

OK, I've Got 24/7 Monitoring Data. Now What?



The most exciting asset management trend in recent years has been the shift from manual, route-based condition monitoring to digital, real-time processes driven by sensors and the Internet of Things (IoT).

Overall, the promise of an IoT revolution for asset management has given us a whole new toolset for reducing wasted effort and tackling common problems. It has delivered an impressive array of platforms and technologies that stand to reimagine all the ways we safeguard and optimize system availability and performance.

**Also, data.
Lots—and lots and lots—of data.**

So much data that it can feel like we've lost the forest for the trees. Potential advantages only become actual benefits if you can build a usable methodology that applies 24/7 data in smart ways that yield meaningful operational improvement. Likewise, data only becomes useful if it's captured reliably—as opposed to just data for data's sake, such as from a faulty sensor—and expressed in ways that help you take meaningful action.

The good news? Just as technology brings complexity, it also brings computing power and innovation for managing it. When it comes to monitoring your installations with 24/7 IoT sensors, there are ways you can sort through the data, grasp intelligent site-specific details to take action on, and hopefully stop feeling like a dog chasing its tail.

Key Findings

- IoT enables new levels of operational efficiency and predictive maintenance.
- Digital twins and physics-driven technologies help eliminate false positives to identify real errors.
- Physics-driven logic identifies if data is sensible from scientific perspectives, and whether it's anomalous or reliable.



The Quagmire of Configuration: Rules, rules, rules

To monitor and manage a system end-to-end, you rely on logical thresholds that you know to be within the boundaries of the system's health, and then you instrument and program the system with sensors that trip whenever a threshold is exceeded or otherwise deviates from its norm.

It seems simple enough. But over time, rules grow old. They become obsolete. They no longer pertain to real-life conditions. Also, when you have lots of rules, they can also become redundant with or contradictory to one another. All in all, over the useful life of a system, the integrity of its monitoring solution can be gradually compromised by ineffective configuration, and no amount of human interaction will ever really seem to fully solve the problem.

For example, you set a rule that says: Raise an alarm if the flow into Valve A does NOT equal the sum of the chiller water flows from Chiller 1 and Chiller 2. And maybe another rule saying: Alert a technician if energy usage in this chiller exceeds x kilowatts. All's well—until you alter the environment in a way that the programmed thresholds no longer make sense. Maybe you add another chiller, and meanwhile, the usage patterns of the facility change. Suddenly, you're getting false alarms, duplicate alerts, and other largely unhelpful kinds of reporting.

What to do? You can disable the alerts in question, but that leaves you exposed to dangerous or otherwise unwanted conditions. You can ignore the alarms, but again you do so at your own risk, and anyway, doesn't that defeat the purpose of monitoring in the first place? Or, of course, you can reconfigure the rules every time the environment changes. But change is constant, and soon it's a full-time job just keeping pace. As rules proliferate and compound, the duty of managing and prioritizing them grows exponentially more complex, reducing your efficiency and often forcing operational tradeoffs for the sake of practicality.

So, you end up choosing between a functional but faulty environment, eroding the value you get from your monitoring system investment; or, you commit your organization to endless churn, detracting from the time and resources you could be spending to analyze and address each issue's actual root cause.

Science and Logic to the Rescue, Once Again

Advancements in machine learning technology can help with this problem. By overlaying detailed schematics about each system and its IoT sensors with known facts about the physics of how each element should behave, you can build toolsets that go a long way to automatically discovering design faults, detecting potentially miscalibrated sensors, and proposing corrective action.

To do this, you need a platform that not only gathers the data from the sensors, but applies basic logic to it in critical ways:

- Is the data sensible from a basic scientific perspective? For example, does the power usage of Pump A increase when the pump speed increases? If not—was the pump speed point accidentally swapped with Pump B when you programmed the automation controller?



- Is the data an anomaly, maybe produced by a one-time environmental event (such as an unusually hot day) or a human factor (such as a visiting contractor accidentally whacking a pipe with his stepladder)?
- Do the readings across two tightly coupled components (such as a valve that water flows through, and a tank where the water is collected) match up in a way that makes physical sense?

And so on. It's the kind of logic that human operators would apply anyway when they arrive to diagnose a problematic situation. But with the right information about the site, computers can use machine learning to do it faster, more proactively, and at a much deeper level than human observation traditionally allows.

Normalizing sensor data across parallel, but unlike, systems

Do any of these sound familiar?

- "An alarm at one of our sites keep tripping—for reasons that don't even apply to that site's particular systems."
- "I copied over the programming from one of our sites to another, but the names don't match so I get faulty results."
- "One of our sites is older than the others, so its installations need to be equipped with all different sensors."
- "My team spent weeks figuring out sensor data. But when our new site comes online, the data will all be different."

Sites that perform similar functions at different locations invariably contain unlike installations. The systems at each site differ in age, configuration, system components, and maintenance history, as well as in the types of facility functions they support. For all of these reasons and quite a few more, the rules we write for one installation are not necessarily useful for the others.

So, when it comes to 24/7 monitoring, we can talk a big game about gaining efficiency across systems by automating error detection and monitoring for common issues. But in reality, when it comes down to it, every site tends to be a special snowflake. How do we gain efficiency for THAT?

The answer again lies in how you approach the data, through an intersection of each installation's schematics with the natural laws of physics. Advanced analytics can "normalize" data across systems by understanding each site's unique characteristics, and then applying rule sets that check the sanity of all readings against scientific and human logic. With the right machine learning algorithms, you can readily and steadily monitor the actual conditions at your sites for the high-priority success and failure conditions you require.

With the right information about the site, computers can use machine learning to do it faster, more proactively, and at a much deeper level than human observation traditionally allows.



Eliminating False Positives so you can Identify Real Errors

This discussion has dwelled on a central irony: 24/7 condition monitoring was intended to help make your sites run like clockwork—but the reality is that it has also introduced new inefficiencies, even as you deploy it to help make your existing inefficiencies a thing of the past.

The true value in addressing data complexity, of course, is that it helps organizations remain true to their aim of maintaining safe, efficient facilities running effective systems with accurate monitoring. This means going beyond traditional tactics that yield unnecessary and meaningless fault reporting and instead focusing on root causes. By doing so, you make your technology investments perform for you in real ways that serve your mission, rather than distract you with false information, make you wonder what faults actually exist, and/or decimate your bottom line.

Rather than relying on basic pre-programmed rules, you can develop algorithms that accommodate the variance in sensor behavior both across and within existing systems. In this way, analytics helps IoT reach its full potential for asset management, allowing visibility into issues that were previously undetectable, and enabling new levels of operational efficiency in the forms of remote monitoring and better predictive maintenance.

See what Tignis can do for your business.

For more information on applying physics-driven analytics to your systems monitoring data, visit www.cohu.com/tignis | Info@cohu.com | Tel: +1 206.745.9866

© 2025 Cohu, Inc.: All rights reserved.

REV20250410

