SPRING 2021

Project Haystack

ISSUE 09

connections

Journal of the Haystack Community - Solutions for Interoperable Device Data

10 ANNIVERSARY

Project Haystack

- Haystack 4 Update
- New Haystack JSON Encoding
- Haystack & Data Interoperability
- Tagging Niagara Station Components
- The Importance of Data Standardization for Retail & Hospitality
- A Look Back at 10 Years of Project Haystack

www.project-haystack.org

Project Haystack

Founding Members

ConserveIt

J2 INNOVATIONS
A Siemens Company

legrand®

LYNXSPRING®

SIEMENS
Ingenuity for life

SkyFoundry

Associate Members

Accu Temp
Building System Integrators

ALLANDER ANALYTICS

ALTURA

Automated Logic

BAS SERVICES GRAPHICS

BRAINBOX AI

bueno
Built Environment Optimisation

buildingfit
Making SkySpark® Work For You.

buildingsIOT

CLOCKWORKS ANALYTICS

CABA
Continental Automated Buildings Association

e-Magic

EMA
ENERGY MANAGEMENT ASSOCIATION

Intellastar

INTELLIGENT BUILDINGS™

KMC CONTROLS

KNX®

kodelabs

RESOLUTE™
BUILDING INTELLIGENCE

TRIDIUM

WideSky

Yorkland Controls
ENVIRONMENTAL SOLUTIONS

CONTENTS

EDITORIAL

PROJECT HAYSTACK'S VIEW

CONTRIBUTED ARTICLES

A LOOK BACK AT THE PAST 10 YEARS

CONTENTS *cont'd*

NEW MEMBER PROFILES

EVENTS

WORKING GROUP UPDATES

TOOLS FOR DEVELOPERS & INTEGRATORS

SOCIAL MEDIA CURATION

DIRECTORIES

Congratulations! 10 Years of Perfecting Data Tagging

by Robin Bestel, Managing Editor, Project Haystack Connections Magazine

Welcome to the Spring 2021 issue of the Project Haystack Connections Magazine. This ninth issue commemorates the 10th Anniversary of Project Haystack, founded in 2011, and the commitment of the community to perfect data tagging and further solidify the value brought by the Project Haystack open-source methodology of semantic tagging for the built environment and beyond.

The member companies and supporters of Project Haystack continue expanding the standards for semantic modeling methodology and building on the tagging libraries for more and more applications. This community-driven, open-source process is engaging more and more companies that work on different facets of specifying and implementation.

The Connections Magazine Spring 2021 issue consists of articles, conversations and updates from Project Haystack members and supporting companies. Here are just a few highlights of this, our Project Haystack 10th Anniversary, commemorative Connections Magazine.

Brian Frank, Technical Lead of Project Haystack, provides an update on Haystack 4: 3.9.10. The Project Haystack website, **www.project-haystack.org**, is in the final stages being transitioned to the new Developer website, **www.project-haystack.dev**, focusing on the new Haystack 4 methodology and assignment of tags.

Contributed articles were submitted by J2 Innovations, ***"HAYSON - New Haystack JSON Encoding"*** and "***The Importance of Data Standardization for Retail and Hospitality***". "***Haystack and Data Interoperability***" is an article written by Marc Petock, Executive Secretary on the Board of Project Haystack and Chief Marketing & Communicaitons Officer of Lynxspring, a Founding Member of Project Haystack. Eric Anderson and Therese Sullivan of Founding Member Tridium, contributed the article ***"Tagging Niagara Station Components"***.

Since our last issue of Connections Magazine, we hope you will join us in welcoming our four new Associate Members just in the past month.

Read their new Member Profiles here: **Allander Analytics**, **Automated Logic**, **Buildings IOT** and **Clockworks Analytics**.

Once again, we are very excited to promote under **Events**, **Haystack Connect 2021**, our biennial conference being held next week, May 4 - 6, 2021. Check out the full Schedule and Speakers lined up for this years, first-ever, virtual conference. You won't want to miss it and thanks to our Sponsors, registration is FREE to all!!

Project Haystack **Working Groups** contributed updates to their work. And as always, we have sections dedicated to **Tools for Developers and Integrators** and **How to Get Involved**, a curation of social media highlighting Project Haystack Members new **Projects**, **Practices** and **Products**, and our **Members Directory**. There is a list of all the **Advertisers**, to whom we thank for their sponsorship that supports publishing our Connections Magazine twice a year.

I personally want to thank everyone I had the pleasure to work with and contributed to this 10th Anniversary, commemorative Connections Magazine. To anyone reading this magazine for the first time, more information about Project Haystack and how to become a member, is available at **marketing.project-haystack.org** and for Developers, please visit **www.project-haystack.org** and **www.project-haystack.dev**.

I look forward to working with everyone again to create the next Connections Magazine in 2021 as we continue celebrating the journey we embarked on 10 years ago to make our buildings more comfortable, easier to manage, and waste nothing!

Introducing FIN 5.1 - The latest update for the world's most innovative software framework for smart buildings, smart equipment and IoT

The latest update release for FIN Framework, coming Summer 2021

Updates include:

- A dashboard builder app – create data-driven dynamic dashboards
- FIN Network - the ability to set up a distributed architecture, and view FIN sites and instances through one view
- BTL certification - B-OWS (Operator Workstation)
- KNX IP Connector for seamless integration
- Edge2Cloud enhancements

Find out more
j2inn.com

J2INNOVATIONS
A Siemens Company

Celebrating Our Growth and Accomplishments

This year, 2021, marks ten years since the beginning of Project Haystack. Wow, ten years. Who would have thought?

Celebrating a 10-year anniversary milestone is a big achievement in itself. No individual can take the credit of all those years of success. It is a celebration for all who have been a part of the movement through all these years.

Any organization's growth and accomplishments depend on people. For Project Haystack, that is a community of volunteers that devote their time and effort to addressing one of the key challenges in achieving the promise of the IoT and the goals of sustainable, intelligent, efficient built environment. Our community not afraid to welcome diverse viewpoints, rethink technology based on new learnings and requirements and to change. It takes a remarkable community. And the Haystack community is no less than remarkable.

As we celebrate Project Haystack's 10-year anniversary, we want to thank the global community – our founding members, associate members and the independent contributors for their commitment and effort. Thank you to the curators of the methodology; thank you to the multitude of working groups; thank you to the industry manufacturers who have made it part of their technology and solutions; thank you to the integrators who have utilized Haystack in well over 30,000 facilities; thank you to the owners, operators, facility managers and government organizations who are using Haystack in the daily operations of their buildings and IoT deployments.

We are very excited to see what the next 10 years will bring. Here's to the next 10!

John Petze
Executive Director, Project Haystack

Marc Petock
Executive Secretary, Project Haystack

Haystack 4 Update: 3.9.10

Haystack 4

Haystack 4 continues to make refinements to our next generation ontology. We have upgraded project-haystack.dev with a new version. There is a **huge** amount of new content to review in this latest version 3.9.10.

See previous version notes:

- 3.9.9 topic 844
- 3.9.8 topic 797
- 3.9.7 topic 743
- 3.9.6 topic 737
- 3.9.5 topic 714
- 3.9.4 topic 699
- 3.9.3 topic 694
- 3.9.2 topic 687

This is the final preview release we will run on project-haystack.dev! Next week we will be updating this website to the latest Haystack 4 documentation.

If you have comments or feedback on this release, please open a new forum post for each topic to organize the conversation.

Documentation

There are 16 new chapters added to the documentation to cover how the ontology is used to model entities. There are new chapters on horizontal entities like sites, spaces, and equip. And we've ported/rewritten many of the old chapters for vertical systems including Meters, AHUs, VAVs, and Plants.

Hayson

The work from WG 792 to define a new JSON dialect code named Hayson has been integrated. The JSON chapter reflects the new specification. All the JSON downloads are now made available in the updated format.

Greenhouse Gases

This release incorporates the work from WG 776 to define several new gases related to global warming. We also define new emission quantities for these greenhouse gases.

Air Quality

In addition to the greenhouse gases, we have flushed out many new air quality substances including:

- ch2o: formaldehyde
- nh3: ammonia
- no2: nitrogen dioxide
- o3: ozone
- pm01: particulate matter 0.1

As part of the WG 776 proposal, we have tweaked how these terms are represented in the ontology. They now subtype from gas (or in case of pmXXX tags from substance). Then we pair these tags with concentration when modeling them as quantities. This models the

terminology more correctly and enables us to pair co2 with emission to distinguish between air quality measurements versus greenhouse emission measurements.

ATES

We have incorporated the work from **WG 734** on aquifer thermal energy storage systems. See the new ATES chapter in docHaystack for an overview.

Protos

As part of the effort to **double down** on usage of children protos for point enumeration, we are now including the exploded protos in machine readable formats. See the new Protos section in the **Downloads** page to download protos in any of the standard formats (Zinc, JSON, CSV, Turtle, etc).

Also as part of this effort, we have reviewed and expanded the children protos of many different equip types. In cases where we didn't make a distinction between sensor/cmd or sensor/sp we have now made this explicit.

State Tag

Previously we attempted to model all the points with a "subject" and a "quantity". This led to a situation with a new state tag to pair with tags like run or enable. However it was awkward and isn't really a viable design to use for all the different point types we need to model. So we have nuked the state tag. See the Motors chapter for how to properly model run/enable command and status points.

VFD Modeling

As part of the review process we followed a similar philosophy with the drive tag which was previously paired with speed and freq to model VFD points. In the last version vfd was defined a subtype of motor, which technically wasn't correct since a VFD is a peer component to the actual motor. So now vfd is just a marker tag you put on a motor equip. Then you also use the vfd marker in conjunction with speed or freq to model the command/status points. See the Motors chapter for how it all turned out.

VAVs

Previously we had started creating some subtypes of VAVs including a coolOnly-vav and fanPowered-vav. But after some discussions, it seems better to handle VAVs more akin to AHUs where we flatten all the options out. Those options still exist, but not as first class subtypes. See the new VAVs chapter for details.

We did introduce a new breaking change to rename vavMode to hvacMode so that it can be used in other situations such as on an AHU, thermostat, etc.

Dual Effective Setpoints

One the issues that came up is how to model thermostats which don't provide a true effective setpoint across all modes, but rather provide separate effective heating/cooling setpoints. I have written some language up how this should work in the new Zones chapter by requiring an hvacMode point in that case. It warrants a deeper review and maybe discussion.

Flow/Thermal Meters

After some discussions, there seems be interest in using the tags flow and thermal on meter equip based upon what points you expect. Review the Flow Meter and Thermal Meter section in the new Meters chapter to see if you like the current design.

Meter Loads

I have written up meter load modeling using the new elecRef, steamRef, etc flow references. I think many of us had an implicit idea that this is how it should work and we no longer should have tags such as elecMeterLoad. In my mind these tags model the flow of electricity/fluid almost like a wiring diagram. You insert meters, valves, loads, etc all together consistently this way. But this might be contentious not separately modeling supply flow from meter flow. So it probably warrants further discussion.

Loops

Based on the requirements of the ATES work, it made sense to revert plantPrimaryLoop and plantSecondaryLoop back to the original names from Haystack 3. We made the definitions a little more generic so they can be used in other situations beside just central plants.

Weather

Note that the weatherRef tag is now weatherStationRef for consistency. We also added an enthalpy quantity and its associated weather point.

ChangeLog

Version 3.9.10 (23 Apr 2021)

- Rename weatherRef to weatherStationRef for consistency
- Annotate equip which typically uses electricity as elec-input
- Rename docSection to docAssociations
- Change protocol to subtype from marker, not entity
- Add nodoc, nosrc, deprecated
- Generate children prototypes as RDF blank nodes
- Update filters chapter for ref list handling
- Exploded point types (attempt only)
- Remove state and stateQuantity in favor of keeping run/enable on their own
- Tweaks to how vfd is used, remove drive tag
- Expand some of the sp/sensor cmd/sensor proto pairs
- Add enthalpy tag, associated weather, duct, and zone points
- Add protos grid files to the distribution
- Add greenhouse gases into ontology from WG 776
- Expand children: ducts, zones
- Port min/max marker tags from Haystack 3
- New alarm point; alarm sensor on motor, chiller, boiler
- New air-velocity quantity
- Rename vavMode to hvacMode, clarify use with dual effective setpoints
- Other new tags: deadband, deviceRef, nh3, no2, o2, o3
- Add heatingOnly as peer to coolingOnly, clarify usage in AHU
- FumeHood children protos
- Add enable cmd as plant level point
- Add cool/heat enable cmd as ahu points
- Revert plantPrimaryLoop/plantSecondaryLoop to primaryLoop/secondaryLoop
- Rework VAVs to not use subtypes (more like AHUs)

- Add ups equipment type
- New Sites chapter
- New Spaces chapter
- New Equips chapter
- New Points chapter
- New Weather chapter
- New Devices chapter
- New Meters chapter
- New AHUs chapter
- New VAVs chapter
- New Plants chapter
- New Zones chapter
- New Motors chapter
- New ElecPanels chapter
- New ATES chapter and tags from WG 734
- New DataCenters chapter

If you have comments or feedback on this release, please open a new forum post for each topic to organize the conversation.

https://project-haystack.org/forum/topic/908

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

HAYSON - New Haystack JSON Encoding

J2 INNOVATIONS
A Siemens Company

"Making Haystack data more easily accessible is key to driving greater adoption of the standard. This is even more important as the standard grows in complexity and naturally becomes more difficult for people to consume. The standard needs to stay focused on the problems it's really trying to solve."

Haystack has enabled us to better understand the connected world of device and equipment data through the use of tagging. Haystack version 4 builds upon this success by providing a machine-readable way for relationships to be formally defined and consumed. There's far more to Haystack than just its tagging model. The standard also consists of:

- **REST APIs:** described as ops in the specification, these are a standard set of REST APIs that enable clients and Haystack servers to interoperate with each other.
- **Granular type system:** the simple building blocks Haystack data is built from. This covers everything from numbers with units to references between records.
- **Filters:** a simple query language for working with Haystack data.
- **Custom data encoding formats:** used for encoding Haystack data. This includes Zinc, Trio and JSON.

So, you can see that Haystack also provides far more than just a way to tag data. It consists of a broad range of technologies from network communication protocols, query languages, to data encoding formats. It's the custom data encoding formats that are covered in this article.

To a Haystack noob* (my sons still call me this whenever I play video games with them – I originally assumed it meant that I'm cool), Zinc and Trio are a little puzzling. Why do these data formats exist? Are they needed? What problems are they trying to solve? Why can't we just use JSON and YAML like the rest of the known world? Despite their obvious advantages in brevity, Zinc and Trio require custom parsers so Haystack data can be understood by clients. This has partly led to a few contributors creating their own haystack libraries covering most of the world's most popular programming languages. Unfortunately, most of these still need to be updated to provide support for all of Haystack version 4's new features.

The obvious fallback is here - just use JSON. Here's a snippet of a grid encoded using the old format:

```
{
  "meta": {"ver":"3.0", "projName":"test"},
  "cols": [
    {"name":"dis", "dis":"Equip Name"},
    {"name":"equip"},
    {"name":"siteRef"},
    {"name":"installed"}
  ],
```

```
  "rows": [

    {"dis":"RTU-1", "equip":"m:",
"siteRef":"r:153c-699a HQ",
"installed":"d:2005-06-01"},

    {"dis":"RTU-2", "equip":"m:",
"siteRef":"r:153c-699a HQ",
"installed":"d:2006-07-12"}

  ]

}
```

Note the granular encoded data values for `equip`, `siteRef` and `installed`. As far as JSON is concerned, they are just strings. But in fact they are another string-encoded data format that requires further parsing – and yes – a software library to decode their true meaning. This can be a real turn-off for a developer who just wants to read some data logs from a Haystack server into their Cloud app.

Making Haystack data more easily accessible is key to driving greater adoption of the standard. This is even more important as the standard grows in complexity and naturally becomes more difficult for people to consume. The standard needs to stay focused on the problems it's really trying to solve.

Enter HAYSON

Haystack now has a new JSON encoding format to help ease some of these woes. Originally nicknamed HAYSON, it started its life in Working group 792. Its aim is to provide a simple JSON encoding format for Haystack data that doesn't require a third-party library.

Here is a snippet:

```
{
  "id": {
    "_kind": "ref",
    "val": "whitehouse"
  },
  "dis": "White House",
  "site": {
    "_kind": "marker"
  },
  "area": {
    "_kind": "number",
    "val": 12323,
```

```
    "unit": "ft²"
  }
}
```

And if you love YAML like me (I hate writing JSON – I much prefer YAML):

```
id:
  _kind: ref
  val: whitehouse
dis: White House
site:
  _kind: marker
area:
  _kind: number
  val: 12323
  unit: ft²
```

Note how each granular value doesn't require parsing. Also note the above dicts are just simple JSON objects. In fact, you can take just about any JSON data structure and it will be valid Haystack data (lists are arrays, strings are strings, etc.).

Here's a more comprehensive example of our original grid using YAML:

```
_kind: grid
meta:
  ver: '3.0'
  projName: test
cols:
- name: dis
  meta:
    dis: Equip Name
- name: equip
- name: siteRef
- name: installed
rows:
- dis: RTU-1
  equip:
    _kind: marker
  siteRef:
    _kind: ref
    val: 153c-699a
    dis: HQ
  installed:
    _kind: date
    val: '2005-06-01'
- dis: RTU-2
  equip:
    _kind: marker
  siteRef:
    _kind: ref
    val: 153c-699a
    dis: HQ
  installed:
    _kind: date
    val: '2006-07-12'
```

Note that the _kind property is being used to denote the haystack kind (although dict doesn't need it because just about all objects are dicts). The underscore is used because tags cannot start with it. It also makes the property stand out when reading the data.

JSON Parsers

Modern web browsers and developer frameworks such as NodeJS, all come with fast native JSON parsers. Performance tests show that parsing a JSON grid is faster than parsing the equivalent data encoded in Zinc or Trio in JavaScript – especially for larger datasets.

Large Grids

One area that Zinc really shines in is the overall size of the document. A large Zinc encoded grid is typically smaller than one encoded in JSON. However, it should be noted that gzip encoding, typically used by all popular web servers, mitigates this problem.

Schema Support

Since JSON forms the backbone of just about all RESTful APIs globally, technologies have evolved around it to make working with JSON easier and more robust.

For example, let's say you're creating a REST API and want to validate the JSON data before it hits your end point. Or let's say you want error detection and autocompletion for your JSON/YAML data in Visual Studio Code.

The new JSON encoding has both JSON and Open API schemas to help with these use cases. The JSON schema has been registered with https://www.schemastore.org/json so all modern IDEs can use JSON/YAML validation for files named 'myFileName.hayson.json' or 'myFileName.hayson.yaml'.

Conclusion

The aim of the new JSON encoding format is to make Haystack data more easily accessible to all and to bring it into alignment with modern web standards and practices.

Thank you to all of you who have made the new JSON encoding format a reality! For more information please check out this repo and our working group.

https://github.com/j2inn/hayson

https://project-haystack.org/forum/topic/792

*in case you didn't know, a noob is defined as 'a person who is inexperienced in a particular sphere or activity, especially computing or the use of the internet'.

Gareth Johnson is the Senior Cloud Architect for J2 Innovations. Gareth has been in the building automation industry for over 20 years and was formerly the Core Architect of Tridium's Niagara Framework.

Believe it—Analytics software you can connect, configure & use all in a single day

It's not easy providing the most functionally-rich, easy-to-use and cost-effective building-analytics software the industry has ever seen, but **we've done it**.

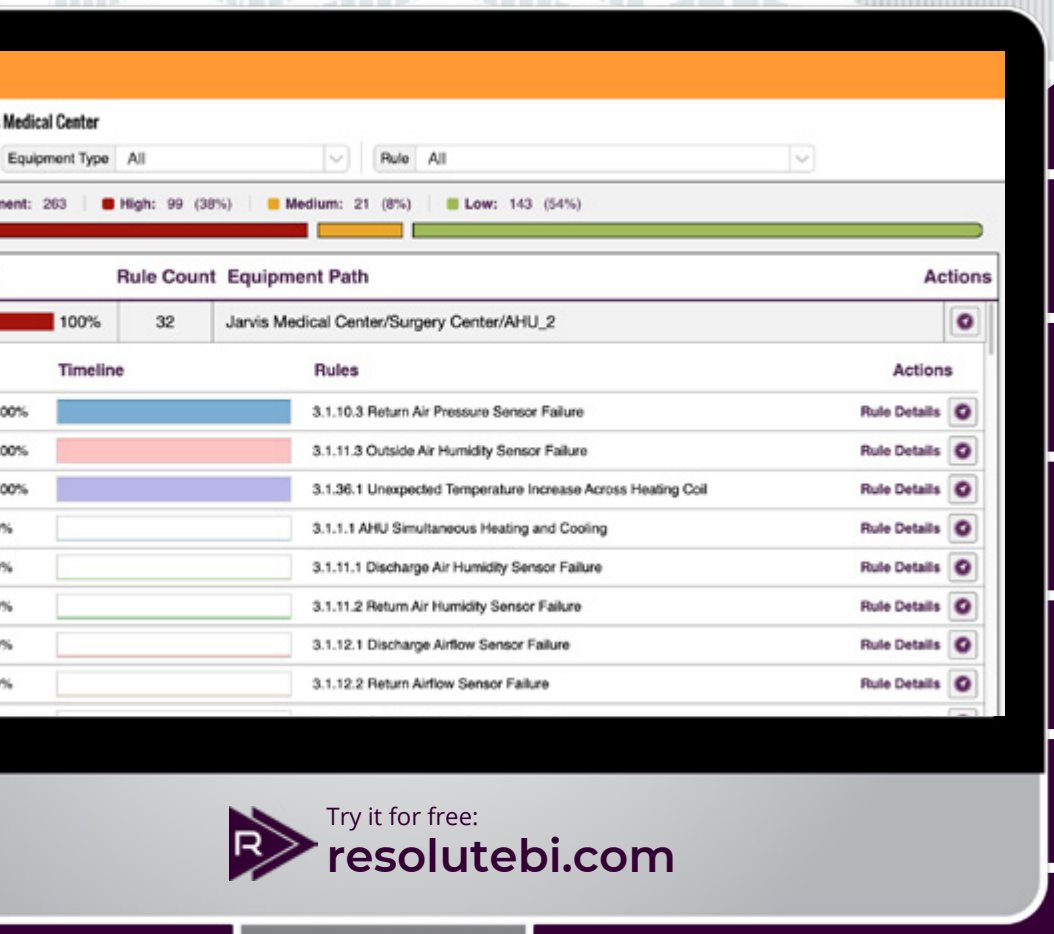

- Automated configuration
- Real-time fault detection & alerts
- Deep-dive equipment & energy analytics
- On-demand performance reports
- Root-cause analysis & prescriptive solutions
- Prioritized action items & tasks

Try it for free:
resolutebi.com

RESOLUTE™
BUILDING INTELLIGENCE

Haystack and Data Interoperability

"Interoperability between systems and data should not be a luxury that only a few have. We have the tools and functionality available and the need for these capabilities to be integrated within buildings has never been stronger. As a data modeling schema, Haystack's semantic modeling, ontology, tagging, naming conventions contextualization, and taxonomies provide a consistent, standardized methodology for naming and describing data points – all developed through a 10-year open-source community driven process."

The exchange and integration of data is something the built and IoT environments have been working on for the past ten years. However, exchange and integration are no longer enough.

At the core of every device and piece of equipment within a building or IoT deployment is an unending, diverse stream of data. With it, comes a fundamental need for better data identification, better exchange, better scalability, better security, and the one area that is often times overlooked—true, open, cost-effective interoperability. Different types of data will continue to shape the market, but the need for a standards-based, open-source methodology that enables data interoperability is the key to optimizing the use and value of data.

What is Data Interoperability?

In its most simple terms, data interoperability is the ability for devices, equipment, sensors and applications, regardless of whether they are made by the same manufacturer or by different manufacturers, to exchange, interpret and use data cohesively across different boundaries to meet the needs of a variety of stakeholders. Data interoperability addresses the ability of systems and services that create, exchange, and consume data to have clear, shared expectations for the contents, context and meaning of that data.

Haystack's Role in Interoperability

As a data modeling schema, Haystack's semantic modeling, ontology, tagging, naming conventions contextualization, and taxonomies provide a consistent, standardized methodology for naming and describing data points – all developed through a 10-year open-source community driven process. It permits the same defined data to be shared across any system, any equipment, any

application, or organizational boundary in order to advance the exchange, interpretation, and use of it.

Haystack provides the foundation to support open interoperability enabling one system to exchange data without manual effort being required to enable the receiving system to interpret the data. Haystack provides structural interoperability by defining the data structure of the information being passed between systems. And Haystack provides open semantic interoperability whereby multiple and disparate systems can exchange, interpret data, and use data to their fullest extent.

Summary

While making use of data is not new, Haystack helps create data ecosystems and enables data sharing and collaboration to be done efficiently, delivering better control, manageability, accuracy, efficiency, and reduced risk. It allows us to maximize the integrity and value of data and respond more effectively to market forces that are changing more rapidly than ever.

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

Tagging Niagara Station Components

TRIDIUM

"This article is about how Niagara users can configure tag dictionaries for easy reuse and faster overall deployment of a tagged project. It will cover how to imply tags and to deploy tags at-scale across projects. With the right tips and tactics, you can make light-work of the job of tagging your Niagara station components. "

It's time to embrace standardized tagging to save time, achieve better outcomes, and future-proof your work for data analytics, machine-learning, and AI. On Tridium's part, we are building more support into Niagara Framework® to make it easier and more intuitive for Niagara users to begin systematically tagging all their station components (Figure 1).

Niagara Framework users are a community of experts in the field of designing customized user interfaces for end-users of smart systems and smart buildings. Tagging is a core skill for accommodating the efficient resolving and rendering of graphical components in a customized UI, as well as for preparing a Niagara station for data analytics. With the move toward standardized point tagging pioneered by open-source semantic-web organizations like Project Haystack, the Niagara Community has moved away from bespoke labeling toward these standards. Today, Niagara visualizations can be driven by tag-based bindings with tags pulled from multiple custom tag dictionaries, as well as from the Project Haystack standard.

This article is about how Niagara users can configure tag dictionaries for easy reuse and faster overall deployment of a tagged project. It will cover how to imply tags and to deploy tags at-scale across projects. With the right tips and tactics, you can make light-work of the job of tagging your Niagara station components.

Direct Versus Implied Tags

All Tag Dictionaries in Niagara contain tag, tag group, and relation definitions. The tag and tag group definitions feed Niagara's HTML5 Tag Manager and Workbench Edit Tags dialog for applying direct tags and tag groups to components. Direct tag groups are an **n:tagGroup** relation from a component to a tag group definition. Figure 2 illustrates the result of adding a direct airflow standby set-point tag group to a component. Niagara's Haystack Tag Dictionary includes tag group definitions for each equip point grouping defined by Project Haystack.

Smart Tag Dictionaries can also include *tag rules* for implying tags, tag groups, and relations. Every direct tag, tag group, or relation uses a component slot, so using implied tags will save memory. Having tag rules centralized in the Smart Tag Dictionaries can make your tagging effort easier to maintain and update. Once you

Figure 1. Introduced in Niagara 4.0, Niagara entity query language (NEQL) made it possible to search for tags within Niagara stations. With every subsequent release, the Niagara development team has made tagging easier to comprehend, deploy and use. Support for Haystack smart relations – equipRef and siteRef – was introduced as of Niagara 4.6.

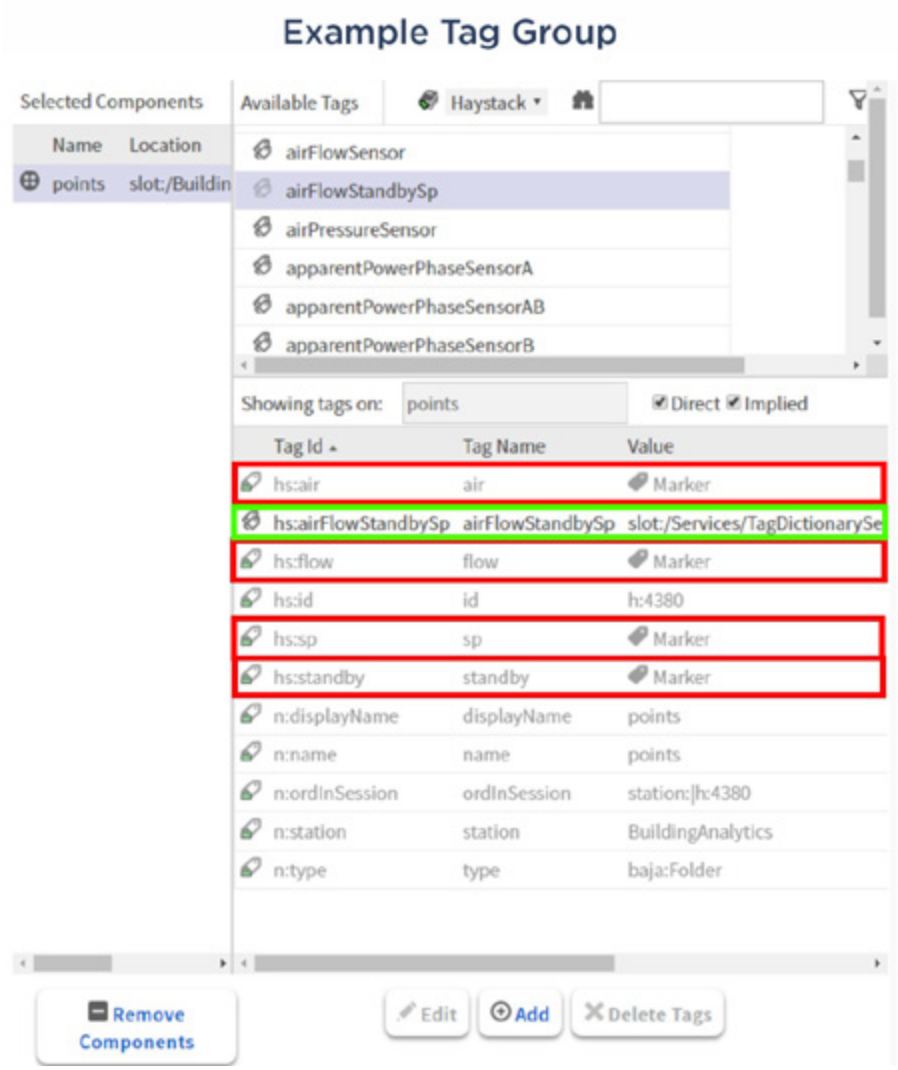

Figure 2. Highlighted in red in this HTML5 Tag Manager view are the tag group's tags. These are implied on the component that has a tag group relation to this tag group. Highlighted in green is an implied tag group id marker tag. Such marker tags are useful when searching a station for components on which this tag group is applied, instead of all components that contain the tag group's tags. A word of caution: if you delete a tag group definition or its dictionary, the relation will also be deleted.

have a working set of tag rules, you can drop them into other stations, and tags will start to be implied immediately.

The trade-off with tag rules is that they *do* need to be computed at run time. To minimize this effort, Niagara4 only evaluates the tag rules *necessary* for the tags being searched– not all tag rules in all dictionaries. However, you should understand this trade-off. Direct tags are faster for searching, but implied tags via tag rules are the way to go from the standpoint of storage efficiency and long-term maintenance.

Tag Rule Conditions

Tag rules are made up of tag rule conditions. Here are some simple ones:

- **'And'** condition is true if all child conditions are true. It is short-circuiting and will stop evaluating subsequent conditions, once the first condition returns false.
- **'Or'** condition is true if any child conditions are true. It is short-circuiting and will stop evaluating once the first condition returns true.
- **'Always'** condition is useful when you want to imply items to every component in the station, regardless of that component's type or any other tags on that component.
- **'IsType'** condition checks to see that the component is the specified type or one of its subclasses, before applying the tag.

The Niagara4 tag dictionary palette offers another class of conditions around 'BooleanFilter.' The basic 'BooleanFilter' applies an NEQL query to each component. If the query applies to the component, then that condition returns true. The 'BooleanFilter' has two subclasses:

- **'HasAncestor'** evaluates the NEQL query on the component itself, and then on its ancestors. As soon as it finds one ancestor that satisfies the query, then that condition returns true.
- **'HasRelation'** evaluates the query on the component itself and then on any component it can reach using a specified relation id.

'IsType' conditions can be evaluated quickly, thus it is best to put those conditions near the top of a tag rule. **'IsType'** will match the specified condition and

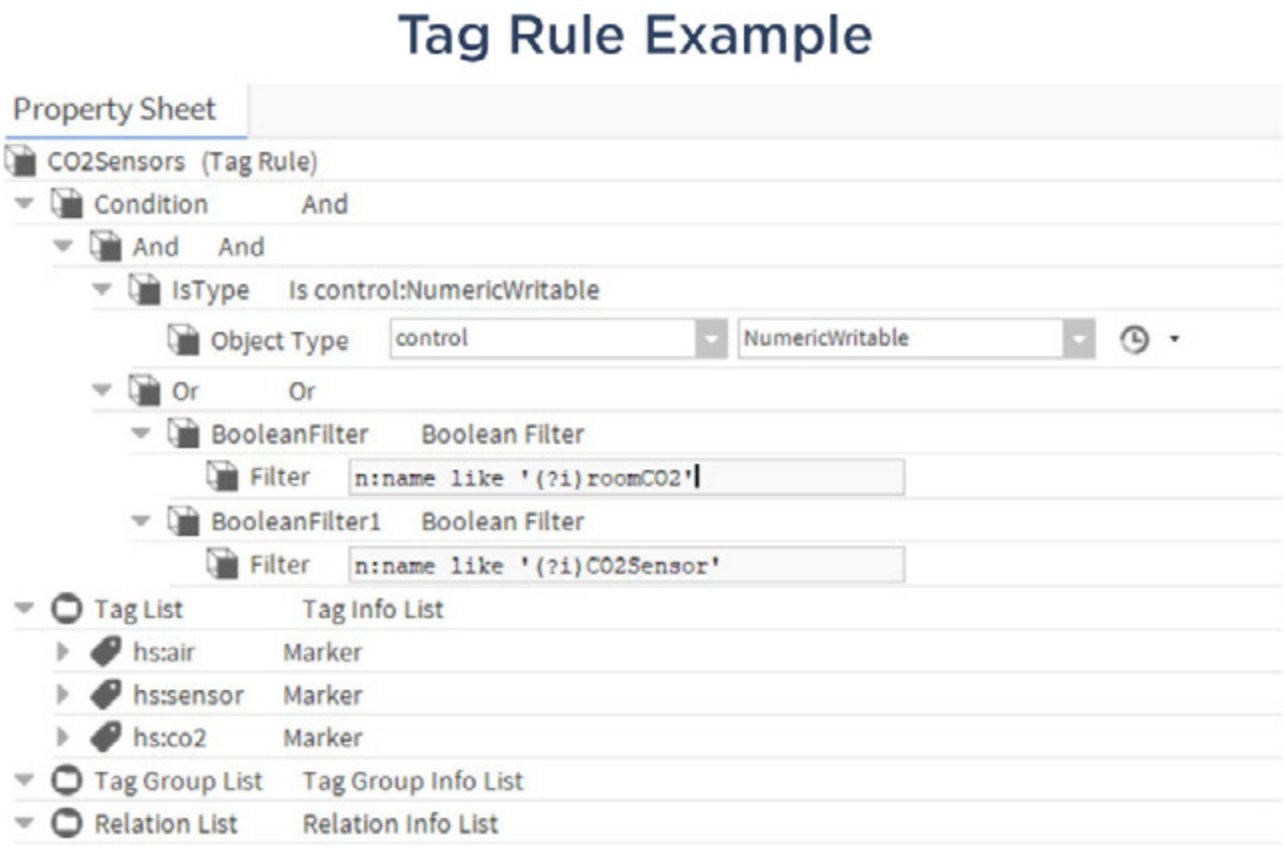

Figure 3. This tag rule has at the top an 'And' condition with two child conditions underneath it. The first is the 'IsType' so it's looking for numeric writables. If it finds a numeric writable then it will evaluate the 'Or' condition. The 'Or' has two Boolean Filters that search for and compare 'n:name' tags. If either of those filters apply to a component, then the implied tags in that tag list – air, sensor, CO2 – will be implied on the component.

subclasses. If you need an exact type match, you can use a **'BooleanFilter'** to apply an NEQL query that uses 'n:type' tags. For example, the query might be written: **"n:type = 'control:numericPoint'"**, if you want just numeric points and not numeric writables. Use the **'like'** operator to compare a string value to a regex expression. You can make that regex case-insensitive by including **'(?i)'** at the beginning. For example, you could form a query like **"n:name like '(?i)roomCO2'"**.

Figure 3 illustrates a Tag Rule example that uses a mix of these condition types in combination.Simple Versus Smart Implied Tags

There are two flavors of implied tags, and Haystack users can benefit from both. *Simple* implied tags have an implied value that is fixed. Examples include marker tags and value tags where the value is specified. Consider the Haystack tag **'hs:phase'**: if this tag is in a rule and it is set to **'BC'**, then when that tag is implied, it will always have the string value 'BC.' Likewise, **'hs:stage'** = 1.00 will always have the value of one.

The other flavor of implied tags are smart implied tags.

The following tags derive their value from the component they're implied on:

- **'n:name'**: the value of the tag is derived from the name of the component.
- **'n:type'**: the value of the tag is derived from the type spec of the component.
- **'hs:kind'**: from the Haystack Tag Dictionary, the value of this tag is derived from the type of point that it is implied on.
- **'hs:tz'**: from the Haystack Tag Dictionary, the value is derived from the time zone of the station.

Figure 4. In this Scoped Tag Rule Example, the ID of the Scoped Tag is 'scoped.' It is in a dictionary with the namespace 'my.' The effect of applying the **'my:scoped'** tag to a component named 'Ancestor' is to place the **'my:scoped'** tag on all of its descendants. The implied tags are marker tags because the **'my:scoped'** tag applied to this ancestor component is also a marker tag.

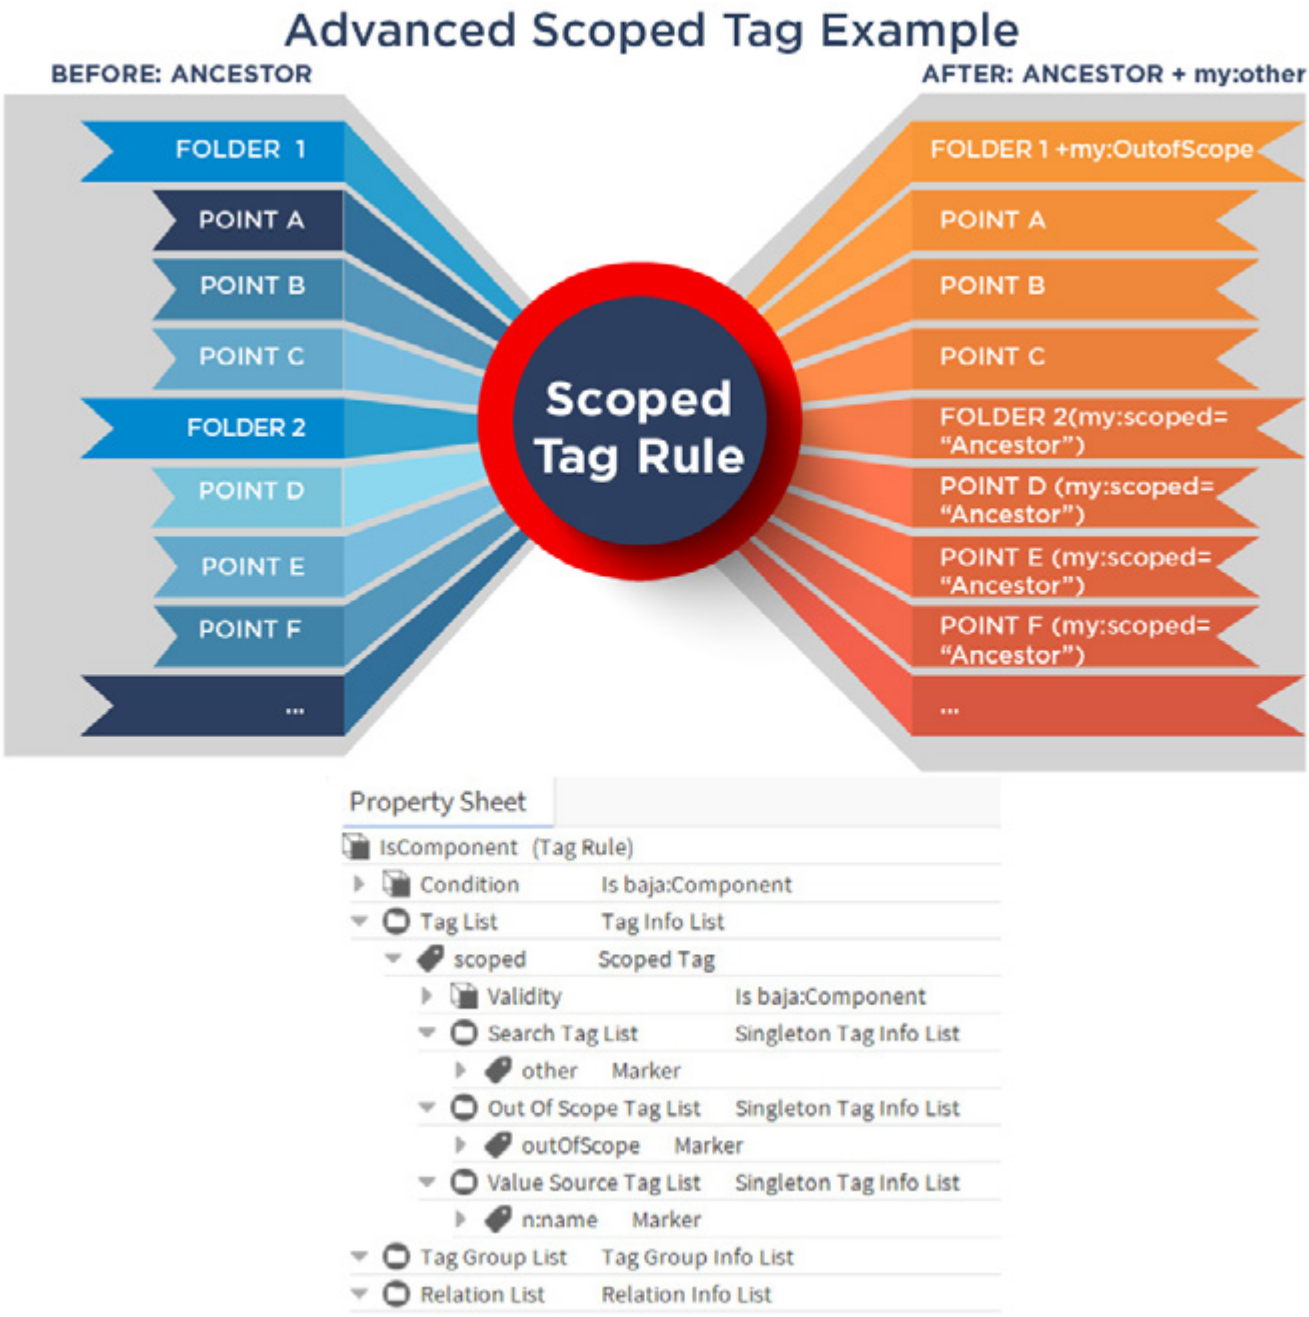

Figure 5. In this advanced Scoped Tag Rule example, instead of searching for a tag with the ID 'my:scoped', the rule searches for a tag with the ID 'my:other.' An out-of-scoped ID is specified: 'my:outOfScope'.. Instead of copying the value of the tag being searched for, the value of an 'n:name' tag will be copied. When a 'my:other' tag is applied to the ancestor component, Folder 2 and its descendants have the 'my:scoped' tag implied on them; the value of those implied tags is copied from the value of the 'n:name' tag, which in this case is the string 'Ancestor'. Because Folder 1 has the 'my:outOfScope' tag on it, the 'my:scoped' tag is not implied on it or any of its descendants.

The following tags may or may not be implied based on the component:

- **'n:input'**: tag will be implied only on non-writable points.
- **'n:output'**: tag will be implied only on writable points.
- **'n:hasPxView'**: tag will only be implied if the component has a PX view associated with it.

The following tags may or may not be implied based on the component and their value is derived from the component:

- **'n:history'**: tag will be implied if a component has an enabled history extension. The value of the tag will be the history ID defined in the extension.
- **'hs:enum'**: from the haystack dictionary, the value of this tag will be implied on Boolean and Enum points. For Boolean points, the value of the tag will be "false, true" or will use the "falseText" and/or "trueText" facet values, if those are present. For Enum points, the tag value will be set using the range facet value, if that is present.
- **'hs:maxVal'** and **'hs:minVal'** functions are similar to **hs:enum** in that they will resolve to the component's maximum and minimum facet values.
- **'hs:unit'** is based on the unit facet set for the component.

Scoped Tags

The 'scoped' tag is another type of *smart implied* tag. The simplest use case for a 'scope' tag is to imply a tag based on whether a given component has an ancestor with a tag that has the same ID as the scoped tag. The value of that implied tag will match the ancestor tag's value. Figure 4 shows a tag rule built using a scoped tag.

The Niagara4 tag dictionary palette has some advanced options for scoped tags as well: You can search for a different tag ID rather than matching the ID of the 'scoped' tag. You can copy the value of a tag other than the one being searched for. Finally, you can specify an '*out-of-scope*' tag ID. If a tag with that ID is found before finding the tag being searched for, then the scoped tag will not be implied (Figure 5).

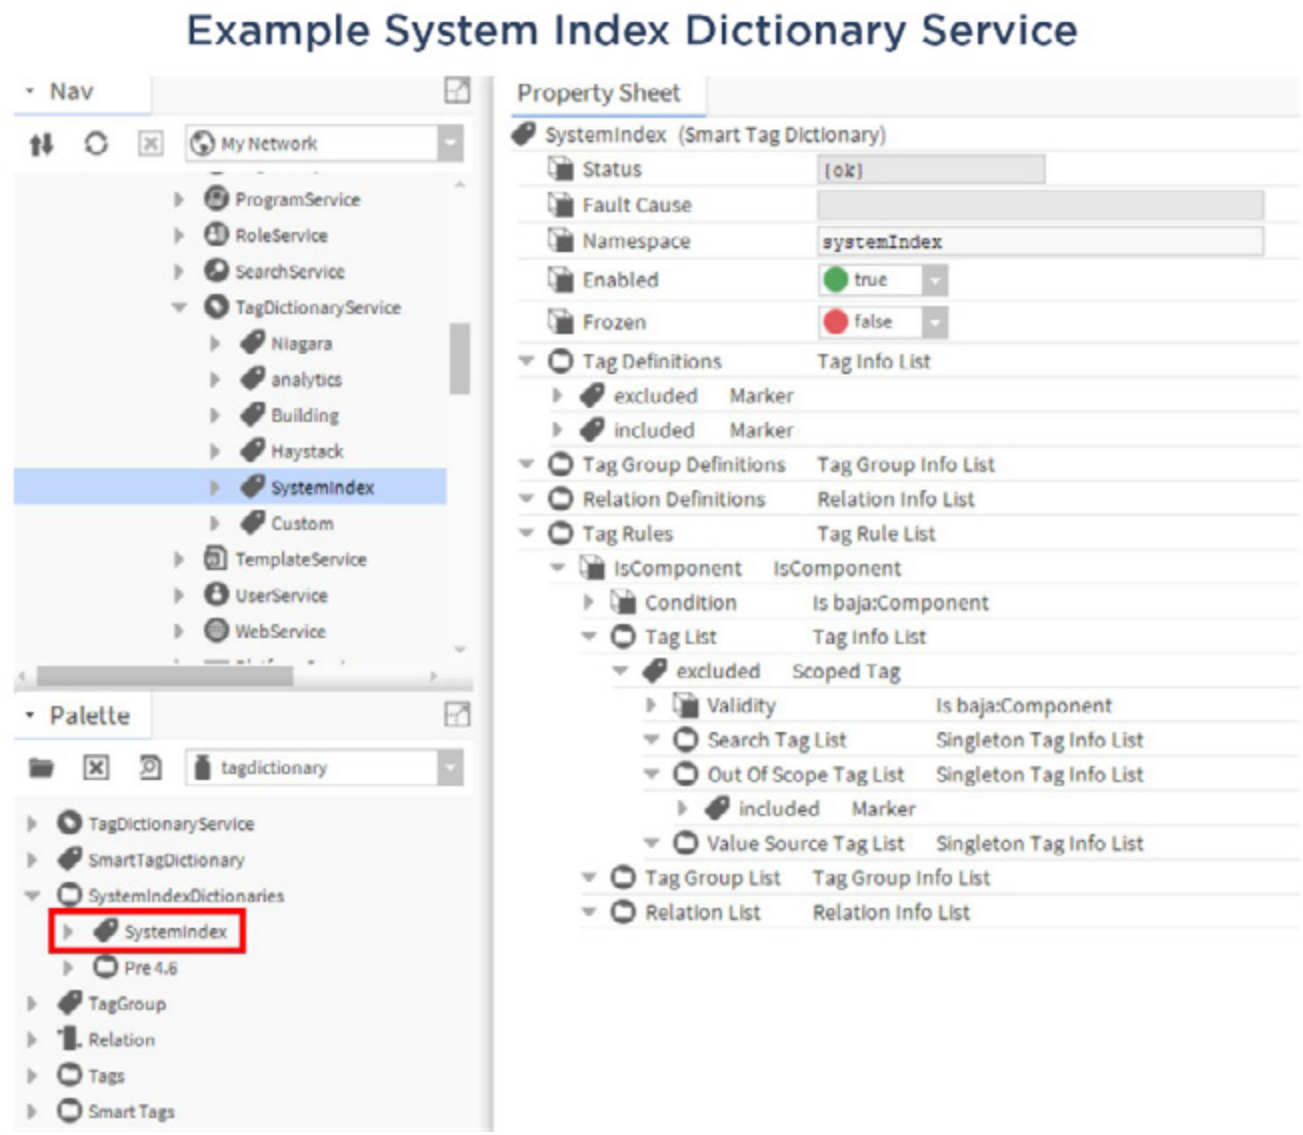

Figure 6. When preparing for a system index, you can apply the 'excluded' marker tag to the top of that part of the tree you want to exclude, and the 'excluded' tag will be implied on all descendants. The system index will ignore those components. If you want to run a system index that includes a subtree of the tree that you are excluding, you can put the **'included'** marker tag on the top of that section.

Scoped Tags for System Indexing

One use case for scoped tags in a Niagara station is to support system indexing. Under the tag dictionary palette, there is a system index dictionary that you can drag into your tag dictionary service. It contains two marker tags: **'excluded'** and **'included'**. The tag dictionary also includes a tag rule that implies an **'excluded'** tag. By default, components with the **'systemIndex:excluded'** tag are not indexed.

In addition to smart implied tags, the Niagara dictionary palette also offers *smart implied relations*. The **'n:child'** and **'n:parent'** smart relations help to navigate a component tree. If you're on a driver network, for example, you can use the **'n:childDevice'** relation to get all the child devices under that network. If you are at one of those devices, you can navigate to the network by traversing the **'n:parentNetwork'** relation. On a device, you can use the **'n:childPoint'** and **'n:childNullProxyPoint'** relations to get from the device to all the child points. The **'n:parentDevice'** relation can get you back from a child point to the parent device it belongs to.

Niagara4's Haystack dictionary offers **'hs:equipRef'** and **'hs:siteRef'** relations. These smart relations will be implied between any non-null proxy points and an ancestor that has the **'hs:equipRef'** tag on it. If that equip component also has an **'hs:siteRef'** relation to a component with the 'hs:site' tag, the **'hs:siteRef'** relation will be implied from those non-null proxy points to that equip's site. These smart implied relations are time savers when building a station by avoiding the tedious addition of many direct relations.

Tag-based PX Graphics and System Database

Niagara *tag-based* PX graphics use all the tag and relation information that one can build into a station as described above. Introduced in Niagara 4.9, *tag-based* PX graphics use bindings based on NEQL queries. Bound components must have the expected tags and relations; however, they don't need to have the same names, nor do they need to be located at the same place in the component tree. Niagara will find them just based on tags and relations. NEQL traverse queries can be used to create bindings to anywhere in the station using relations. The great advantage of *tag-based* PX graphics is that they are easier to reuse across stations.

Another feature in Niagara that uses tag and relation information is System Database. Niagara *System Database* (systemDB) is our name for the stored result of System Indexing – the periodic update of the tag and relation information for selected entities from each station in your Niagara network. By default, Niagara indexes all networks, devices, points, schedules, point/device folders. A System Index will also pull up any components with a PX view using the new **'n:hasPxView'** implied tag. SystemDB is currently single-tier, and it enhances any features that use NEQL queries. You can search all the stations in your Niagara network and build hierarchies against the results, given that they have been indexed up to systemDB. As of 4.10, if your Niagara Supervisor has a systemDB and you've indexed your stations to it, you can leverage virtual *tag-based* PX graphics against that system database.

Summary

Since the Niagara Framework is tag-agnostic, users have the utmost flexibility when applying tags to Niagara stations and building systems. The result is visualizations that are reusable and manufacturer- and equipment-agnostic. Tag-based bindings streamline on-site efficiency, save time and provide more flexibility when deploying graphic templates on new and existing Niagara stations.

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.

Therese Sullivan is Director of Marketing at Tridium. Therese has written extensively on the convergence of IT/OT and the Buildings IoT. She is a former editor of Haystack Connections and is a Contributing Editor for AutomatedBuildings.com.

Contributed Article

The Importance of Data Standardization for Retail and Hospitality

J2 INNOVATIONS
A Siemens Company

"Project Haystack provides a great solution to this problem; adding meta data, in the form of semantic tags, enables analytics software to compare all the sites irrespective of what controls system or protocol is used. The challenge is implementing Haystack tags in a cost-effective way."

We are all familiar with "chains" – companies that operate multiple shops, restaurants, or hotels across one or more countries, with a standardized proposition for the consumer: a consistent brand experience in terms of style and quality no matter which site they visit. When it comes to managing the buildings of these sites, however, there are a variety of situations that are anything but standard!

Chains most frequently lease or buy buildings (or parts of buildings) that have already been built and therefore come with existing HVAC equipment. Whilst most chains have standardized approaches to fitting out new sites, their equipment procurement policies change over time, and it is frequently not economically viable to rip out all the equipment that is already there. The diversity is compounded because such companies often grow by acquisition, so inherit multiple sites with a different equipment fit-out. The result is often a wide range of HVAC equipment and their associated controls as well as other services that need to be managed. As more and more of such buildings are being managed remotely from Head Office, or an out-source facilities management (FM) provider, there is a big problem with how best to standardize the data coming from each site. Comparative analytics can be used to identify and prioritize opportunities to improve efficiency and reduce costs. But in some countries, there are also regulatory compliance that requires companies to report on their total energy use, carbon emissions and set targets for reductions.

Project Haystack Provides the Solution

Project Haystack provides a great solution to this problem; adding meta data, in the form of semantic tags, enables analytics software to compare all the sites irrespective of

what controls system or protocol is used. The challenge is implementing Haystack tags in a cost-effective way.

Retailers, restauranteurs, and hoteliers are well known for being extremely cost-conscious when it comes to "behind the scenes" investments that are not directly seen by their customers, so deployment costs for any solution need to be low.

The current generation of controls typically installed to manage such buildings have not been designed to support this sort of data standardization. Whilst they may support the BACnet open protocol, this was never designed for communicating data across the internet and does not support the implementation of a full data model or semantic tagging. Even the new breed of IoT style solutions which either propose wireless devices separate from existing controls, or extract data from existing systems are generally doing this in a non-standardized way.

A New Generation of Controls

What is needed is a new generation of controls solutions which properly support Haystack tagging, with software tools that streamline the tagging process, since this will significantly reduce costs compared to other options. Currently the cost of implementing tagging on existing sites using the few solutions that do support tagging can be a barrier to adoption. Some good news for the widespread adoption of Haystack across existing building portfolios is that computer hardware costs have fallen to a level that the cost of embedded hardware platforms to run such Haystack compatible software on site is no longer a deterrent to deployment. This is especially if the software also provides secure edge to cloud connectivity since at a minimum, even for a non-Haystack solution, some sort of gateway and modem would be required. There are several vendors who offer Raspberry Pi type controllers which cost less than $150. Another feature necessary for cost-

effective multi-site deployments is the ability to remotely configure the solution. If the cost of sending a skilled controls specialist technician to site can be avoided, then commissioning costs can be significantly reduced.

A further way to reduce implementation cost is through the use of templates; this enables the tagging process to be completely automated provided the controls installers know in advance the equipment to be managed, since then a template can be created for each device type that pre-tags all the data points.

Once all the site data has been Haystack tagged it can be visualized on dashboards and analytics rules can be applied across all the sites to identify trends and highlight outliers. This empowers the facilities management team in their prioritization of corrective actions, and decisions about future investments to reduce energy and operational costs.

In short, multi-site operators who want to improve the performance of their portfolios, need to standardize their data using Haystack to enable comparative analytics and optimization of each site's performance. This will require the use of a new generation of software at the Edge that fully supports Haystack tagging and secure remote connectivity, so enable easy remote management and to get the site data to cloud-hosted analytics applications.

Chris Irwin joined J2 Innovations to develop sales in Europe, the Middle East and Asia. Chris comes with a wealth of experience in the building automation market and with skills in strategic business development and strategic marketing..

Building System Threats from Attackers, Bad Employees and a Lack of Understanding

INTELLIGENT BUILDINGS™

"In all three scenarios forensic data was either not available or was of little to no use. Project Haystack is ready-made to help provide this information by creating tagging that identifies points to be collected for post event forensics and better yet, analyzed real-time to possibly identify nefarious activity to prevent negative impact to a control system."

Building Cybersecurity has become a "real thing". And that is good, but this too introduces issues. In this article I will give real-world examples of three incidents that Intelligent Buildings has firsthand knowledge. One of which was an attack and two with good intentions and a lack of understanding.

In all three scenarios forensic data was either not available or was of little to no use. Forensic data is key to understanding what and how it happened and provide information to aid teams to prevent events like these in the future. Project Haystack is ready-made to help provide this information by creating tagging that identifies points to be collected for post event forensics and better yet, analyzed real-time to possibly identify nefarious activity to prevent negative impact to a control system.

These stories are based on actual events with modifications to mask the company that the incident happened. No company or people names are used.

92 Days to Recover

Event Type: Hack

A day before this event, a building engineer checked his personal email on the application server. He received an email that appeared to be from a fellow engineer. The building engineer that checked his email thought it was unusual that the employee had sent an email to him to his personal email, but the email had a link to a site that appeared to be something that might have come from this employee. The link did not take the engineer anywhere. It just appeared that the other employee had not copied the link correctly. The building engineer talked to the person that sent it later in the day and told them the link he sent didn't work, for which the person let him they had not sent him anything, and if had it would have gone to his work email because he didn't know his personal email. None of these inconsistencies caused the engineer to notify anyone about what happened. They did, however, back up the application server to an external hard drive just in case. The day ended without incident.

The following day about mid-morning, another engineer needed to make setpoint adjustments because tenants on the 10th floor complained that their area was too cold. The engineer went to the application server to make the adjustments and noticed that a window was opened he had never seen. There was a message that said, "Your Important files are encrypted. Many of your documents, photos, videos, databases, and other files are no longer accessible because they have been encrypted. Maybe you are busy looking for a way to recover your files but do not waste your time. Nobody can recover your files without our decryption service." There were also instructions on how to recover and how to pay for the recovery. There was also a countdown timer letting them know that all their files would be lost.

The engineer that clicked the link the day before was also in the room at the time and let the other engineer know that he had made a backup the day before and that they should call the vendor and have them reinstall the operating system and the application. The vendor was able to reinstall the operating system and installed the application and the files necessary to get the application up and running. This took a couple of days.

Not too long after the system was back up and running, there were some anomalies that occurred with some of the equipment in the central plant. The variable frequency drives (VFDs) seemed to be sometimes running slightly faster and sometimes slightly slower; however, this settled out and did not occur for the rest of the day.

Several days later, staff arrived and went about their typical day. The first indication of a problem was the engineer noticed that the central plant controllers were offline. They also noticed no alarms were showing for the central plant which there should have been at least the alarms associated with the offline state of the controllers.

Investigation of the plant found the chillers were off. Further investigation found that several of the VFDs were inoperable. The staff attempted to restart the main chiller manually but were not able to. They next tried to restart the other chiller manually and were unable to start it either. They began investigating mechanically the cause, at which time they found extensive damage to the pumps due to what they believed was cavitation.

At this point, the association between the ransomware attack was not even considered.

It was determined that several of the pumps were going to need to be replaced, and each of the chillers required to be fully inspected to find out if there was any damage to them. The controls vendor was also contacted to investigate the system to check the application out and get the controllers back online, and determine why alarms did not go out.

The staff determined that several of the pumps needed to be replaced and several of the VFDs. They found that the VFDs appeared to have been run beyond their critical speed. The vendor happened to be listening and informed the engineers that the only way to happen is if someone disabled the safety feature that would prevent it.

This changed the thought process from possible equipment failure (not likely due to the number of devices affected and the extent) to the system's possible hack. The direction of the investigation took the focus that this was probably what happened, but how?

The first thought was that another unrelated hack had happened, but this seemed unlikely because lighting doesn't strike twice in the spot, right? The vendor went through the backup made the day before the attack and found that infection occurred that day of the backup, and it most likely came from the email that the engineer clicked when he checked his email on the application server.

Now, most of what is stated here are strictly speculation because there was no forensic data. Because the system had been reloaded and no logs were retained, who and when they accessed the system and what was done is still unanswered to this date.

Here is what is believed to have occurred.

- The vendor found that a RAT or remote access trojan was also in the payload of the ransomware
- Once RAT (remote access trojan) was embedded and the hacker or hackers planted the ransomware
    - There most likely beacon that notified the hackers that the RAT was installed
- When the system was reloaded, the beacon notified the hackers that it was active again, and they realized they would not get paid, so they decided to damage parts of the system
- They remotely access the system and played with the VFDs to see what they could do but waited until they felt no one was watching the system and entered it sometime after 7 PM (the VFD access is via BACnet and does not require a username and password)
- They disabled the alarms
- They attacked the VFDs, which most likely caused cavitation, and destroyed the pumps
- They bricked the central plant controllers (the plant controller access is via BACnet and does not require a username and password)

The Results

The damage to the system included:

- Roughly half the VFDs needed to be replaced
- The chillers had to be dismantled somewhat to inspect for damage
- Several pumps had to be replaced due to the damage from cavitation
- The central plant controllers had to be replaced because hacker rendered them unrecoverable
- A new PC for the application was purchased to ensure no residual infection was present
- Because the backup was corrupted and the only backup was a year old, the application host programming had to be redone
    - Not from scratch, but the system had had additional devices and programming done in the past nine months due to expansions and revisions
- The entire process to fully recover took 92 days and thousands of manhours
    - Space cooling and heating units had to be rented to maintain tenant comfort during this time

Hackstack: In this event there were anomalies that could have been identified during the events mentioned such as the unusual VFD and pump activity.

Connectivity Lost to Over 100 Buildings

Event Type: No Policy/Lack of IT

A company embarked on a nationwide initiative to aggressively improve the cybersecurity of all its facilities. IT began the process and spent two years and several million dollars trying to understand how systems were set up, what systems were and were not connected if there were any policies and procedures in place, who was responsible for what, how to manage systems, and inventory all the devices in facilities. IT learned the basics of managing users but not the full extent or implications of doing such.

As a rule, the organization did not have internal resources with a broad-reaching understanding of the various systems and what it took to manage them. However, in one region, there was an employee that was the go-to guy. This individual was responsible for the most significant number of sites. He had been instrumental in creating a regional system that centralized command and control of over 200 sites. This employee was given free rein to implement and change as he saw fit to create a unified and standardized system across all the sites within his realm of responsibility. His work appeared to be a model that all other regions to follow. However, due to the trust and control, he was given, the organization did exercise oversite of this individual, so the work this individual did was not documented, nor were the details of how the systems were configured.

This individual had to be let go. IT was notified, as is policy, to remove the employee from all system access (business systems like email and network access). As mentioned earlier, IT had learned the basics of account management of control system devices. Still, IT had not learned or understood the implications of removing users without understanding the roles and underlying functions of what the user account influenced.

IT began removing his user from the controllers and the application server. Unbeknownst to anyone, his user was also the user that created the machine-to-machine connection between the controllers and the application server for command and control. Before anyone was aware, over 100 sites lost communication to the central application server, which meant that the interface allows facility management to access any of the sites.

Because no one knew this employee's password and for very valid legal reasons, the employee could not be contacted; simply putting another user in place of him would not fix the issue.

What was required was to work with the manufacturer of the controller to recover communication between the application server and the hundreds of controllers affected. There were hundreds because for every site, there could to anywhere from one to 20 controllers at each location. This meant there were easily over 500 controllers that needed to be "touched."

The manufacturer was able to re-establish communication to all the controllers after 6,000 man-hours of work. This did not include the hours upon hours of work required by facility staff to manually control the 100-plus site until centralized control was restored. This number has yet to be determined but could easily be in the thousands of hours.

Haystack: The employee's user should have never been used in the address book. While not a normal practice, tagging users would help monitor for things like idle users, incorrectly configured users, and in this case, comparing the address book user to system users would have flagged this and notified admins of the compliancy violation at the start which could have prevented this incident. Not to mention educating IT on the potential impacts of their actions.

Over 6,000 Devices Knocked Offline

Event Type: Policy Enforced/Lack of IT Understanding

A large commercial real estate building of over 100 floors was built with all the latest Smart Building technology and Cybersecurity for IT and OT was built into the foundation of all these systems. To further insure that cybersecurity specification were met, they hired Intelligent Buildings to cyber commission over 10,000 devices in the facility.

The systems had been fully commissioned and were effectively turned over to the owner. The HVAC vendor was making final adjustments to the system and noticed that they could not connect to a device they had been working on just before lunch. They began checking other devices and found that none of the devices they checked were responding. The technician went and directly connected to a device nearby, and it was unresponsive.

The tech power cycled the device and was able to connect to it directly. The tech then tried connecting to the same device over the network and was able to connect. However, none of the rest of the devices throughout the building were responsive. This number was over 100 devices. This did not include the field devices.

In the meantime, the vendor responsible for the power monitoring noticed that the devices monitoring the racks were unresponsive. The number of devices totaled over 2,000. The vendor tech had to climb a later to investigate. The tech power cycled the device, but in this case, they had to connect via a serial cable to confirm that the device came back online and ensure those configuration parameters that the vendor set was still in the device.

At this point, other vendors (lighting, elevator, etc.) noticed they had unresponsive devices as well.

The vendor began reaching out to the GC and Project Management to inform them the systems were now ***not*** ready for occupation. The discussion started as to what was to be done. Delay occupying the building was on the table but not really an option legally.

Each of the vendors had been working independently and were not aware there was an issue across other systems. So now the question was what happened?

Vendors started contacting IT to see if they could identify a cause, if there was one, from an IT perspective. IT did not see anything that they knew could have caused the issue.

Light bulb moment…

The cyber commissioning company was contacted to determine what might have happened. CCC found out about the vulnerability scan. They asked IT about the scans that were used. The type that was used is a known device killer. It effectively acts as a DoS attack on the devices. The devices cannot handle the interrogation this type of scan performs.

The vendors had to go to each device and manually reboot their devices one at a time. The vendor of the data center rack power monitoring devices had to climb a ladder over 2,000 times because each device had to be manually rebooted and connected. This vendor made the statement, "if this happens again, we'll give you a ladder, laptop, and serial cable, and you can do it yourself.

It is believed that the cost of this one scan cost in the high six figures.

Hackstack: Unfortunately, there is not much that tagging could have done to prevent this. However, post event, having the forensic data categorized with tagging could have shortened the length of time it took to identify the source of the incident.

The Project Haystack Cybersecurity Working Group has a journey ahead of it to get cybersecurity embedded into the fabric of our industry. But it is worth the effort.

If you are interested in joining this effort, please let me know. We'd be happy for you to join us.

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

Fred Gordy is an industry expert of cybersecurity for building control and power monitoring systems. His control systems knowledge gives him insight on challenges of interlacing traditional IT environments with control systems for a cohesive and secure operational technology platform.

Project Haystack 10 ANNIVERSARY

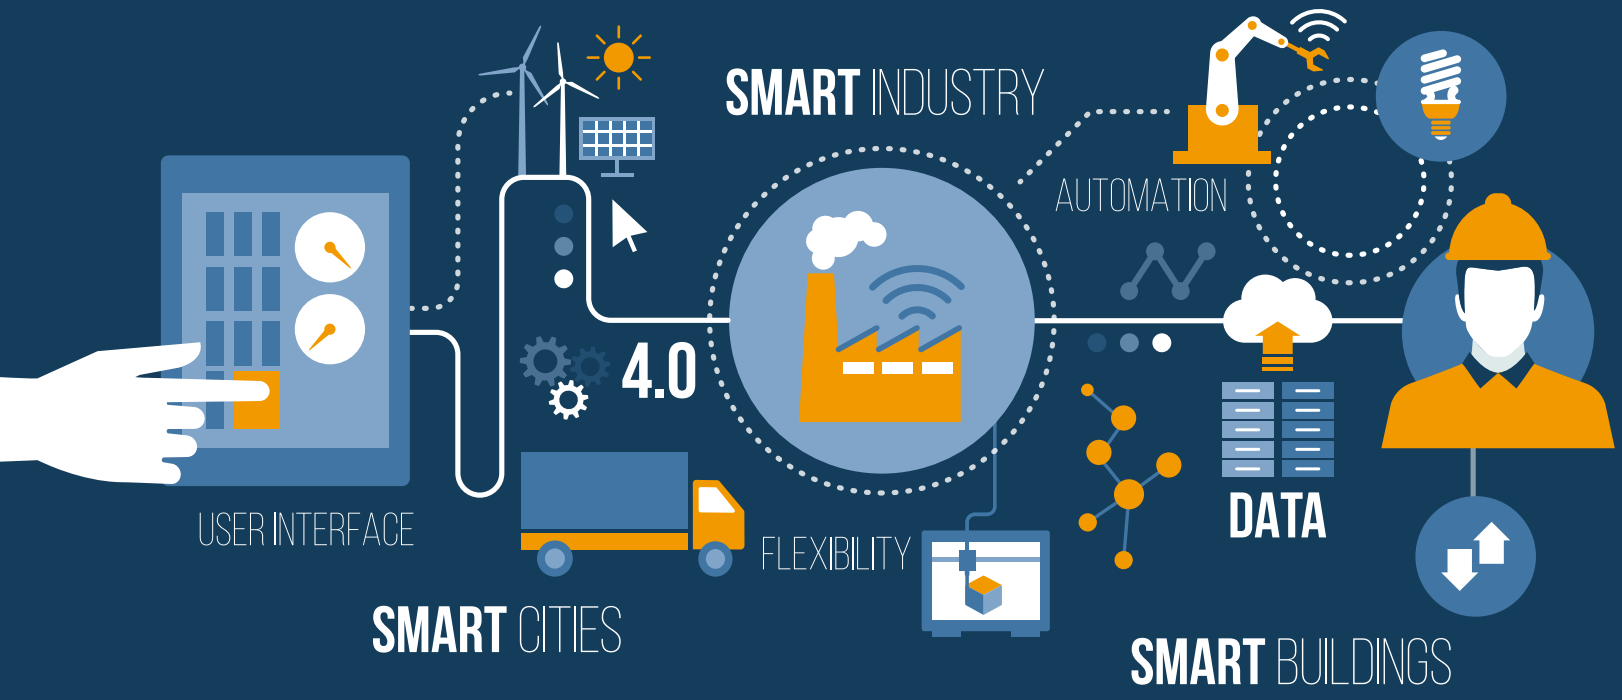

A Look Back at 10 Years!

Finding the Needle

The emerging field of analytics is the key to turning the data from smart systems into actionable intelligence.

We finally got our wish! Modern control systems and smart devices give us access to tremendous amounts of data – environmental conditions, energy use, equipment operation and many other facets of building operations. In this data are the keys to better building performance – trends, correlations, exceptions, deviations, control loop inefficiencies, and equipment faults – but how do we find them?

The emerging field of analytics is the key to turning the data from smart systems into actionable intelligence. Analytic tools – software that has the capability to automatically analyze building, energy and equipment data to find patterns of interest – can uncover control issues and equipment faults showing us the path to better facility operation.

Data Meaning – Semantics

In order to begin the analytics process we have to know the meaning of our data. For example, if we get a data item from a BAS and it has a value of 77.6 we can't do any effective analysis until we understand whether it is 77.6 degrees F, or PSI, or RPM, or kW, etc. "Units" is one good example of meaning that we need to begin the analytics process, but it is by no means the only one.

Continuing with our example, perhaps the point with the value of 77.6 is named zn3-wwfl4. If I am familiar with the system and the naming conventions used when it was installed I may be able to determine that means Zone 3, West Wing, Floor 4. Now I have a bit of information to work with. If I know the building well I may also be able to tell that zn3-wwfl4:

- Is a zone temperature
- Is an exterior zone
- Is south facing
- Is supplied by a VAV box
- Is served by AHU-1
- Is operated on occupancy schedule #1 which is 7:30 AM - 6:30 PM
- Has an occupied cooling setpoint of 74 degrees F

This "data about data" is called meta-data. This meta-data enables me to understand the impact of the current value of 77.6 – I can now see that it is over temperature during occupied hours and the occupant is probably getting uncomfortable. Without the necessary understanding of the data I can't determine the impact of the current value and its relationship to proper system operation. In order to provide effective analytics I need to "map" this meta-data to point zn3-wwfl4.

Interestingly, with all of the power they have gained over the last decade, most building automation systems provide poor semantic modeling of the operational data they contain. The systems provide us with a name and a value but little other information about the specific item. The result is that a labor intensive, process is typically required to "map" the data before any analytics can begin.

Mapping Meta-Data

So how can we capture all of this information and associate it with the data items in our automation systems and smart devices? We cannot do it simply by trying to use standardized point names. Clearly, in even our simple example we have more data that can be effectively embodied in a point name. Add to that the fact that we may want to add numerous other meta-data items over time and it's obvious we need another approach. There are a number of elements to an effective solution.

1. Separate the point name from the representation of meta-data. Use tags to represent the meta-data and associate those tags with the point name to provide the semantics that will describe the point.
2. Utilize a standardized library of tags to provide consistency of meta-data terminology, which will enable automated tools to interpret data meaning.
3. Follow a consistent naming convention for data points so that the mapping of tags can be consistently applied and automated.

Given our earlier example, a record representing the point with its associated meta-data might look like:

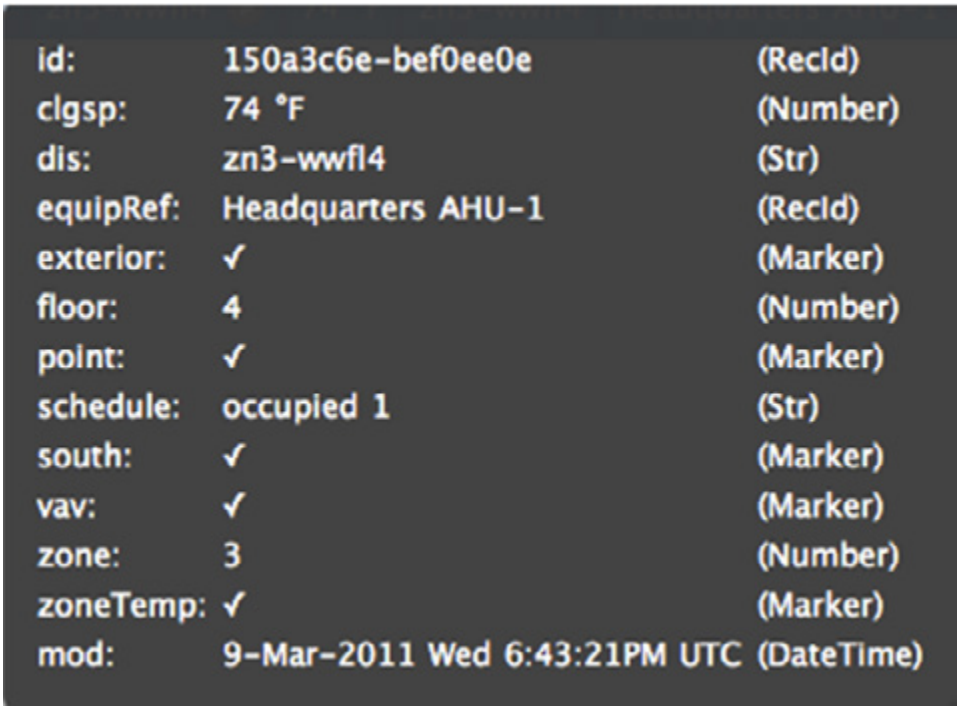

Project Haystack

Helping the industry move forward with meta-data and naming conventions is where Project-Haystack comes in. Project Haystack is an open source initiative to develop tag naming conventions and taxonomies for modeling of building equipment and operational data. The project is developing standardized data models and tag libraries for sites, equipment, and points related to energy, HVAC, lighting, refrigeration and other environmental systems. Substantial libraries of tag names and proposed taxonomy models are already in place.

A key factor in the need for Haystack is the reality that we have millions of points in thousands of systems out there and those point names are not going to change. It's simply not an option – and it isn't necessary. What we need is a standardized model for applying meta-data to enable us to associate meaning with those point names. Project Haystack will provide this.

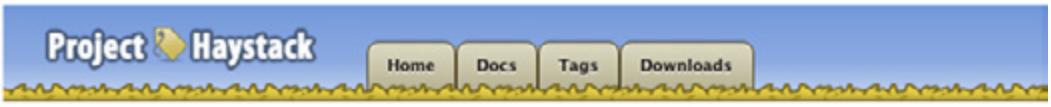

ahu	Air Handler Unit which heats and/or cools air.
ahuRef	Associate an entity such as a vav with an ahu.
area	Floor area of a site measured in ft^2 or m^2.
boiler	Boiler equip used to generate hot water or steam for heating.
boilerPlant	Group of one or more boilers.
boilerPlantRef	Associate an entity such as an ahu or boiler with its boilerPlant.
chiller	Chiller equip used to generate cold water for cooling.
chillerPlant	Group of one or more chillers.
chillerPlantRef	Associate a record such as an ahu or chiller with its chillerPlant.
coldDeck	Associated with the cold air ductwork in a dualDuct AHU or VAV.
cool	Cooling coil as bool or numeric point used with ahu equip.
dis	Short display name for an entity.
dischargeFlow	Discharge air volumetric flow point of an ahu or vav measured in "cfm" or "L/s".
dischargeHumidity	Discharge air humidity point of an ahu or vav measured from 0% to 100%.
dischargePressure	Discharge air static pressure point of an ahu or vav measured in "inH_2O" or "kPa".
dischargeTemp	Discharge air temperature point of an ahu or vav measured in °C or °F.
dualDuct	Indicates separate hot air and cold air ductwork for ahu or vav.
elecKw	Electrical demand point measured in "kW" for elecMeter.
elecKwh	Electrical consumption point measured in "kWh" for elecMeter.
elecMeter	Equip which meters electricity.
elecMeterLoad	Equip or point which consumes energy as electrial load.
equip	Equipment asset.
equipRef	Association with an equip entity.
exhaustDamper	Exhaust air damper point for an ahu measured from 0% to 100%.
exhaustFan	Exhaust air fan point used with an ahu.
fan	Supply fan point used with ahu and vav equip.

Who Should Participate

Project Haystack encompasses the entire value chain of building systems and related intelligent devices. Owners and consultants can specify that Haystack conventions be used in their building automation systems to ensure cost effective application of analytics tools (and other software applications), and management of their buildings for years to come. System integrators and manufacturers who integrate Haystack support into their projects and products are positioned for the future of value added services and can streamline the process of turning data into actionable intelligence.

Response to the project has been overwhelmingly positive since its launch in early March. Building owners, equipment vendors, ESCOs, systems integrators, and Governmental organizations are all coming together to help move this essential project forward.

Pragmatic use of semantic naming conventions and taxonomies are essential to make analytics more cost effective and to enable us to and derive value from all of the operational data we now have access to. Please consider joining the effort – its open to everyone interested in the continuing journey of building efficiency. You can check it out and learn how to participate at **http://project-haystack.org**.

John Petze, C.E.M., Partner, SkyFoundry

John Petze, C.E.M., is a partner in SkyFoundry, the developers of SkySpark™, an analytics platform for building, energy and equipment data. John has over 25 years of experience in building automation, energy management and M2M, having served in senior level positions for manufacturers of hardware and software products including Andover Controls, Tridium, and Cisco Systems. At SkyFoundry he rejoins Brian Frank, co-founder and chief architect of Tridium's Niagara Framework, as they look to bring the next generation of information analytics to the "Internet of Things".

Something Happening Here

Haystack Connect Event.

The theme song for the first ever Chattanooga http://haystackconnect.org/ event was Buffalo Springfield's - For What It's Worth playing on the words; "There's something happening here What it is ain't exactly clear." The song gets a little darker after that but the not the event. It continued on a very positive note and ended on a high one.

So from every successful event there needs to be a takeaway that changes our point of view. For me it was captured when a couple of hayseeds took the stage and planted the seeds of change and showed us all how a connected Haystack using data modeling could alter industry dynamics as we know them. A great production with an amazing cast in this historical and mildly hysterical production.....smile.

Just a couple of Hayseeds planting ideas about using data modeling and an open source initiative to change how we view everything. This is what our world needs to understand about the true value of the power of data modeling.

Eric and Ken of ControlTrends.org, a fellow media sponsor of the event captured this all on video so I can show you what I am talking about.

They wrote; One of the highlights of Haystack Connect was the presentation by Jason Briggs and Scott Muench of J2 Innovations. Their demonstration shows the ease, power and flexibility of Haystack. Great job guys!

http://controltrends.org/2013/05/see-the-power-of-haystack-connect/

Take the time to view it all, as it starts with building on Niagara Framework but that is not where it ends. You will be suprised to see how the haystack concept can be used almost everywhere and at several levels.

In addition several others from the event have provided us interviews to let you catch up on what you missed from their point of view.

Haystack Connect Event Review The ability to automatically interpret the meaning of data enables best of breed applications to be implemented, quickly and at lower cost. This gives owners more choice and a faster path to financial return. - Marc Petock, Vice President Marketing, Lynxspring

Back from Haystack Connect What was great about Haystack Connect was that it was organized by small, mid-sized and emerging companies.- George Thomas, President, Contemporary Controls

Simple and Open Wireless Control for Smart Buildings There is no doubt building operators are looking for a more intelligent and elegant building control solution to make their lives easier. - Danny Yu, CEO, Daintree Networks

Lighting Controls, Made Simple By putting a microprocessor in the ballast of the fixture you will be able to turn on one lamp at a time and provide multilevel illumination without using a dimmer. - James F. Loughrey, Inventor, Entrepreneur, Logica Lighting Controls, LLC

Most of this paperless event was captured on the agenda/presentation/event software called Pointview.

A completely paperless event should be applauded but now is a reality for all events moving forward. All that is required is a cloud service like Pointview plus a excellent wifi with seamless connection from your room to event provided by the venue. In addition to a no hand out program, "it is on your phone" the event moved towards the end of paper cards exchange allowing us to bump each others phone and exchange contact info while running the app. I was impressed

when this worked between an iphone and my android phone. Great stuff. We all need to walk the walk and talk the talk in everything we do.

Go to the **http://2013.haystackconnect.org/** web site and click on the Agenda and then Agenda Summary at a Glance. Go to presentations during the core days of Tuesday and Wednesday. All of the presentations using Power Point have been converted to pdf and are linked to the web site. Example; Data Modeling.

Presentation Files in pdf format are listed on the bottom right hand side of each breakout session. Just click on them to download and see what you missed.

A great event and as George Thomas said in his interview; "You need to give credit to John Petze, Scott Muench and Marc Petock for their efforts in organizing the event."

Good, no, Great job guys! I am already looking forward to my next meeting with you in New York. We are now organizing the meeting and building the agenda online with Linkedin. Meeting is open to all. A great place to see and hear what others are doing. Please join us.

Connection Communities Meeting New York - AHRExpo Please come to a meeting to help us all make stone soup in New York. I have just confirmed a room for our meeting with AHRExpo - The Connected Communities Collaboration meeting with various industry speakers will be Wednesday, Jan. 22, 2014 1:30 at **http://www.ahrexpo.com/** Room Location TBA

Help guide us to find the best recipe for stone soup for the industry, see John and Marc's comments below.

http://www.automatedbuildings.com/news/may13/reviews/130425025901cccny.html

Ken Sinclair, Publisher, automatedbuildings.com

Ken Sinclair is the publisher of the online magazine and web resource AutomatedBuildings.com. Since 1999 he has been providing the news and connection to the community of change agents that are creating our present definition of smart, intelligent, integrated, connected, green, and converged large buildings. The virtual magazine and web resource provides a searchable platform for discussion and exchange while creating opportunities for B2B for all new and existing stakeholders.

Haystack Connect

- HOME
- REGISTRATION
- VISION
- VENUE & EVENT
- SPONSORSHIP PACKAGES
- SPONSORS & EXHIBITORS

Smart Devices. Smart Buildings.

Smart Business.

Connecting Companies, Communities and People who are redefining Smart and Connected Systems.

Learn More

2013 Haystack Connect

APR29 – MAY2

THE CHATTANOOGAN
CHATTANOOGA, TN

MON April 29th - THU May 2nd, 2013
Registration only $595 (including meals for 4 Days)

REGISTER | AGENDA | ANDROID APP ON Google play | Download on the App Store

VIEW MORE VIDEOS

APRIL	APRIL / MAY	MAY
29th	30th / 1st	2nd
SPONSOR MEETINGS	MAIN MEETING SESSIONS	SPONSOR MEETINGS

VISION

THE HAYSTACK VISION

Technology continues to drive rapid change in smart building technologies, energy management and operational efficiency solutions...

Read More

AN EVENT

BY THE COMMUNITY, FOR THE COMMUNITY

The building automation, energy, and M2M industries have a variety of events. Some are under the control of a single manufacturer...

Read More

SPONSORSHIP OPPORTUNITIES

Haystack Connect

Haystack Connect offers a range of sponsorships and exhibitor packages....

Read More

DaintreeNetworks | EasyIO | J2INNOVATIONS | lynx | SkyFoundry | Intelegence

Platinum Sponsors

VIEW ALL

info@haystackconnect.org

Haystack Connect 2015

May 18 - May 20

Smart Data. Smart Devices. Smart Buildings.

Smart Business.

Join us **May 18th - May 20th** for the Haystack Connect conference: The place where the community of automation and IoT professionals come together to learn and share the latest techniques for connecting systems and using data to advance the efficiency of buildings, equipment systems and processes.

Venue

Haystack Connect 2015 will be held at the Cheyenne Mountain Resort in Colorado Springs Colorado. This unique southern Colorado resort destination is equal parts world-class meeting facility and four-diamond resort. The multifaceted Cheyenne MountainResort is a place where the possibilities are as endless as the views. For meetings, spectacular indoor and outdoor venues create the backdrop for imaginative events, allowing attendees to enjoy the beautiful setting while remaining productive and inspired. Haystack Connect will take the entire facility providing attendees with a unique, intimate and productive event.

Haystack Connect Review

A summary of my observations and takeaways "My Back from the Stack".

Great job by the Haystack community in organizing a feel good event. The event had more than 50% of attendees under 40, that was amazing. This will help the Haystack movement grow faster than anything else we can do. They also created a new word "Haystackable". Good fun was had by all and you could feel the revolution in the air.

The complete conference is now online with linkage to pdf of most presentations on the right hand bottom of each interactive session page served off the PointView menu. Conference Agenda and Session Details below formated by PointView.

http://2015.haystackconnect.org

For a great video and commentary of event go to **http://controltrends.org**. Eric and Ken did an amazing job of exposing with video interviews the people who made this event work.

ControlTalk Now Live at Haystack Connect Haystack Connect as it resumed its mission to adroitly usher in the world of self-defining data — in a collaborative way that no single-minded entity could possibly achieve. Marc Petock and three special guests, Brian Frank, Anno Scholten, and Richard McElhinney join ControlTrends for a pre-reception discussion of why some 240 of the top global automation and IoT professionals are attending the second Haystack Connect meeting, and what the continued growth and success of Project Haystack represents for the future of building data.

Christopher Naismith of SES Consulting and I after our session ATTRACTING SELF-LEARNING ASSETS met with several folks to discuss self learning companies and we learned lots. An eye opener for me and the other over 40 folks was:

Millennium folks by definition are self learners. We simply need to define what they need to learn. We also need to explore how to best digitalize the thoughts of the boomer's and connect to today's flavor of the day social media. Our discussions found a few folks using "slack" as an internal communication tool. Christopher has started a group on slack... call Hayslack.....too cute right? Let Chris know if you want in and we will get you an invite. christopher@sesconsulting.com Most folks in the group to date, are consultants trying to make sense of all this stuff.

I just posted linkage to this Haystack Guide specification on the Haystack group: **http://project-haystack.org/download/file/Guide-Specification.docx**

Be sure to join the **haystackconnect LinkedIn Group** to be part of Connecting Companies, Communities and People who are redefining Smart and Connected Systems. Smart Devices. Smart Buildings. Smart Business.

The "betweeners", those folks between millenials and the baby boomer's are the connection piece I have been looking for to provide the connection of the millenial folks to the baby boomer's now leaving the industry. They have an amazing talent to be able to talk with trust to all generations.

Track 2-Session 1: Next Gen Hardware I liked this thought from session's discussions, Cellphone industry has sold every one in the world at least one cell phone now they want to do it again for each M2M & IoT device. Imagine the numbers. Intel jumped in to discussion to advise that they are as much as possible pushing functionality into a chip level. It was a mind expanding discussion to grasp texting and having social media discussions with our field sensing and metering devices and again proved that Haystack naming is imperative.

New generation hardware will move data to the field from the cloud. This is a new concept using cheap memory like

those used for cameras will allow field devices to store more information than ever possible before. Field hardware will have faster communication speed 5 Gigabyte plus, and will be designed as big data at the edge devices. This is a big heads up and head shake for big data in the future as data will become an integral part of sensing and control.

Three lessons learned:

1. Next generation devices will have more memory, lots more memory, moving big data to the edge and storing it there in raw format for use as required. Devices will store their own data for deeper understand by all.
2. These next generation devices will have greater network speeds to allow big data to be polled from the edge. This seems to be the best devices to bring forward legacy devices. Example Contemporary Controls' new product, review power point to understand this evolution. Original control systems were not data gathering systems memory and communication speeds were not an issue. Cost of memory and network speed has dropped drastically.
3. Next Next gen - 2 to 3 years will see cell phone like devices with gps, camera, at the edge of big data using smart phone evolving power and wireless networks such as NFC, blue tooth to wireless gather data this will allow the integration of people information and connection of devices to social media. Machine to machine will talk and text and advise us of how they are doing or if the need help via social media smile it will be a new world 4 sure. This is where thinking needs to be for systems moving forward and future proof

How will traditional control companies deal with this? They need to grow quickly embracing these concepts to be leaders not followers.

Haystack code is easy to understand and can be read by machines. Take the time to watch Brian Franks presentation.

This is what haystack code looks like;

id: @cheyenne
dis: "Cheyenne Resort"
site
geoStreet: "3225 Broadmoor Valley Rd"
geoCity: "Colorado Springs"
geoState: "CO"
geoPostalCode: "80906"
geoCountry: "US"
geoCoord: C(38.78,-104.83)
area: 106500ft²
yearBuilt: 2003
primaryFunction: "Hotel"
tz: "Denver

Rob Johnson's presentation summarizes the state of the BACnet Extended Data Model (XD) concept. BACnet XD is a significant step forward for BACnet, incorporating tagging and referencing concepts, and allows incorporation of multiple tagging dictionaries.

If you were there you know you need to be back in two years at the next Haystack Connect Event. If you were not you need to learn more about his exciting movement. The need is clear and we all need to ask those non believers; "How do you propose to manage big data tagging?" With the rapidly evolving trend of the movement of big data to the edge devices, every product now needs to address how it will interact with a machine readable agreed on tagging scheme.

We were very pleased to have our first ever contributing editor of AutomatedBuildings.com at the event. Tom Hartman was on a self learning journey to observe our changing industry.

The one piece of paper, paperless conference started with this story of stone soup so I will end my review with it. It is a good summary as well as opening. Thanks John.

The Fable of the Stone Soup – Updated

A system integrator weary from hours of manual effort to map data from different systems and devices sat down to assess his fate. There could be no value delivered to his customer, or money to be made for himself, if every project required this painful manual process. If only, he thought, everyone would define their data in such as way that it could be easily interpreted – automatically even – by his trusty laptop – oh but that's just a dream he thought.

Sometime later he happened upon an energy engineer struggling to do analysis of meter data to identify demand peaks across a portfolio of buildings who said – if only the meters had been named in such a way that I could tell which was which and what was what. It would save me so much time and save my customer so much money – oh but that is just a dream he thought.

As the system integrator traveled the country going from project to project the same problem faced him at every turn. There has to be a solution he thought. But what can I do? I am just a lone engineer... with deep experience in variable volume air distribution systems. If only there were others like me – maybe would could work together on a solution... Oh but that's only a dream. No one cooperates on standards in our industry.

Unlike the software industry where open source initiatives to solve industry-wide needs are the norm not the exception. (with sarcasm)

Our system integrator traveled the world and shared his story with everyone he met.

The data center expert who said – I don't know much about VAV – but I sure know PUE.

The Chiller plant genius who said I don't know much about energy rates, but I can predict a chiller's fate.

And the energy analyst who said I've never touched an actuator but I know my energy factors.

And together they came, each with knowledge of piece of equipment, device or system. They pooled their knowledge and realized that all of this seemingly unorganized equipment could be modeled in a uniform way. And with the help of some wizards of the bytes and bits they created a solution they knew would be a hit.

Easy to use, applicable to any type of system, device or data, open source, free to use, extensible, human and machine readable, and most important – driven by the community to meet their real world needs.

And that my friends is the story of Project-Haystack – the solution for making device data self-describing. They key to unlocking the value in our operational data.

And one helluva a good reason to have a party.

Rita H. Wouhaybi

Systems Architect, Intel

Rita is a Systems Architect with the Industrial & Energy Solutions Division in the Internet of Things Group at Intel Corporation.

Milan Milenkovic

Principal, IoTsense

As a principal engineer at Intel prior to founding IoTsense, Mr. Milenkovic led a number of projects involving IoT, distributed systems, smart buildings and energy management.

Rawlson O'Neil King

Communications Director, CABA

Rawlson O'Neil King is an expert on green buildings, smart homes and smart communities.

Data Freedom for Smart Buildings and Beyond

Takeaways from Haystack Connect 2017.

Hundreds of building equipment, controls and software vendors descended on a sleepy resort in Tampa last week with a simple goal: solving the interoperability challenge facing the Internet of Things. To be fair, Haystack Connect's stated purpose is enhancing data fluidity and freedom with an open source metadata framework that allows building stakeholders to unlock new value from networked devices and equipment. Don't be fooled: Project Haystack will accomplish far more than this, and the companies that came together in Tampa know it.

Excitement was in the air at the bi-annual gathering of Project Haystack supporters. If the mood of 2013's Connect was one of "I'm curious" and 2015's was "I'm interested," the vibe at 2017 was unmistakably "I'm ready." The organizers and vendors I spoke with made clear that they have moved beyond simple excitement over the vision that Project Haystack paints. Equipment and controls manufacturers, software vendors and systems integrators are today creating value for themselves and clients with the framework. Intel's recent announcement that they were joining the organization as a Founding Member further validates the feeling that this framework is ready to extend the open-source standard throughout the buildings market and beyond.

What exactly is all the fuss about?.

Data on Data

Those of us who don't work in the buildings industry might tend to overlook that the built environments in which we spend much of our lives are in fact highly engineered systems that are designed, developed, operated and maintained by a variety of players. As digitization has pervaded every other aspect of our lives, so too has it commandeered buildings. Networked sensors, devices and systems throughout buildings now feed data to analytics tools designed to enhance the efficiency of operations and the services offered to tenants.

However, the variety and fragmented nature of the buildings suppliers landscape challenges the ability of any supplier to deliver a valuable, seamless experience to end users. Central to this challenge is that systems from different vendors describe data in different ways. Semantic data models, as they are known in the digital world, provide standardized ways of describing where data was generated and what information it contains. This data about the data, or "metadata," make it far easier for various users to create value from data through visualization and analytics tools.I just posted linkage to this Haystack Guide specification on the Haystack group: http://project-haystack.org/download/file/Guide-Specification.docx

Read the Full Article here: **https://www.linkedin.com/pulse/data-freedom-smart-buildings-beyond-takeaways-from-haystack-adam-hise**

Adam Hise, Managing Director of Storage Risk Solutions at Ascend Analytics

Previously, Adam was a Senior Associate at Harbor Research at the time of this writing. Now, he is focused on advancing energy storage and renewable energy revenue risk solutions with Ascend's suite of industry leading analytic capabilities, following the integration of EnBalance Storage in January of 2021.

2019 Haystack Connect

Smart Data. Smart Devices. Smart Buildings. Smart Business.

May 13-15, 2019 / San Diego, California

The place for the Project Haystack community to network, share, create synergy, and generate business opportunities.

Haystack4

Haystack 4 Is Coming – What it Is and Why it Matters

Tuesday May 14, 2019

7:00am-8:30am	Breakfast & Exhibits
7:30am-4:00pm	Registration
8:30am-10:15am	**GENERAL SESSION** • **Welcome & Opening Remarks**: John Petze, Executive Director, Project Haystack, and Marc Petock, Executive Secretary, Project Haystack • **Well We Don't Know What We Don't Know – Bringing Organization to the Data Chaos to Derive Insight and Generate Actions. Keynote Presentation**: Jim Fletcher, Strategy Partner, Momenta Partners • **Benefits of Haystack Tagging for OEMs**: Ryan Schlotfeldt, Vice President of Sales, J2 Innovations; Marc Petock, Chief Marketing & Communications Officer, Lynxspring; James Johnson, Sales Engineer, Tridium; Todd Lash, Director and Segment Head – Systems and Room Automation, Siemens Industry; Jason Houck, CIO, IoT Warez • **Haystack 4.0 – What It Is and Why It's Important**: Brian Frank, Founder & CEO, SkyFoundry

Jim Fletcher

Strategy Partner, Momenta Partners

Jim Fletcher is an independent advisor, working with Momenta Partners as a Strategy Partner, focused on working with clients worldwide to drive Digital Transformation across a range of industries. Jim is a former IBM Distinguished Engineer, CTO for the Watson IOT Platform, IBM Master Inventor, and member of the IBM Academy, with extensive experience especially around Connected Industry (Cloud and IOT) with a strong focus on Edge Computing, and the role of analytics in driving business transformation. Jim holds over 60 patents, and is a regular speaker at conferences worldwide on a range of emerging topics. Jim is a member of the NC State University Executive in Residence program, and was recently inducted into the NC State University Computer Science Hall of Fame.

Interoperable Blues Band (IBB)

One of the most fun events at Haystack Connect!

Lenny "Fuzzy" Rankins

Megan Wolf

The Newest Associate Member Companies Joining the Mission

ALLANDER ANALYTICS

Based in the United Kingdom, Allander Analytics is a software development company founded in 2017 in collaboration with experienced energy managers, building control specialists, and software engineers.

Our goal is to modernize the outdated systems in place across the industry by developing and utilizing cutting-edge software that can fully reverse-engineer buildings based on energy performance.

We aim to provide a full understanding of how buildings operate and effectively demonstrate to our users where savings can be made.

Building Book: Data Analysis and Analytics

Our mission is simple; Help our customers reduce their energy consumption, reduce their carbon footprint, and save money. With these goals in mind, we created Building Book; A Cloud-based Software as a Service (SaaS) application where our users can seamlessly connect with their buildings and analyze performance in real-time.

Building Book users can also easily upload historical energy consumption data to analyze past performance. Using AI and machine learning, Building Book can create energy models for buildings that can be used to identify problems and alert users to issues that could otherwise go unnoticed.

Project Haystack Integration

Project Haystack enables us to take advantage of the standardized conventions of haystack tags when interfacing with a building's IoT systems using these protocols. This allows us to quickly map buildings of any size and type within our software and begin recording real-time data for analysis.

Project Haystack also gives us a common platform for users to take advantage of, when filtering through vast amounts of data within our system.

We are delighted to be part of the Project Haystack community.

You can learn more about us at allanderanalytics.co.uk.

Automated Logic

Automated Logic was founded in 1977 and manufactures high performance building management solutions that make buildings smarter, more energy efficient, and more comfortable. Automated Logic is a part of Carrier Global Corporation (NYSE: CARR), a leading global provider of healthy, safe and sustainable building and cold chain solutions.

We Make Buildings Better.

Our WebCTRL® building automation system (BAS) is known for its superior graphics and intuitive user interface. It allows building operators to manage all of their building systems - including heating, air conditioning, lighting, fire, and security – anytime, anywhere through a variety of internet devices.

Our EnergyReports™ package supplements the WebCTRL BAS by pulling building energy data into the cloud for in-depth analysis. EnergyReports provides customizable dashboards, energy cost analysis, carbon dioxide emission calculations and benchmark comparisons of individual buildings as well as portfolios of buildings. The EnergyReports platform helps connect building managers and occupants to building energy consumption.

Our IntelliSuite™ Analytics Solution adds a comprehensive set of analytics capabilities. Real-time dashboards, fault detection and diagnostics (FDD), predictive analytics, and advisory/optimization services help operators optimize building and equipment performance.

Automated Logic's worldwide network of authorized partners includes over 230 field offices, with proven experience in building automation, energy management, and controls.

Supporting open standards is a key element of Automated Logic's strategy to lead the future of building automation, as we continue to advance our open architecture so that we are positioned to offer the most scalable, interoperable, and secure building automation solutions in the industry. Our participation in Project Haystack reinforces this commitment as it provides standard naming conventions for building and equipment data, helping to provide cost-effective analysis and management of that data now and in the future.

Learn more at **automatedlogic.com**.

buildingsIOT

Buildings IOT is changing the way the built environment understands, reacts and adapts through technology. Our software increases the longevity of building assets, improves comfort for building occupants and helps building owners achieve greater efficiency. We work with contractors, integrators, engineers, owners, operators and all real estate stakeholders to provide new solutions for buildings that include offices, hospitals, data centers, shopping malls, universities and government institutions. Our customers include some of the most prestigious and well-known organizations in the world.

We believe smart buildings don't just come from one decision made at the beginning of a project –they're made continuously over time by teams from construction to operations. Smart buildings come alive through thoughtful designs and meaningful integrations and they're backed by data that is reliable, informative and useful.

With dedicated teams of building system experts, master systems integrators, user experience designers, software developers and project execution wizards, Buildings IOT delivers smart buildings that consistently live up to their promises.

Learn more at **buildingsiot.com**.

CLOCKWORKS ANALYTICS

Clockworks Analytics was founded in 2010 within MIT's Building Science Department with a mission to make a massive positive impact in the built environment.

Automated Analytics. Smarter Facilities.

The Clockworks Analytics HVAC Fault Detection and Diagnostics (FDD) platform plugs into existing BMS and metering systems and analyzes thousands of data points to prioritize the highest impact building issues related to energy performance, indoor air quality and equipment operation. Our unique data information model goes beyond simple fault detection by identifying the relationships between issues, diagnosing the root cause, and providing clear, recommended actions.

Our enterprise customers and our controls and mechanical service partners across the world have saved over $20M and completed 20k tasks using Clockworks' analytics-based monitoring to proactively address building health issues, identify energy savings, and avoid equipment failures.

Root Cause Diagnostics at Scale

The Clockworks' expert system is based on an information model that includes point tagging, along with standardized metadata for sequences of operation and mechanical configuration. This allows the platform to scale across building, system and equipment types without any custom rule-writing to account for the thousands of nuanced scenarios that appear in all building portfolios.

We are excited to bring our expertise to Project Haystack and drive greater interoperability between vendors like us and the Haystack community, as well as continue our mission to help the industry-at-large turn operational data into real-time insight for healthier buildings in a scalable way.

Learn more at clockworksanalytics.com.

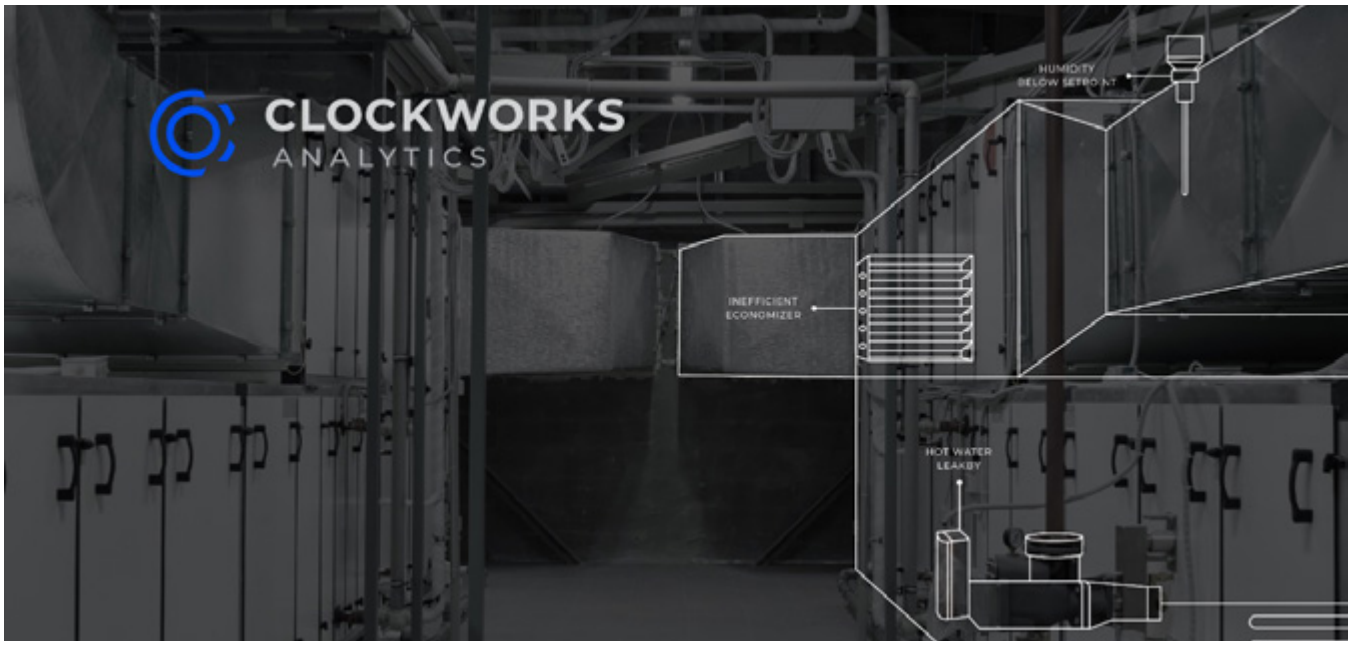

REGISTER TODAY! Haystack Connect 2021

Haystack Connect will be a

Virtual Conference May 4 - 6, 2021

Haystack Connect 2021 is organized and produced by the Project Haystack Organization—an open source community of people and companies who share the vision that a connected, collaborative community can move the industry forward in ways that no single supplier can! The event builds on the inspiration and mission of the community to address the challenges of making smart device data work seamlessly across applications of all types through the adoption of a standard approach to semantic modeling of equipment systems and their data.

This year marks the 10th Anniversary of the Project Haystack Organization's formation. To commemorate the occasion and allow the greatest number of people to attend, Project Haystack is making registration for Haystack Connect 2021 FREE for all attendees. Guarantee your spot by registering now, and don't forget to check out sponsorship and speaking opportunities.

To Register, visit: www.haystackconnect.org.

Haystack in Practice – Haystack Tagging in Real World Applications

- Haystack's impact in the engineering community and design phase—Haystack in project specifications
- Evolution of new data-oriented business offerings by the engineering community
- Haystack beyond building environmental systems—Haystack as applied to industrial, data centers, agriculture, and others.
- How Haystack is transforming software and hardware products
- The Haystack/Analytics Connection: Haystack's role in streamlining analytics of IoT data
- End users' benefits and experiences implementing data management strategies and Haystack's role
- How Haystack drives business value for systems integrators, OEMs
- In the Trenches—experience applying Haystack tagging in projects small and large, and across vertical applications
- Haystack in digital twins

Technical Track

- Conclusion of the Haystack 4 public review and final release
- Important new open-source announcements for Haystack 4
- Upcoming enhancements for validation and data shaping
- Working group sessions
- Crash course – get up to speed on Haystack 4

New Products and Services Using Haystack – "Pitchfest"

- Up to twelve 10-minute sessions are being offered on a first come first serve basis

HAYSTACK CONNECT 2021 AGENDA

Visit www.haystackconnect.org for exact session times.

TUESDAY, MAY 4, 2021

GENERAL SESSION

- **Welcome & Opening Remarks:** John Petze, Executive Director, Project Haystack, and Marc Petock, Executive Secretary, Project Haystack
- **Keynote Address:** Steve Holzer, Principal, HolzerTime
- **State of the Union for Haystack 4:** Brian Frank, Technical Lead, Project Haystack

TRACK 1 – HAYSTACK IN PRACTICE

- **Recognizing Haystack Permanence – What's the Next Step?:** Matt Schwartz, Associate Principal, Altura Associates
- **Making Data Work for You - How Does Haystack Help to Ease Concerns?:** Jamie Lee, Product Manager, Siemens Industry, Inc.
- **Increasing Interoperability with Clockworks Analytics Information Model:** Nick Gayeski, Chief Executive Officer, Clockworks Analytics

TRACK 2 – TECHNICAL TRACK

- **Developer's View of Haystack 4:** Brian Frank, Technical Lead, Project Haystack
- **Haystack Core TypeScript Library:** Jason Briggs, Chief Technology Officer, J2 Innovations, and Gareth Johnson, Senior Cloud Architect, J2 innovations
- **Haxall – A New Haystack Open Source Framework**: Brian Frank, CEO and Founder, SkyFoundry

WEDNESDAY, MAY 5, 2021

TRACK 1 – HAYSTACK IN PRACTICE

- **Haystack in Practice - Success Story for Coster Group:** Alex Rohweder, Chief Executive Officer, J2 Innovations; Matteo Pierone, Global Strategy Head and EMEA Professional Services Director, J2 Innovations; Davide Manca, R&D Manager, Coster Group; and Fabio Antoci, R&D Engineer, Coster Group
- **Haystack as Part of the Solution to Climate Change:** Omar Tabba, Vice President, Products & Solutions, BrainBox AI
- **Ontology Alignment Project:** Jon Schoenfeld, Director of Energy and Analytics, Buildings IOT

TRACK 2 – TECHNICAL TRACK

- **Dynamic Integration in Digital Twins:** Anto Budiardjo, CEO, Padi, Inc.
- **Haystack in TwinWorX Digital Twins Solution:** Eugene Woo, Chief Technology Officer, e-Magic Inc.

- **Applying Haystack Tagging to Variable Refrigeration Systems:** Yuya Saito, Associate, Innovation Program, Daikin Open Innovation Lab Silicon Valley (DSV)
- **WG#798 Haystack JSON Encoding:** Gareth Johnson, Senior Cloud Architect, J2 innovations, and Radu Racariu, Senior Programmer, J2 Innovations
- **Haystack Labs: Prototyping Data Validation with Functional Point Groups and SHACL:** Cory Mosiman, Software Engineer, PassiveLogic

THURSDAY, MAY 6, 2021

TRACK 3 – PITCHFEST

- **Working Integration of a Connection Profile Broker Using Padi:** Anto Budiardjo, CEO, Padi, Inc.
- **How Can I Learn How to Use Project Haystack Tags?:** Emma Eynon, Director, Fantom Factory
- **Simplifying Implementation of Haystack:** Philippe Prados, Data Scientist / Architect, Engie Digital & OCTO Technology
- **MicroBMS Demo Including Edge2Cloud Connectivity:** Scott Muench, Vice President of Customer Experience, J2 Innovations, and Jeremy Wolfe, Vice President Sales Americas, J2 Innovations
- **Automated Tagging:** Jean-Simon Venne, Co-Founder and CTO, BrainBox AI
- **Edgy Haystack:** Alper Üzmezler, Managing Partner, BAS Services & Graphics, LLC

TRACK 1 – HAYSTACK IN PRACTICE

- **Applying Haystack 4 in a Residential Analytics and Control Application:** Adam Wallen, Training Lead, SkyFoundry
- **Haystack Benefits in the K-12 Education Market**: Jamie Lee, Product Manager, Siemens Industry, Inc., and Steve Crowe, Integration Team Project Manager, Resolute Building Intelligence
- **Analysis of Tagged Energy Data Via Machine Learning**: Jan Široký, Head of Research Department, Energocentrum PLUS

TRACK 2 – TECHNICAL TRACK

- **Tridium's Haystack Tag Dictionary:** Eric Anderson, Niagara Software Engineer, Tridium
- **Haystack Tagging for Aquifer Thermal Energy Storage Systems:** Jaap Balvers, Team Lead Building Analytics, BAM Energy Systems
- **Haystack+MQTT+Sparkplug:** Richard McElhinney, Chief Software Architect, Conserve It

Celebrating 10 Years
2021 Haystack Connect
Smart Data. Smart Devices. Smart Buildings. Smart Business.

Thank You Sponsors

COMMUNITY CHAMPION

J2 INNOVATIONS
A Siemens Company

PLATINUM

ALTURA
e-Magic
SkyFoundry
SIEMENS
TRIDIUM

GOLD

buildingfit.
buildingsIOT
BRAINBOX AI
CLOCKWORKS ANALYTICS
KMC CONTROLS
LYNXSPRING
TOSIBOX

SILVER

Intellastar
Conservefit
Fantom Factory
novant
STACKHUB
Yorkland Controls ENVIRONMENTAL SOLUTIONS

MEDIA SPONSORS

Automated Buildings
CABA
CONTROL TRENDS
EMA ENERGY MANAGEMENT ASSOCIATION
Monday Live!
NEXUS
SBA

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 Type System WG

Champion:
Brian Frank,
SkyFoundry

Haystack WG 551 was the vehicle used to design Haystack 4. In March of 2019 we transitioned this working group to the review process stage. Due to the scope of Haystack 4, the review process was run using an independent website on https://project-haystack.dev/ which allowed us to maintain both the Haystack 3 and Haystack 4 online simultaneously.

It has now been slightly over two years during the review process, and we are drawing the review to a close. Over the coming months, we will transition project-haystack.org to the new Haystack 4 documentations and it will become the official primary specification for Project Haystack.

If you are interested in more details about the Haystack 4 rollout plan or wish to contribute feedback, visit forum post **https://project-haystack.org/forum/topic/903**.

Also don't miss the session on Haystack 4 during Haystack Connect in May!

#609 AHU Standing WG

Champion:
Jay Herron,
BuildingFit

The AHU Working Group is meeting the week of April 26, 2021 to discuss the following items. Evaluating and proposing solutions for these items is our main priority:

- Filter equipment and points
- Min/max points
- Air pressure specialization
- Two-pipe heating/cooling system notation
- AHU tagging updates for denoting fan placement
- Reheat, non-powered VAV equipment tagging.

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

#837 Haystack Labs Standing WG

Champion:
Cory Mosiman,
PassiveLogic,
(previously at
NREL)

The original intention of the Haystack Labs working group was to explore new directions for the Project Haystack community, identify the needs of existing Haystack users, and try to understand the major differences between Haystack and some of the other Semantic Interoperability efforts.

One of the consistent themes that has been popping up in our discussions is the need for some sort of validation mechanism. The question of 'Does my Haystack match your Haystack' has been a pretty consistent theme and is often a source of contention. Whether Haystack is being exchanged between analytics applications, a team is looking to perform QAQC of their own internal Haystack implementation, or a team is looking to validate that their Haystack implementation meets a client's specification, data validation appears to be a key requirement. There are a few different aspects to consider in the context of data validation with Haystack.

In an effort to keep the work practical and focused, we have been using ASHRAE G36 equipment types to guide our implementation. We break down the question of 'Does your Haystack match my Haystack?' into the following aspects of validation:

1. Does the implementation utilize tags that are part of the Haystack library or are there custom tags?
2. Does an entity implement a very specific set of tags? Think `{discharge air temp sensor point}`.
3. Does a specific entity have a set of entities related to it? Think of a VAV box that we want to ensure has at a minimum a `{damper cmd point}` and a `{discharge air flow sensor point}` connected to it via an 'equipRef'.

The first validation aspect is identifying whether the tags implemented on an entity are part of the Haystack standard library of tags or whether they are custom tags. This is relatively easily solved with Haystack 4, since the standard library of tags is now well serialized and versioned (thanks Brian!). This aspect of validation is typically most important when exchanging Haystack data between different applications.

The second validation aspect to consider is identifying that a specific set of tags has been implemented on an entity. Some refer to these as tagsets, others as concepts, others as classes. Oftentimes we think of these as very specific 'types' of things, maybe as a type of equipment or type of point. Haystack 4 defines a bunch of these and refers to them as protos, which are basically predefined tagset templates. These should really help users across the Haystack community implement tagsets in a more standardized manner. We definitely want feedback from the community to identify whether protos are missing or where it is unclear what a proto means / how it should be used. Many of the protos are analogous to canonical classes defined in the Brick schema or Google Digital Buildings project.

The third validation aspect to consider is identifying that an entity has a specific entity type / a specific set of entity types related to it in some specific way. The easiest example to consider is that a specific set of point types equipRef to a specific equipment type. Haystack 4 defines the children mechanism, which helps answer the question of "what points might be contained by a vav box".

#837

Haystack Labs Standing WG (cont'd)

There is an important distinction that should be made here that is easy to miss. Protos and children both have a similar goal of defining typical things that 'might be'. These are useful for quickly building things out in a generic way, which is often what users / applications are doing when they implement Haystack. Validation is really the process of defining exactly what 'must be' and seeing whether or not an implementation of Haystack meets this definition. We see these as very complementary efforts but with different end goals.

We are currently referring to these definitions of what 'must be' as validation profiles or validation shapes. They are similar to the 'Equip Functions' used by the Ontology Alignment Project or 'Abstract Types' defined in the Google Digital Buildings project. Our goal is that these validation profiles will be part of the Haystack standard, that we can define specific community agreed upon validation profiles, that validation profiles will be highly composable, and that we will define a standard process for validating conformance to a validation profile. If you are familiar with the SHACL specification and shapes, we have been using this as a first pass for investigating what this might look like.

All in all, it has been a fun year in the Haystack Labs WG and I'm really thankful to the folks who are participating, contributing ideas, and asking questions. The Haystack surveys that have recently gone out were also put together by the Haystack Labs WG and our plan is to share the results of the survey with the community down the road. If any of this work sounds interesting to you, we are always looking for others to join and help us flesh out these ideas. Please consider joining and contributing.

Check out the Working Group here:

https://project-haystack.org/forum/topic/837

#734 ATES Systems WG

Champion: Jaap Balvers, BAM Building Analytics – BAM Energy Systems

The ATES Systems WG (#734) has compiled a proposal for standardizing the tagging definitions for aquifer thermal energy storage systems. These energy-efficient heating and cooling systems use thermal energy from ground water to provide heating and cooling to a building. The proposal contains a number of new equip definitions (e.g. ates, well, flowInverter) and some new tags that can even be used in other situations (e.g. cool, warm).

A write-up and tagging definitions are available and recently we began adding the proposed changes to the Haystack repository. We expect implementation in the next version of Haystack. At Haystack Connect 2021 the WG will present an overview of the new Haystack definitions and some challenges encountered along the way.

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

Project Haystack Working Groups List

WG	Topic	Champion
#514	Dry Bulb Points and The 'air' Tag	Jay Herron
#551	Haystack Type System WG	Brian Frank
#626	RESET Standard and Air Quality Tags	Cory Mosiman
#496	Lab/Fume Hood Working Group	Gabe Sandoval
#501	Flow Modeling working group	Karine Lavigne
#503	Access Security Working Group	Justin Tashker
#505	Refrigeration System	Nathan Rona
#506	Unitary Equipment Working Group	Eric Loew
#553	Reference Model	Patrick Coffey
#492	New Data Center Tag Working Group	Ron Snyder
#530	BIM/Haystack Working Group	Chris Renter
#667	Cybersecurity Working Group	F Gordy
#701	Data Center Tags	Jason Ganiatsas
#709	Haystack RDF Export - Working Group	Matthew Giannini
#776	Working Group: Greenhouse Gas	Matthew Giannini
#734	Working Group: ATES Systems	Jaap Balvers
#497	Chiller Plant Enhancements Working Group	Sean Stackhouse
#595	Invitation to Project-Sandstar Working Group	Alper Üzmezler
#705	Lighting Systems WG	Jeremy Yon
#792	Haystack JSON Encoding WG	Gareth David Johnson
#609	AHU Standing WG	Jay Herron
#837	Haystack Labs Standing WG	Cory Mosiman

To learn more or to join a Working Group, visit https://project-haystack.org/forum/wg

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 White Paper 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 Specifications. Now available in English, French, Chinese, Japanese and Turkish.

Download CABA White Paper 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.

Social Media Curation

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

Projects

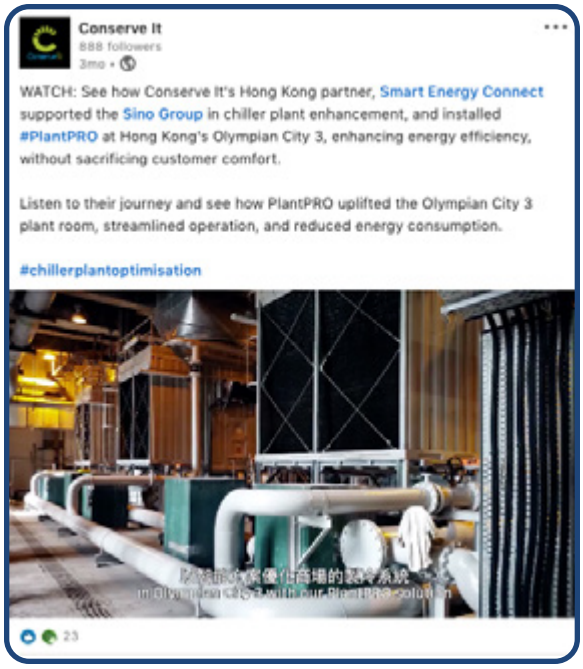

Installed PlantPRO at Hong Kong's Olympian City 3.

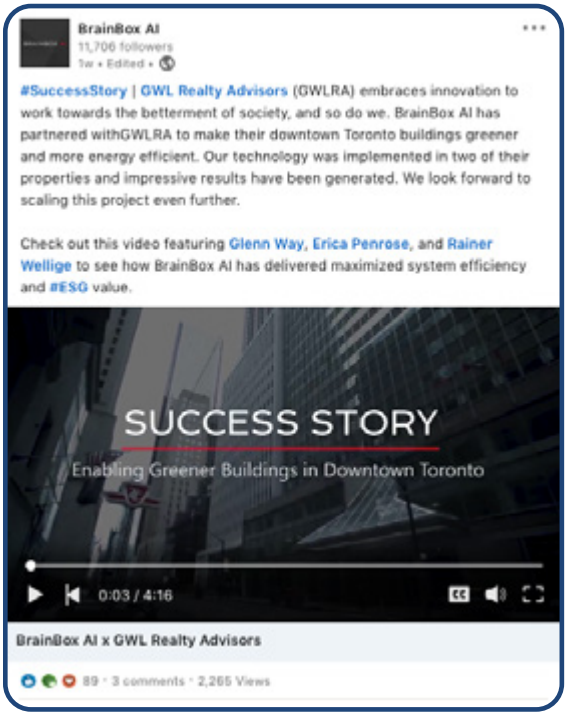

Making Downtown Toronto Buildings Greener and More Efficient.

A Journey to Net-Zero by 2050.

Committing to Zero Carbon Across the Entire Built Environment.

Benefits of Bringing Intelligent Automation into Government Buildings.

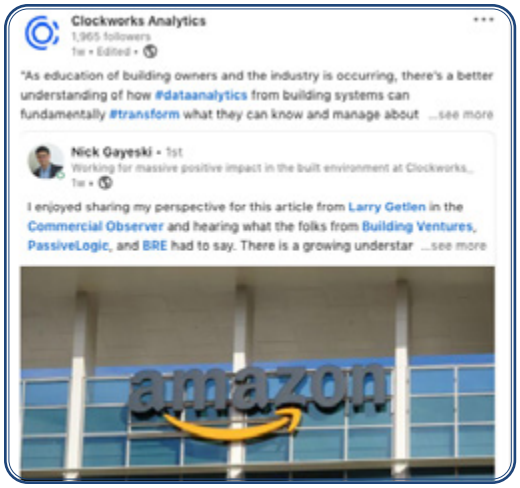

Better Education of Building Owners and the Industry.

Distributed Architectures? Edge-to-Cloud IoT Architectures?

Evolving an Archaic Preventative Operational Model and Making it Data-Driven.

Products

Zone Level Temperature, Air Quality and Energy Management Control.

One of the Key Tenants of the IoT is Scalability. But What Does That Mean?

What Are You Looking For in an Edge Controller?

Just Released Haystack Core - A TypeScript Library That Supports Haystack 4.

Actionable Insights to Improve Building Performance.

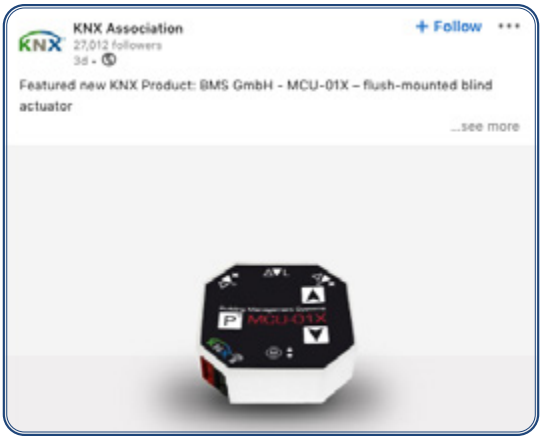

Featured New Product: A Flush-Mounted Blind Actuator.

A Wireless and Batteryless Switch.

Project Haystack

connections

Become an Advertiser

The Project 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.

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.

Allander Analytics designs and develops industry-leading energy management and data visualization software. Our Building Book platform enables users to model, analyze and report on the energy consumption of their buildings. Using the latest technologies and the power of the cloud, we provide real-time analytics alerting customers to anomalies and opportunities within their data.

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.

Automated Logic is a global provider of high-performance, integrated building management solutions that make buildings smarter, more energy efficient, and more comfortable. Automated Logic's worldwide network of authorized partners includes over 230 field offices, with proven experience in building automation, energy management, and controls. It is also part of Carrier Global Corporation (NYSE: CARR), a leading global provider of healthy, safe, and sustainable building and cold chain solutions.

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.

BrainBox AI is at the forefront of the green building revolution with its unique technology combining artificial intelligence and cloud computing to create a fully autonomous commercial HVAC solution. BrainBox AI overlays deep learning algorithms on existing HVAC functionality to automate the modulation of each component, reducing a building's total energy spend by up to 25% while improving occupant comfort by 60%. The solution leverages AI to predict building energy consumption at a very granular level and enables our autonomous HVAC system to operate the building pre-emptively.

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.

Associate Members

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.

At Buildings IOT, we're changing the way the built environment understands, reacts and adapts through technology. Our software and services increase the longevity of building assets, improve the comfort of building occupants and help building owners achieve greater efficiency. We develop and deploy cloud-based building analytics software, we implement complex Integrated Building Management Systems, we design and install controls systems, we maintain building assets and we provide IT managed services. We excel at all of our efforts because we know buildings.

The Clockworks Analytics HVAC Fault Detection and Diagnostics (FDD) platform plugs into existing BMS and metering systems and analyzes thousands of data points to prioritize the building issues related to energy performance, indoor air quality and equipment operation. Our unique information model goes beyond simple fault detection by identifying the relationships between issues, diagnosing the root cause, and providing clear recommended actions. Clockworks' analytics-based monitoring allows you to proactively address building health issues, save energy and avoid reactive failures tomorrow.

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.

e-Magic Inc. specializes in providing expertise and software for the design, development, and integration of large scale industrial IoT and Azure Digital Twins solutions globally Applications include Centralized Operations, Smart Buildings, Facilities and Cities, Smart Manufacturing, Industrial production and AI for prediction and optimization. Our solutions have been installed in a wide range of industrial sectors including: buildings, facilities, manufacturing, utilities, mining and metals, cement, oil and gas, food and beverage, chemical, petrochemical and pulp and paper.

EMA is a trade association dedicated to providing education, training, and certification in the field of building and facility energy efficiency. Its Energy Management Professional certification (EMP) has achieved accreditation by ANSI and is recognized by the Department of Energy's Better Buildings Workforce® program.

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.

Associate Members

Intelligent Buildings, 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.

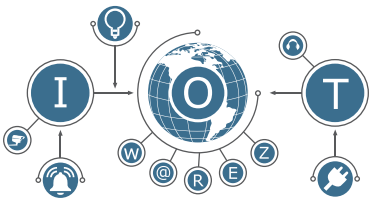

IoT Warez develops custom software that helps technologies communicate together. From state of the art data centers to environmentally conscious facilities, our software development team is capable of building solutions that connect anything and everything. IoT Warez offers a suite of hosted software 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.

KODE Labs has developed an enterprise level platform which streamlines the routine tasks of discovery, templating, tagging and data configuration and provides complete monitoring and control of building systems. The KODE Labs platform offers a data focused interface that surfaces the critical insights necessary to maximize operational efficiency across your portfolio.

Resolute provides a building-performance analytics and reporting solution that integrates with the Niagara Framework®, enabling quick and reliable use of real-time data, analytics-driven insight and on-demand reports to better manage buildings and achieve quantifiable performance gains. Leveraging the power of Project Haystack standardized data and tagging models and the Niagara open-source connectivity protocol, our solution allows direct connectivity to the Resolute Cloud™ from a Niagara instance - regardless of brand - without the need for additional 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.

Associate Members

Through the implementation of WideSky®, we aim to unlock the value of your energy, environmental and building data. Our scalable, intelligent solutions can improve profits and sustainability of your business. The qualified and experienced WideSky team has decades of operational and information technology experience. Coupled with our partner network, we can implement future-proofed, well-supported solutions tailored to your business on a global scale.

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.

For all the latest Project Haystack marketing activities visit marketing.project-haystack.org.

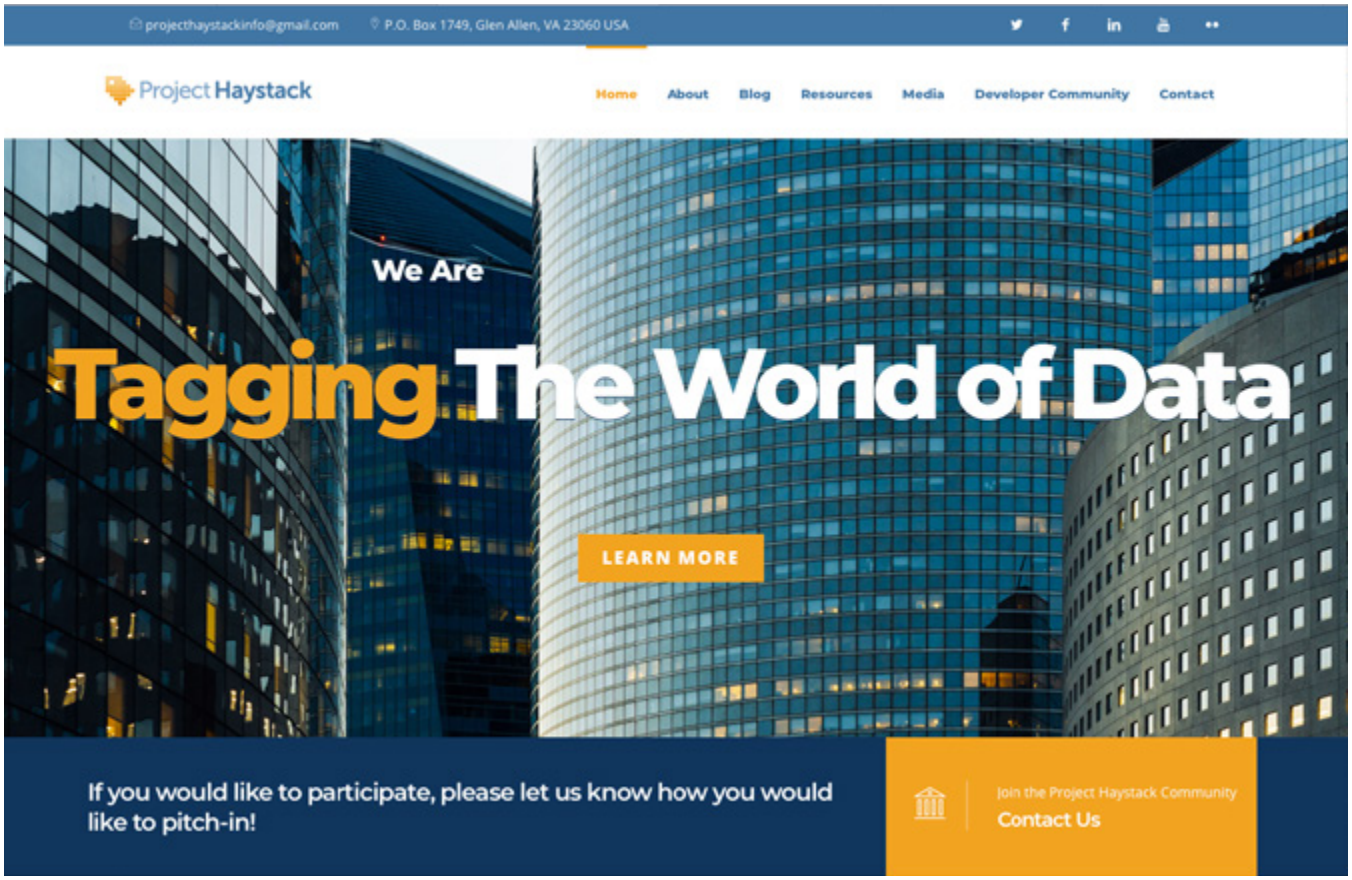

Project Haystack

www.project-haystack.org