smart Grid solutions.InferStack connects to the in-building systems to provide Energy Monitoring and Analysis, Analytics for Fault Detection and Diagnostic, Control for Plant Optimization--all features to make a smart building and reduce energy consumption and waste.

Intelligent Buildings, LLC, a nationally recognized smart real estate advisory services company, provides planning and implementation of next generation strategy for new buildings, existing portfolios and smart communities. Their work includes "The Smartest Building in America", the largest energy analytics project in North America, the smart buildings standards for the U.S. and Canadian governments, conception and management of a Clinton Global Initiative and the recently released Intelligent Buildings CyberSafe service.

Associate Members

IoT Warez develops custom software that helps technologies communicate together.From state of the art data centers to environmentally conscious facilities, our software development team is capable of building solutions that connect anything and everything. IoT Warez offers a suite of hosted software options that provide customized solutions. Our platform-as-a-service connects multiple brands of software into one platform that can be remotely managed from a smart device.

KMC Control is an American manufacturer of open, secure, and scalable building automation solutions. From secure hardware devices to smart and connected software, KMC delivers embedded intelligence and optimized control.. It is committed to providing industry-leading Internet of Things-enabled automation solutions with leading tech suppliers to increase comfort, convenience and to help reduce energy usage.

KNX Association represents KNX technology now used in applications for lighting and blind control, security systems, HVAC, monitoring, alarming, water control, energy management, smart metering as well as household appliances, audio/video and more. KNX provides a single, manufacturer-independent design and commissioning tool (ETS), with a complete set of supported communication media and configuration modes. It is approved as a European and an International standard.

Kodaro expands building system connectivity through dynamic software developed for the Internet of Things. It helps contractors, controls companies and end-users find value in building data gathered from the edge to the cloud. It develops software to create more connectivity between systems, giving increased access to better data, not bigger data. Kodaro's goal is to provide actionable analytic information, developed from real-world expertise with all building systems.

SensorFact® is a cloud-based data acquisition and storage service for sensor data. It allows for sending sensor data from one location, through their pointCollex technology or directly to their pointCollex API, to a client account in sensorFact. Once there, clients can name, organize, tag, monitor, and choose which sensors data to store long-term. In addition, sensor data is available to share or integrate with other systems.

Founded in 2013 and headquartered in Israel, SmartGreen is a provider of advanced solutions that optimize the operational and energy performance of commercial and industrial buildings and facilities. Our flagship product, OptiNergy™ is an innovative cloud-based platform that combines machine learning algorithms, big data analytics, and Internet of Things (IoT) technology and devices.

Tridium is a world leader in business application frameworks — advancing truly open environments that harness the power of the Internet of Things. Our innovations have fundamentally changed the way people connect and control devices and systems. Our products allow people and machines to communicate and collaborate like never before. They empower manufacturers to develop intelligent equipment systems and smart devices for enterprise and edge assets.

VRT has been pioneering the provision and support of industrial information solutions since the mid 1980s. Its main business is implementing solutions based on real-time information to improve operational efficiency and safety, and to reduce risks related to business continuity. To meet the increasing demands in the areas of smart buildings and smart cities, VRT has developed its own cloud technology-based IoT management platform, WideSky®.

Yorkland Controls has roots in distributing and warehousing heating control products such as Flame Safeguard and Burner and Boiler Management Systems, and has expanded into new markets including Building Automation, Lighting, Security and Energy Services. It works to promote the advantages of controls to the industries and markets that it serves and to demystify available technology for its customers.

www.project-haystack.org