How the Special Snowflake Problem Impacts Asset Management—and What You Can Do

Asset managers face ongoing pressure to maximize the reliability and efficiency of their installations. Monitoring a single site is challenging enough, but when management spans two or more sites, additional complexities are typically introduced.

The core problem? No two sites are the same. They differ not only in geographic location, but also usually in age, schematic configuration, maintenance history, and specific functional requirements.

Key Findings

- The "special snowflake" issue puts hard limits on your ability to monitor and optimize all systems at all sites at the same time.
- Effective point mapping for 24/7 monitoring solutions leads to better reporting and analytics.
- 24/7 monitoring data received from parallel but unlike facilities can be distilling into action using digital twins and physics-driven technologies.

As a result, each site becomes a special snowflake—a site unique unto itself, with its own idiosyncratic history and methodology for asset management processes. This in turn puts a hard limit on your ability to monitor and optimize all systems at all sites at the same time. The networks of sensors that automation personnel configure at each site might be similar at a glance but very different underneath, making it hard to gather the data needed for detecting issues and assessing the health of all systems across all sites.

What's needed is a plan that weaves together the capabilities of technology and the essential facts of all systems. Getting to this means looking at three fundamental aspects of the installations—how components are mapped, how they fit together in their unique building structures, and how the laws of physics pertain to the elements that flow within them. Only when you have these pieces of information at hand can you start to realize the full power of today's asset management technologies and understand, at a nuanced level, the performance of individual systems.

The causes and limitations of "one-off" site monitoring

Put another way, the central challenge for many asset managers is the need to deal with 24/7 monitoring data they're receiving from parallel but unlike facilities, and somehow distill that data into action they can take through predictive maintenance and meaningful alerts. Most find ways to effectively manage each site individually, but wish for a way to view the state and progress of processes at all sites through a single piece of software.

How do diverse installations get to this state of misalignment? Sometimes it's because the technology itself is mismatched—older sensors at one site, or a newer proprietary monitoring tool at another. Sometimes it's the human factor, where the best attempts to organize the data for one site completely diverge from the way data is organized at the others. Sometimes it's the result of cost-cutting along the way, such as trying to save time by configuring the new control system at one facility by simply copying over the programming from another.

Other reasons that sites grow unique from one another range from differing ages of construction and generations of equipment to the variety of vendors and software versions used at time of installation. There's also the unavoidable human factors—different sites are typically subjected to different operating philosophies, usage patterns, and levels of wear and tear over time.

Whatever the reason, asset managers end up with one-off, dedicated solutions for each site. Rather than orchestrating all sites' activities into a larger system, they end up specializing in the various quirks of each location, and doing their best to manage accordingly. This approach can get the job done, but it doesn't scale, and it misses out on specific risks and opportunities associated with each site's unique conditions.

As technology evolves in ways that can help asset managers gain efficiency, it's important to look at how doing a bit of work up front can help streamline and strengthen the data gathering, reporting, and analytics processes.

Normalizing installations so you can monitor them in parallel

To understand better what's necessary for improving analytics, let's try working backwards from the ideal solution.

Imagine looking at one software tool that gives you a comprehensive view into all relevant activities at all of your sites. A dashboard shows the health of various systems, alerts where conditions are suspicious, and opportunities for predictive maintenance. You can drill down to get details on one part of a system, as well as construct queries that tell you specific information. Powering all of this is an analytics engine that sees comprehensively the details of each site, and reports only the most accurate information based on unique site variables.

In order to achieve these kinds of system-wide analytics, you need a solution that does three key things. First, it readily identifies every part of every installation—what the part is, what sensors are attached to it, and what it's physically expected to do. Second, it maps all of this knowledge to an underlying data model that can turn each component and connector into a digitized entity, capable of having metadata attached to it so that analytics can be performed. Lastly, it supports algorithms (simple or complex) that can query each entity about its characteristics.

In other words, when you select a query that says, *Show me all pumps on the chilled water side of my chiller plant that are consuming more energy than expected for the operating speed they report*, the solution can differentiate which components at which sites are chilled water pumps; determine the energy usage metadata of each one; verify those results against basic laws of physics to ensure they're sensible; and deliver the final report to your user interface.

Notice the verification step. That's key to sorting out the true results from any false positives. And it can only be accomplished when algorithms track the physical behavior of individual, known system components that are understood by the solution itself to be part of an overlying schematic design.

Which brings us back to the snowflake problem. Getting to complex, reliable analysis across multiple sites is a matter of normalizing those sites—abstracting their commonalities, while also accounting for whatever factors make each individual site, installation, or component unique. Specifically, you need to normalize:

- **Tagging** – the identifier schemes you use when labeling system components and connectors
- **Structure** – the way you capture, understand, and maintain system schematics at a digital level
- **Behavior** – the expected flow (of water, heat, electrical current, or whatever) based on physical laws

Table 1 organizes these normalization activities into three roughly sequential steps. The details of the table are discussed in the remainder of this article.

	To capture...	Normalize your...	Using...
1	Precision	Point mapping	Robust, industry-standard asset tags
2	Structure	System schematics	Digital twin technology
3	Insights	System behaviors	Machine learning algorithms based on digital twin details and laws of physics

Table 1. Three steps to normalizing your assets across installations.

The power of effective point mapping

24/7 monitoring solutions, including building automation systems and distributed control systems, use data from sensors to indicate the value of various points. A point is the target process value of a system component, usually measured over time, such as the desired temperature change in a cooling tower over the course of an hour. Point mapping refers to the way these points are reflected digitally, in a way that can be consumed by software and manipulated for reporting and analytics purposes.

As sensors produce data, the point values come into the monitoring solution and are stored in a database. The tricky thing about databases is that they are as simple or as complex as you want them to be. It's easy enough to store data that doesn't make any sense or isn't very helpful. What you want is data that is crisp and specific to the components and processes you're trying to monitor. In other words, it's one thing to look at a value, and it's quite another thing to have that value provide meaning to what you're trying to accomplish.

Having a sensible, internally consistent naming convention ultimately sets you up to apply meaningful tagging.

For example, let's say the monitoring system in a multi-chiller environment reports a value of 7 for a point labeled chCdwDt. A human operator might parse this label to determine what it's referring to, and figure out it's a chiller (ch) reporting its condensed water differential temperature (CdwDt). But he's not necessarily going to know the details: What chiller does it refer to? There's no number included with the label. Some chillers have both internal and external temperature sensors—which one is it reporting? Is the value of 7 in degrees Celsius or some other measure? And so on.

Most asset managers and operations personnel recognize that the way to get around this is to use a naming convention that standardizes the labels you use for identifying system components and points. This is the correct approach, but it usually doesn't get taken far enough. Ad hoc naming conventions produce various limitations:

- **Lack of depth.** One operator programming a system might use a naming convention he finds perfectly logical, but it might not adequately account for all of the shades of difference among the system components that end up being useful for monitoring purposes.
- **Lack of consistency.** No matter how solid the scheme is that one operator follows when programming a system, the next operator to come along—whether it's at the same site or a new one—isn't necessarily going to follow the same convention.
- **Lack of clarity.** Some older monitoring systems limit the character length of the names you can apply, which only increases the ambiguity factor of these other limitations.

The lack of these criteria ultimately produces data sets that are unhelpful and incomplete. At best, you'll end up with a total count of your assets and the sensor readings associated with each one; but you won't gain an understanding of how they're connected to and impact one another—at least, not on the digital level.

Having a sensible, internally consistent naming convention ultimately sets you up to apply meaningful tagging. Tags are the classification data you apply to each point in a system, capturing the relevant qualities that make them similar or different, so you can run abstract queries such as the one in the previous example. A given pump might have the tags "pump" and "chilled" and "water" so that when you run a query against these three tags, your monitoring system can tell you information about chilled water pumps, or any similar point that might be relevant.

And so, effective monitoring and analysis relies on a well-thought out, standardized tagging scheme for all point mapping across installations at all of your sites. The tags you assign serve to constrain the query results delivered by your monitoring solution, so that only relevant data is used for reporting and analytics. If the data is diffuse, nonspecific, inconsistent, or incoherent, your attempts to use the data for optimizing operations will suffer and fail.

An effective tagging scheme will account for all components and connectors in your monitoring system, the details about each one, and the parameters for point measurement that are required for performing machine learning and data analysis. The good news is that publicly documented industry standards, such as Haystack, exist for just this purpose. Find the standard that best suits your needs, and take the time to learn about implementing it so that the instrumentation of sensors in your environments can be ideally oriented for efficiency gains.

Digital twin: Turning your schematics into ones and zeroes

Once you have the fundaments of effective point mapping taken care of, you still need a way for your monitoring solution to understand how all of the points are connected, and what impacts they have on one another. For that, you need a digital medium that understands the tags and uses them to model each entire installation, so that the network of physical activity we observe at a human level is replicated in what is called a digital twin.

You might think you have this already, if your system schematics are captured in a portable document file (PDF) or similar electronic format. But a PDF, while digital, is basically only a photograph—it portrays what human beings perceive, not the deeper modular connections that need to be read by analytics systems. A digital twin, by contrast, is a fully evolved digital replica of the whole system—its assets, components, processes, and connectors, along with the physical and systemic characteristics of each entity in its scope. This detailed understanding enables a much more granular level of data manipulation, and sets the stage for advanced analytics.

Figure 1 shows at a high level how the digital twin evolves as a replica of your built environment.

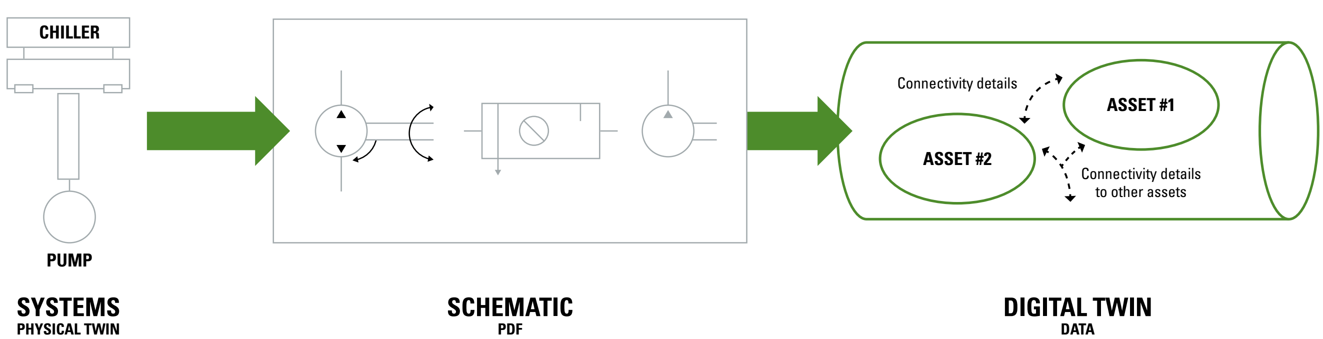

Figure 1. The digital twin replicates the built environment in a detailed electronic format.

The beauty of this approach is that it not only makes sense of your systems for all of their intricate parts and relationships, but it also sets you up to parse those relationships in new and meaningful ways. Machine learning does its magic by aggregating all of the data it sees and turning that data into useful and insightful patterns that map to an organization's business needs. By creating a digital twin of your installations, you can engage data scientists to write algorithms that operate on the nuances of every piece of every system, accounting for the unique connectivity of certain parts while taking advantage of a more holistic, integrated system view.

When it comes to finding the right algorithms to produce your insights, physics itself plays a key role. Too often, monitoring systems generate false positives based on faulty information that violates clear laws of physics—in particular, the physics of flows. Having the digital twin means your data scientists can inform your solution electronically about the expected physical behaviors within a system or installation, and your operators can generate rules and configure alerts based on this new level of representation.

Combined with effective point mapping and the application of new algorithms that understand physics and flows, digital twins form the basis for creating unified analytics queries across all of your dissimilar installations, providing the next-level analysis you require for optimizing and maintaining asset performance. By working with an analytics vendor and platform that can store your entities with comprehensive tagging and query those tags intelligently using a digital twin, you can upgrade your monitoring processes to benefit from the latest machine learning and advanced analytics technologies.

About Tignis

Seattle-based Tignis provides unique physics-driven analytics for connected mechanical systems, utilizing digital twin and machine learning technologies. Tignis increases the reliability of connected mechanical systems by automatically monitoring and learning, continuously detecting threats to reliability—even on diverse and complex systems, and precisely identifying and predicting operational impacts. Tignis enables you to simplify system monitoring processes, filter out the "noise" of false positives, and gain a more durable, digital foundation for understanding and mapping the processes and priorities you care about day-to-day.

For more information on applying physics-driven analytics to your systems monitoring data, visit www.tignis.com

info@tignis.com | (206) 745-9866