WHITE PAPER

The Evolution of Industrial AI for Reliability

and its role AMS product and service offerings

AMS™

EMERSON™

Table of Contents

Introduction

Artificial Intelligence (AI) has been at work in the condition monitoring industry long before it was recognized as such or became as ubiquitous as it is today. In this briefing, we explore the genesis of AI, with specific focus on the condition monitoring industry, the forces that initially gave rise to it, and how the market continues to evolve along with AI, pervasive IIoT sensing, and new commercial structures for packaging offerings designed to deliver condition monitoring as a service. We also explore the ways that AI has been present in Emerson's AMS offerings for several decades already and how that continues to evolve at present.

The Tipping Point

Not every popular or seemingly new idea in the industrial market is necessarily a new one. In ***The Tipping Point***[1], author Malcolm Gladwell describes the seemingly magic moment when an idea that may have been around for a long time goes viral – almost always without aid of large marketing campaigns like Super Bowl commercials and celebrity spokespersons. Word of mouth is instead the underlying mechanism, aided today by social media. Word of mouth may seem antiquated, but is alive and well, even in the age of social media which is simply today's way of an individual "telling someone else" – but instead of one at a time, perhaps hundreds or thousands of followers at a time. Indeed, word of mouth is so powerful because – unlike advertising – we all tend to place more credence in millions of people talking about an idea than in somebody being paid to say nice things about a product or service.

For AI, that tipping point can be traced to November 30, 2022[2]. Many people will remember the first time they logged in to ask ChatGPT a question, only to be amazed that there seemed to be a highly articulate, seasoned expert on the other end of the chat with instant access to the world's accumulated knowledge. We understood in that moment that this was fundamentally different than someone looking up information and merely regurgitating it in cut-and-paste fashion. This was somehow acting as a real human would act – consulting multiple sources of information, coalescing it into a coherent, thoughtful response, drawing conclusions and extrapolations, and then explaining it clearly like we were talking to a wise professor, mentoring us in response to an honest question. In fact, the only clue that this couldn't possibly be human was the fact that it could respond with 200-300 well-composed words in just seconds – unlike an actual human at the other end of the chat who would need at least a few minutes to ***type*** the information, let alone ***research*** it.

[1] **The Tipping Point: How Little Things Can Make a Big Difference by Malcolm Gladwell | Goodreads**

[2] Ávila, N. "**The Tipping Point of AI**" Forbes, February 16, 2024. Retrieved March 24, 2025.

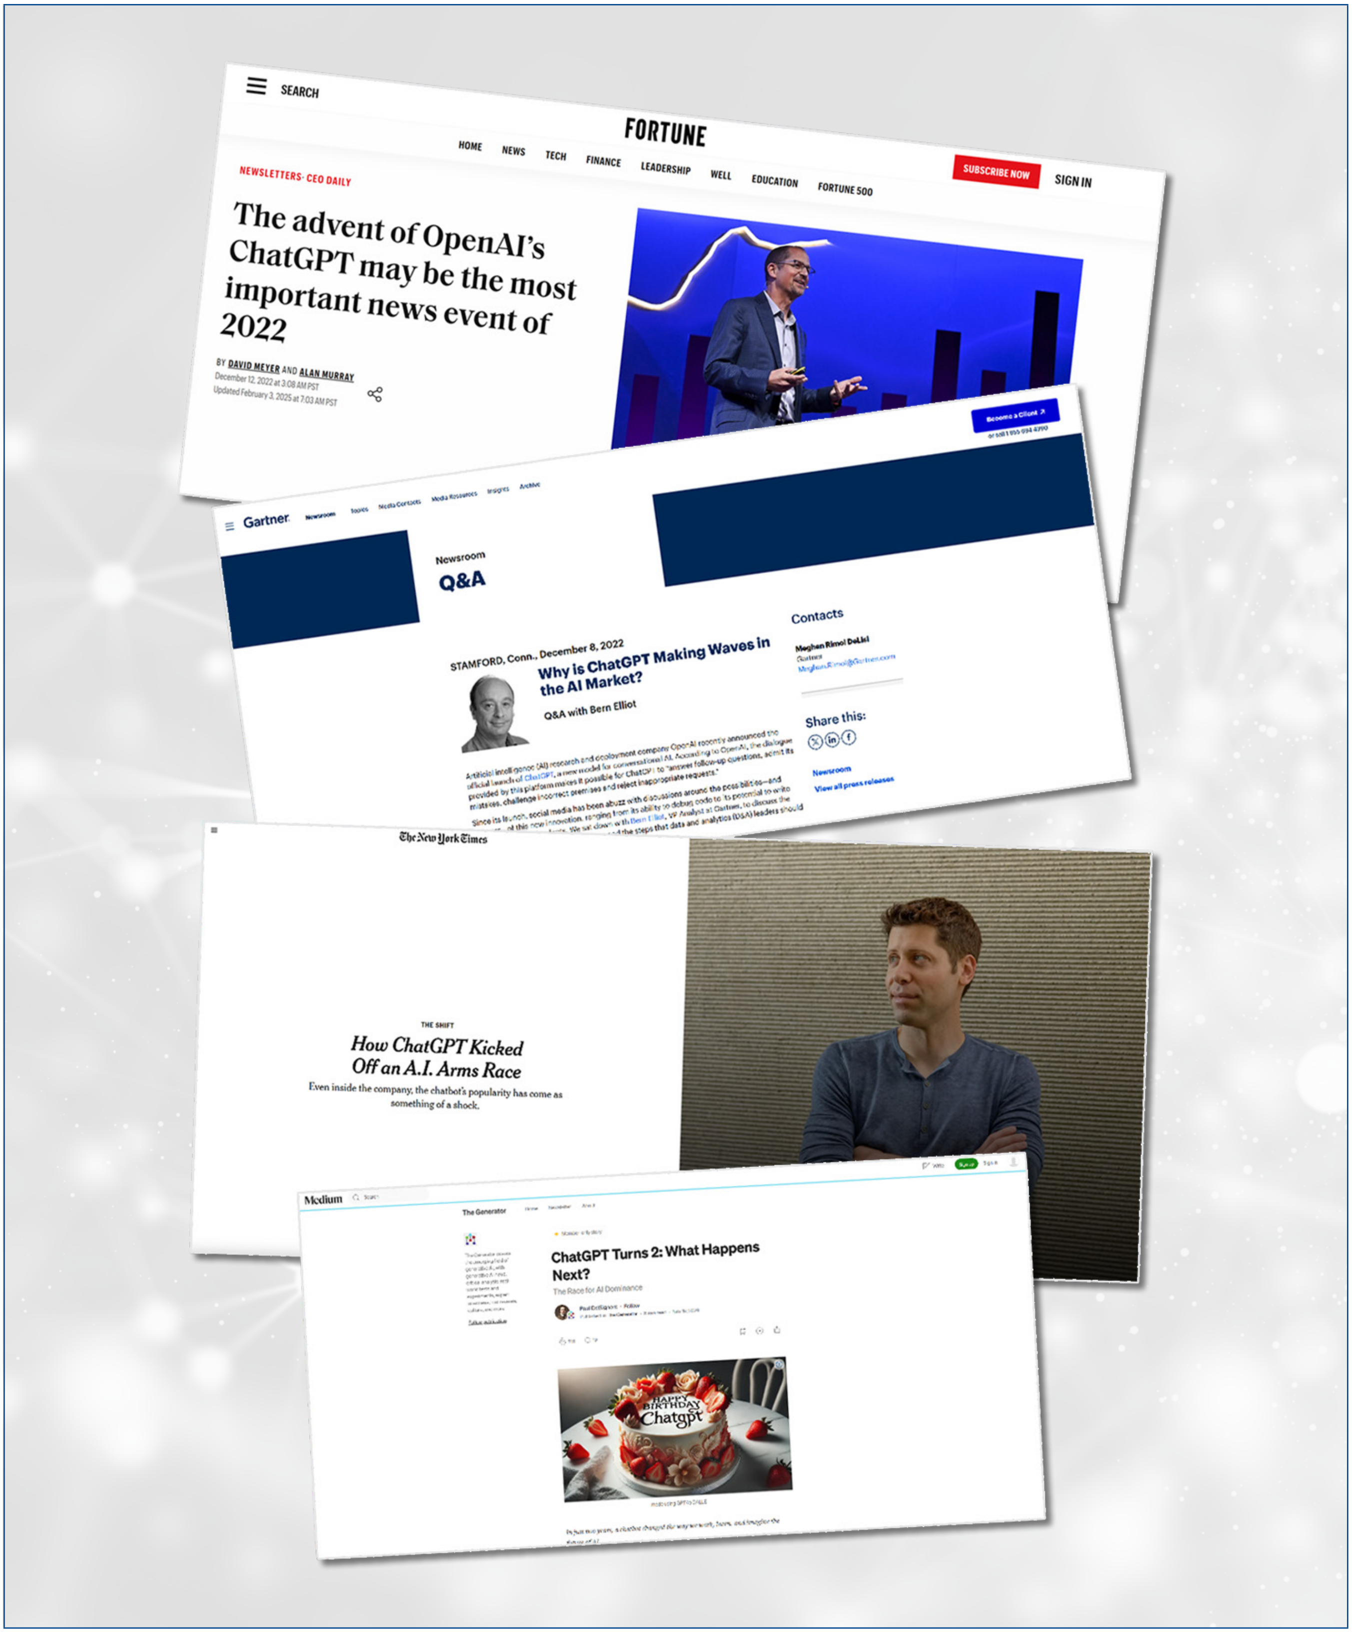

The release of ChatGPT in November 2022 is widely acknowledged as the AI tipping point.

Although the phrase is overused, ***"this changes everything"*** was never more true. And, like all true tipping points, we told someone else what we'd experienced, adding to the already blazing ChatGPT wildfire.

Seldom has a technology so quickly transformed the competitive landscape across every industry, from stock photo libraries to market research, to criminal justice, to medical diagnostics – to name just a few. Hundreds more examples could be named, and it is indeed hard to find an industry that has not been touched by AI[3]. To understand just how far we've come in such a short period of time in terms of moving from a curiosity to a must-have, consider that when 1,500 senior business leaders were asked about AI in 2017, only 17% even knew what it was within the context of their own companies, with fewer still able to cogently understand how it might impact their business beyond just vague generalities[4]. But now, just 8 years later, you'd be hard pressed to find even one of those 1,500 senior business leaders that isn't talking about AI in imperative terms.

Ironically, AI has been present for at least two decades in many of the products we use – many times without us realizing it – because we didn't explicitly use the term "AI" to describe it. Think Siri on your iPhone. Or your favorite map program when driving and its ability to reroute you when it detects traffic is becoming congested. Or the text message you get when your credit card company detects unusual activity and wants to be sure it's actually you before approving a purchase. As we will see, AI has been hard at work under the hood in industrial products as well, but has often had to overcome substantial (and perhaps well-earned) skepticism when revealed.

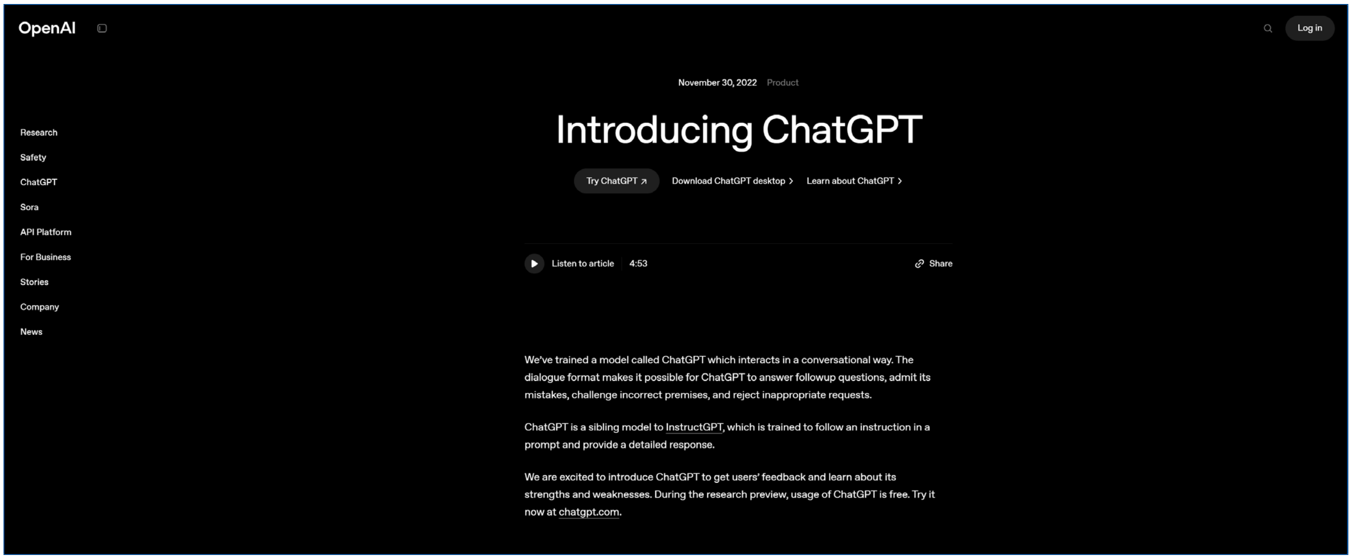

On the day it was launched, November 30, 2022, very little in the understated user interface for ChatGPT hinted at the profound impact it was going to have.

[3] Allen, J.R., West, D.M. "**How Artificial Intelligence is Transforming the World – Chapter II: Applications in Diverse Sectors**" Brookings Institute, April 24, 2018. Retrieved March 24, 2025.

[4] Davenport, T., Loucks, J., Schatsky, D., "**Bullish on the Business Value of Cognitive: Leaders in cognitive and AI weigh in on what's working and what's next**" (Deloitte, 2017), p. 3.

Coming to terms with AI

Before we delve deeper into the topic, it is important to understand the background of AI itself and the terminology used to describe it.

AI can be broadly characterized into two areas as follows[5,6]:

- **Artificial General Intelligence** (AGI)[7]

  The textbook definition of AGI is artificial intelligence that **equals** that of a human and is able to reason, plan, learn, represent knowledge, use common sense, and speak conversationally using natural language[8]. From a practical standpoint, industrial AGI is generally considered to be a very powerful engine that can be configured to do specific tasks but does not come out of the box pre-configured to do anything. An analogous concept is Microsoft® Excel. All by itself, the application can do nothing. Instead, a workbook with appropriate equations and operators must be developed, and then data must be introduced for the workbook to process. In the same way, AGI is extremely powerful, but must be configured and then 'trained' using data.

AGI tends to conjure up images of white-coated scientists devoted to the care and feeding of a monstrous computer with a mind of its own and processing terabytes of code. In reality, most AI is instead what is known as "narrow" and may often be embedded right in the product if not in the cloud.

[5] SOURCE: "**Understanding the Different Types of AI**" IBM, October 12, 2023. Retrieved March 24, 2025.

[6] In addition to the two categories listed here, there is also so-called Super Artificial Intelligence (SAI) but it is strictly theoretical. According to IBM, and unlike AGI which merely **equals** human intelligence, "*if ever realized, Super AI would think, reason, learn, make judgements and possess cognitive abilities that **surpass** those of human beings.*" Source: ibid.

[7] Also known as strong AI or broad AI. However, "strong" can have unintended connotations by implying that every other kind of AI is weak and thus relatively useless. Ironically, strong AI represents only **potential** strength – not **actual** strength – as all AI must eventually be configured to be somewhat narrow to solve actual problems.

[8] This list is loosely based on topics appearing in the following AI textbooks: **Russell & Norvig 2003**, **Luger & Stubblefield 2004**, **Poole, Mackworth & Goebel 1998** and **Nilsson 1998**.

- **Artificial Narrow Intelligence** (aka: purpose-built AI, narrow AI, weak AI)[9,10]

  This type of AI can perhaps be most easily understood by again considering our earlier example of Microsoft Excel. Narrow or purpose-built AI is like a pre-configured workbook in Excel that already knows what to do with the data. While it still needs to train itself on the data to distinguish normal and abnormal, the equations and operators are already present and do not have to be developed. It also generally already knows something about the type of data it will be operating on. For example, if it is operating on machinery vibration data, it will know whether the data is acceleration, velocity, or displacement, and whether the originating machine is a motor or pump, and perhaps even whether a pump is handling water or a fluid of a different viscosity. It will also typically know what speed the machine is running at so that vibration frequencies can be related to this running speed as a reference. Purpose-built AI is much like clay that has been formed into something very specific, with a particular application in mind that will affect size, shape, and finish of the piece. However, it can be argued that there is really no such thing as AGI and that all AGI is eventually turned into purpose-built AI before it can do anything useful[11].

The AI integrated into today's passenger vehicles exemplifies artificial narrow intelligence. Commercially available AI-powered features include traffic jam assist, predictive adaptive suspension, voice-controlled assistants, and AI-enhanced night vision The AI embedded in today's passenger cars is an example of artificial narrow intelligence. The vehicle shown here, an Audi A8, contains AI-powered features such as traffic jam pilot, predictive adaptive suspension, voice-controlled AI assistant, and AI-assisted night vision.[12]

---

[9] Also known as weak AI. For the same reasons as in footnote 8, the term can have unintended connotations by implying that other techniques to solve the same problem are stronger. The term "narrow" is generally a better descriptor.

[10] Indeed, so-called narrow AI can be very strong in that it can solve very specialized problems very well. Examples include voice assistants like Siri or Alexa, facial recognition software used in security systems, and recommendation engines used by platforms like Netflix or Amazon.

SOURCE: "**Types of Artificial Intelligence**". geeksforgeeks.com. Retrieved March 21, 2025.

[11] This is precisely the argument set forth in the IBM article referenced in footnote 6.

[12] SOURCE: Dupont, T. "**10 Cars with the Most Advanced AI Features Right Now**". *How-To Geek*. Mar 5, 2025.

Within the realm of narrow or purpose-built AI, we can further subclassify as follows:

- **Physics-based** (aka: first principles)

  These characterize the earliest forms[13,14] of AI and are still very useful. They seek to mimic the thought process used by a human expert by breaking a problem down into smaller parts and then asking a series of "if, then" questions to interpret the data. In the world of machinery, for example, we might ask whether the dominant vibration coincides with the machine's running speed or some multiple of running speed. We might then ask what types of bearings the machine has so we can determine whether a frequency coincides with the bearing geometry for a rolling element bearing (outer race, inner race, ball, cage, etc.) or perhaps a fluid whirling frequency for a fluid-film bearing. The idea here is that the data is analyzed according to rules that a human has explicitly programmed based on knowledge of the underlying principles of the application – such as how a particular type of machine vibrates under various malfunctions. Other terms used to describe these types of systems are "rules-based" and "expert". Within the Emerson AMS portfolio, our Nspectr capabilities (originally introduced in our MasterTrend software three decades ago) fall into this category[14].

The Emerson AMS Asset Monitor is an example of a device that uses physics-based AI to determine not only whether an asset is healthy or sick, but to identify what is specifically wrong with the asset, such as a lubrication problem, a bearing problem, an imbalance problem, a misalignment problem, a gear tooth problem, a turbulent flow problem, and many others. Because this AI is embedded in the device and does not require a connection to a server – whether on-premises or in the cloud – we refer to this as edge analytics or edge AI.

[13] Our Nspectr tool within MasterTrend software dates back to the mid-1990s. Since that time, we have had three successive generations of software: RBMware, Machinery Health Manager (MHM), and today's AMS Machine Works.

[14] Other companies in the condition monitoring space were likewise releasing their own expert systems during this period with the earliest of these appearing in 1992. That system was extremely crude by today's standards as it ran under DOS, was not networked, could not automatically analyze incoming data, and could only perform on-demand audits of data fed to it in special batches using a special filetype. It was also unable to analyze machinery with rolling element bearings.

- **Data science-based**

  These tend to fall into the realm of statistical analysis, pattern recognition, and the ability to distinguish normal data from abnormal data. Examples might include a system with multiple variables. If variables A and B trend in the same direction 99% of the time, the software might be able to figure out that when A and B trend in opposite directions (reverse correlation), this is highly abnormal based purely on statistical analysis. The software thus might not be able to ascertain why the variables are trending in opposite directions – it would merely be able to flag this condition as highly anomalous behavior.

Machine Learning

Another term that is heard quite frequently is **machine learning**[15,16]. However, this can be especially confusing in the field of machinery condition monitoring because it may not be apparent that the machine in this case is not the one on which we are measuring condition. The machine is instead the software program – sometimes called the AI engine. Machine learning is simply describing that the AI engine is able to refine itself without explicit programming and thus becomes better and better over time when exposed to more data. In the world of machinery condition monitoring, this might mean more data from the same mechanical assets, or more data from a larger population of assets[17]. Either way, more data results in an algorithm that is continually learning and forming a closed feedback loop to improve itself over time.

To better understand machine learning, let's expand on our previous example where variables A and B are trending in opposite directions. When no machine learning is present, the system is purely reactive. It can only tell the user that something is wrong. In contrast, a system with machine learning, over time, might discover that A and B only trend in opposite directions when C is at or above some threshold level because it is looking for patterns and relationships in data rather than following a rote set of rules that never change. Indeed, it might well discover that the behavior of A and B is perfectly normal when C exceeds a threshold, and that this behavior of C is itself perfectly normal, like a machine running in the reverse direction because the process allows it to run either forwards or backwards. The only anomaly might then be when C is underneath its threshold and yet A and B are trending in opposite directions. When a system is able to make such discoveries and improve itself over time without being explicitly programmed to make a particular improvement, this is machine learning. In contrast, if we independently discovered the relationship between A, B, and C and explicitly programmed a hard-coded rule into the system, we would be in the realm of a reactive system rather than a learning system. Such a system cannot learn – it can only react – and if it is to react differently, it must be explicitly programmed as such.

As a general rule, physics-based systems require explicit revision to its algorithms to learn whereas systems relying on data science can instead (but do not have to) incorporate machine learning. Systems that do not incorporate machine learning are sometimes also referred to as reactive systems[18].

---

[15] The IBM article of footnote 6 refers to this as "limited memory AI" and explains that it can indeed learn over time.

[16] Another term frequently heard within the context of machine learning is "neural net(work)". Strictly speaking, the terms are not interchangeable as a neural network is a special type of machine learning that mimics the way the human brain works.

[17] This is indeed one reason why hosted, cloud-based AI solutions can be appealing: because they can have access to the data from many thousands of similar assets and thus improve their algorithms more quickly and robustly than if limited to a smaller data set.

[18] Refer again to the IBM article of footnote 6.

Generative AI

Without question, no other type of AI has had as much influence on the tipping point of pushing AI into the vernacular of the culture as has Generative AI (GAI). The reason is quite simple: GAI makes AI ***tangible*** for people. No more is it some type of abstract concept that is meaningful mostly to data scientists and technology nerds; instead, you can see in real time what AI can do.

But what exactly is generative AI?

In a nutshell, GAI produces things we normally associate with human creativity such as storytelling, drawing, having an impromptu conversation, and composing music – to name just a few. We earlier alluded to the epiphany that many people had when they used ChatGPT - or similar programs - for the first time. Suddenly, ***something*** (or was it ***someone***?) was actually talking to them using natural language, idioms, and perhaps even slang. And if that wasn't enough, they perhaps requested their first AI-generated art. In this case, ***seeing*** literally ***became*** believing.

Later, we will explain the surprising impact that generative AI has had in removing objections to the use of AI within the condition monitoring industry.

"The Next Rembrandt" is an AI-generated painting that brought the legendary Dutch painter back to life. By analyzing Rembrandt's body of work and using machine learning algorithms, this painting was created, imitating the artist's style and subject matter[19].

[19] Image and descriptive verbiage are from the article "**The Ten Most Influential Works of AI Art**" by Michael Filimowicz, PhD (**Higher Neurons – Medium Magazine**, June 4, 2023).

The roots of AI: a need for productivity

If we think about Emerson's roots in the condition monitoring space, they extend back to the acquisition[20] of the CSI brand of vibration monitoring technologies (today, part of the AMS brand) and its emergence as a leader in the portable data collector space in the mid-1980s[21]. Route-based portable data collectors revolutionized the industry by allowing vibration monitoring / condition monitoring to become economically viable for less-critical machines. Prior to this, vibration monitoring was relegated mostly to permanent monitoring systems on critical machinery and this category of products indeed characterizes our acquisition of epro GmbH in 2009[22,23] for inclusion in our AMS portfolio. This background will provide helpful context for how the industry has evolved and why AI has become such an indispensable component in solving customer problems today.

Route-based portable data collection, as mentioned above, revolutionized the industry by allowing condition-based maintenance to be economically practiced on machines that could not justify a permanent monitoring system. A typical 250 KBPD refinery, for example, might have 100 machines that are classified as critical and thus suitable candidates for continuous, online monitoring, whereas there are nearly 10,000 rotating machines in such a plant. This means that literally 99% of the machines in the plant were not viable candidates for a condition-based maintenance approach until the advent of portable data collectors. Portables meant that – for the first time – ***any*** machine could be addressed with periodically collected vibration and other data as long as someone could walk out to the machine, place a handheld sensor on the bearings at various locations, and collect the data.

While this did indeed revolutionize the industry and dramatically improve machinery reliability, it had two primary drawbacks:

1) It required semi-skilled, manual labor to ***collect*** the data

2) It required highly-skilled, manual labor to ***interpret*** the data

For customers in machinery-intensive industries, the ability to justify an in-house team of specialists to both collect and interpret data was not a problem. Providers for such customers focused on selling the hardware and software infrastructure, training customers how to use it, and providing support. The prevailing model was clearly DIY for customers; here, providers were focused on ***delivering*** technology. The customer, in contrast, was primarily responsible for ***using*** the technology and ultimately obtaining the value from it.

[20] "**Company News: Emerson Electric to buy Computational Systems**" Oct 18, 1997. *New York Times*. Retrieved March 21, 2025.

[21] CSI was founded in 1984 and quickly brought two products to market: the 2100 machinery analyzer / route-based data collector and companion MasterTrend software.

[22] "**Emerson buys online machinery monitoring capability**" *CONTROL Magazine*. April 7, 2009. Retrieved March 21, 2025.

[23] Epro traces its roots to 1970. The decade between 1960-1970 marked the emergence of permanently installed vibration sensors and continuous monitoring systems as generally accepted good engineering practice on critical rotating machinery.

In the DIY model, the provider is focused only on delivering tools and technology to the customer – not results. The customer, in turn, is ultimately responsible for results because they are the user of the provider's tools / technology.

Other providers quickly realized, however, that not all customers wanted to treat condition monitoring and vibration analysis as a core competence. They capitalized on an opportunity with such customers to run their program for them – both data ***collection*** and data ***analysis*** – and deliver condition monitoring as a service (CMaaS).The typical model was thus monthly or quarterly visits from a service provider that would collect the data from the machines identified by the customer and then analyze the data and issue a report. The report would typically group machines into those that required urgent attention, those that needed to be watched but were not yet critically approaching failure, and those that were entirely healthy. Customers were then responsible for addressing those machines that had been flagged as requiring maintenance, such as balancing, replacement of bearings, changing or adding lubricant, etc. The customer thus required a maintenance department – but not a condition monitoring department – because they had essentially outsourced their CM department. In some cases, these customers were so large and had so many machines that the provider might literally have an office on the customer's premises and be almost indistinguishable from actual employees. However, regardless of how tightly the provider was embedded at the customer, they were still fundamentally a contractor – not an employee.

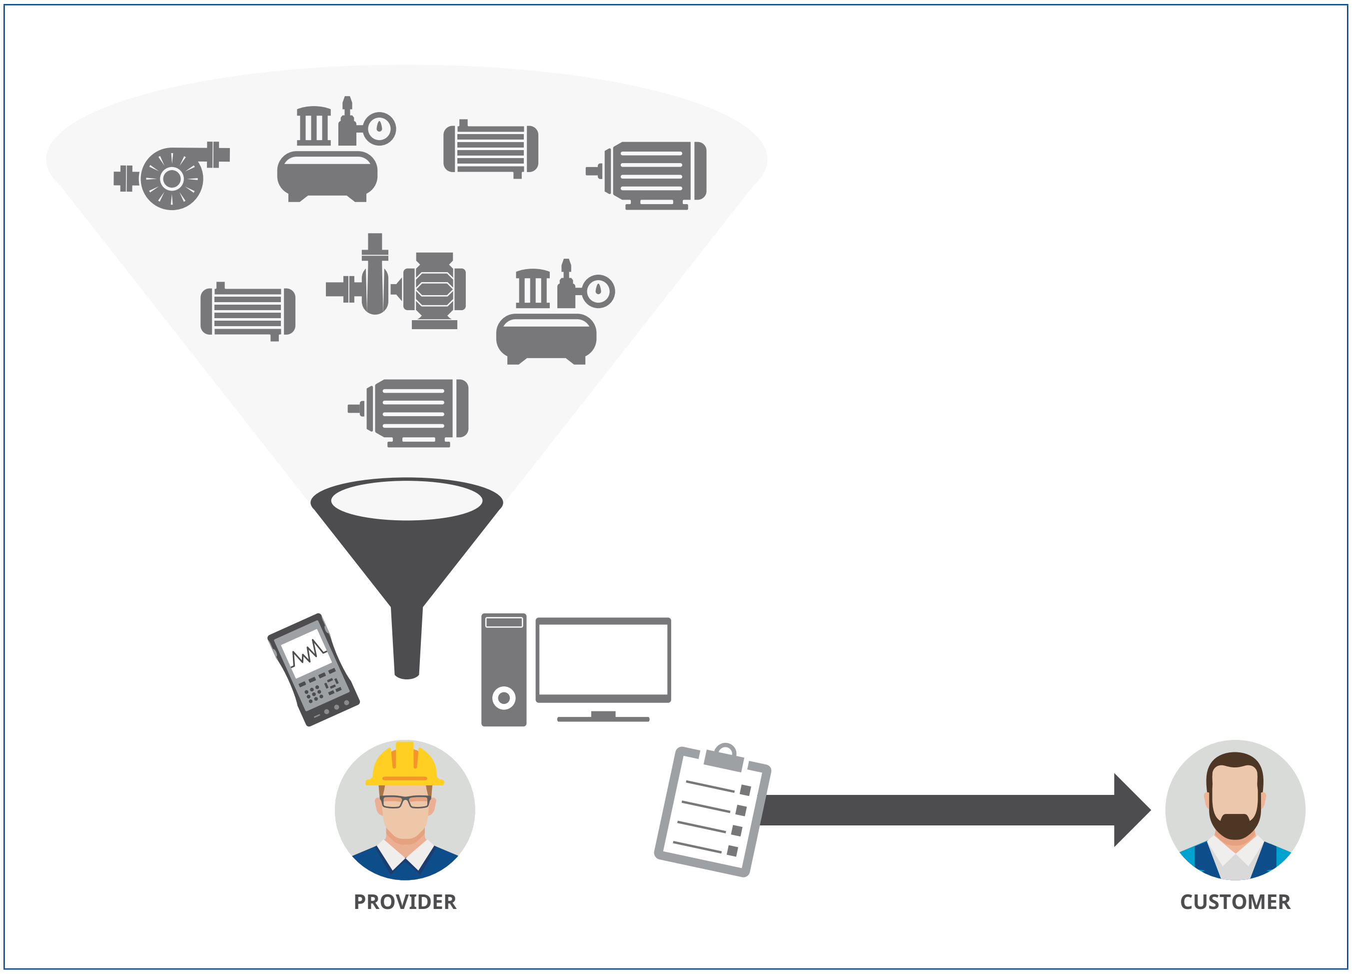

In the outsourced CMaaS model, the provider is focused on delivering actionable information. The customer is responsible simply for acting on that information. The provider in this model owns the infrastructure (hardware, software, computers, etc.) and also owns the task of collecting the data and reducing it into a report containing the actionable information.

While this outsourced CMaaS model opened up an even wider percentage of the market by including customers who did not want to ***use*** the technology and merely wanted the ***results*** of the technology, it was still not without its own set of problems.

- There were no economies of scale; monitoring twice as many machines generally meant twice as much labor. And labor costs were increasing.
- Although highly skilled labor is required to interpret data, it was tedious and costly to use this labor to sift through reams of data looking for problems rather than addressing only that data that clearly indicated problems – and then offering expert advice on those anomalies.

Thus, the outsourced model created challenges for ***providers*** to be more productive with their highly skilled, and thus highly compensated, personnel.

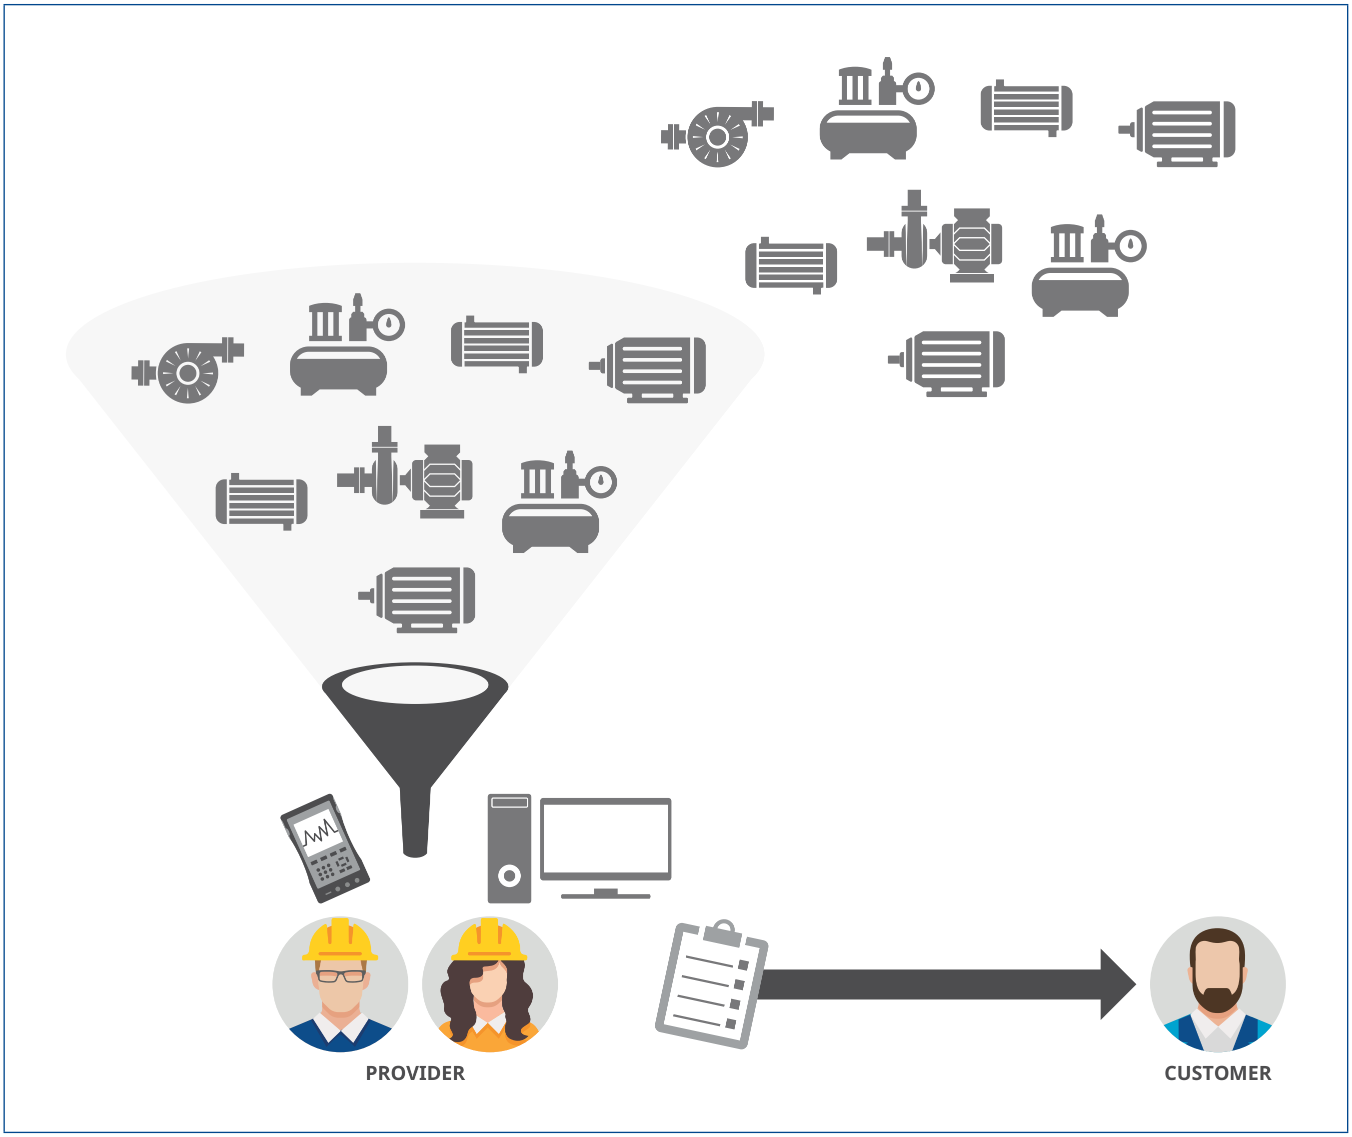

A drawback of the outsourced CMaaS model is that there are few economies of scale. If a provider is to monitor twice as many assets, they must effectively double the labor required to collect it, analyze it, and reduce it to actionable information.

It was natural, then, that much of the impetus for AI arose from the providers themselves, striving to offer CMaaS in ways that had economies of scale and could be more competitive, thus opening the opportunity to justify monitoring for a broader population of assets.

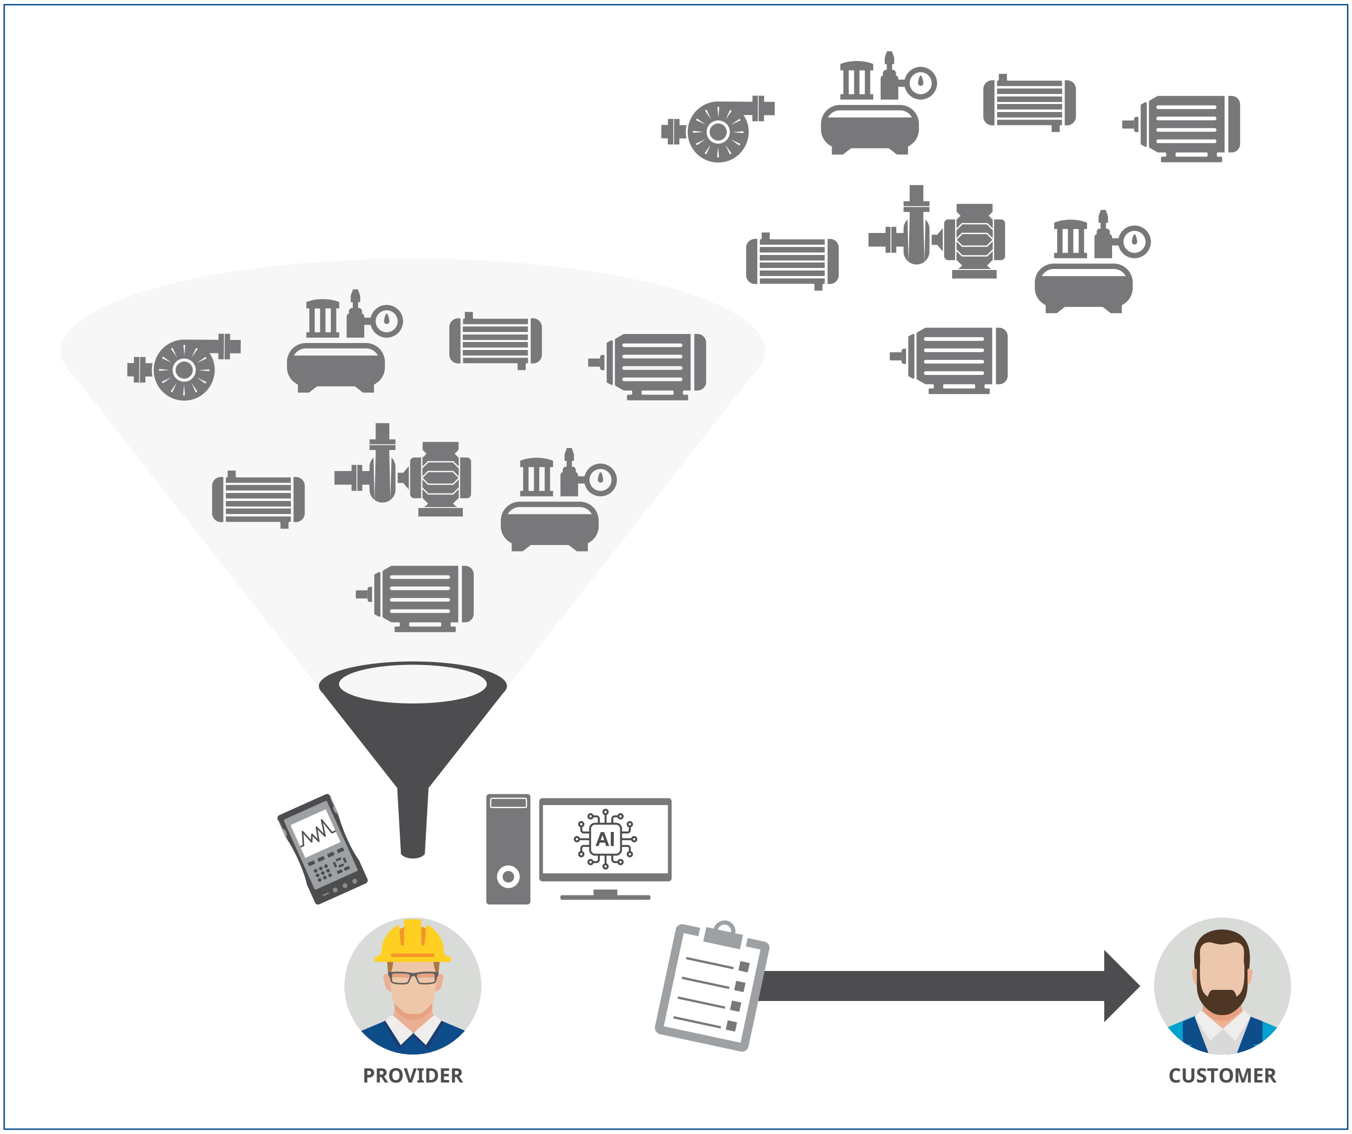

By embedding AI into the data analysis part of the service, the linear relationship between labor and number of assets monitored was no longer inescapable. Here, we have shown that the same number of assets are being monitored as before, but because of AI, only one person is required instead of two.

Notice, however, that although AI has addressed the data analysis part of the equation, it has not addressed the manual labor required to collect the data. Twice as many assets still means twice as much labor.

It was within this context that the next big transformative technology was to present itself: pervasive, wireless IIoT sensing.

Wireless: from niche to mainstream

For years, industry realized that the single largest cost of a permanent data acquisition system was the wiring. The quest for lower wiring costs initially led to wired sensor networks where a single cable could snake its way throughout a plant and sensors could be connected to it, like tributaries to a river. However, these costs were still too restrictive in most cases. The holy grail was wireless, but customers generally had expectations that wireless sensors would 1) be inexpensive, 2) have battery lives of 3 years or more, 3) be capable of not just trended parameters but full waveform data, and 4) have at least reasonably the same frequency response, sampling rates, and other capabilities as would be present with a wired sensor into a hand-held data collector.

Emerson's first foray into this arena actually occurred more than 20 years ago with the release of our 4100RF series wireless system. While it exploited all the available technology of its day, the industry simply had not advanced to the point where industrial wireless could do all the things customers wanted and would allow it to go past its own tipping point of being suitable for more than niche applications. The market may have been ready for wireless to replace a walk-around approach as the dominant scenario, but the technology was not.

In large part, the quest for wireless was driven by the providers themselves: if they were ever going to be able to offer the actual data collection as a service with any economies of scale, there had to be a way to automate the data acquisition – and the data thus acquired had to be comparable to manually gathered data using a wired sensor. Only then could providers hope to bring the cost of outsourced programs down into ranges where more customers could afford to use condition monitoring – and existing customers could afford to monitor more assets in their CM programs.

The 4100RF offering was ahead of its time in terms of size. Indeed, it was what most customers envisioned: a sensor not much larger than a saltshaker and a gateway device the size of a conventional hard drive with an antenna on top. However, while the form factor was in line with expectations, the cost and capabilities left a gap between what industry needed and what technology was able to deliver. It would be two more generations of product before Emerson would be able to bring today's AMS Wireless Vibration Monitor to market with all the capabilities (and the price point) expected of a true IIoT wireless vibration sensor.

Our 4100RF wireless system was one of the industry's earliest offerings and was award-winning technology. But it simply wasn't able to address most applications at the cost and capability intersection that would be necessary to displace most manually collected data scenarios. As a result, it enjoyed only limited, niche-focused success.

The size, capabilities, and price of the AMS Wireless Vibration Monitor allow it to be used in most applications previously addressed with handheld data collection. It is a triaxial vibration sensor with embedded temperature sensing as well as the ability to calculate speed via examination of the vibration signal. Like our AMS Asset Monitor, the AMS Wireless Vibration Monitor also features embedded narrow AI and is capable of delivering highly reliable prescriptive analytics to users. As such, it addresses two important needs in a CMaaS solution: automated data collection and automated data analysis / fault detection.

The implications of a true wireless IIoT vibration sensor like the AMS Wireless Vibration Monitor are revolutionary in terms of the economies of scale that it makes possible and how this transforms the CMaaS space. Indeed, one of the ways to assess the merits of a breakthrough idea is to look at the number of providers crowding the market after its introduction. Next to AI, few innovations have been a disruptive as pervasive, wireless IIoT sensors. A cursory search on the term "wireless IIoT sensors" generates a staggering 250,000 hits. Even narrowing this to "wireless vibration sensors" does little to reduce the stagger factor: it still results in an astonishing 80,000 hits.

Next to AI, few innovations have been as **disruptive** as pervasive, **wireless IIoT sensors**.

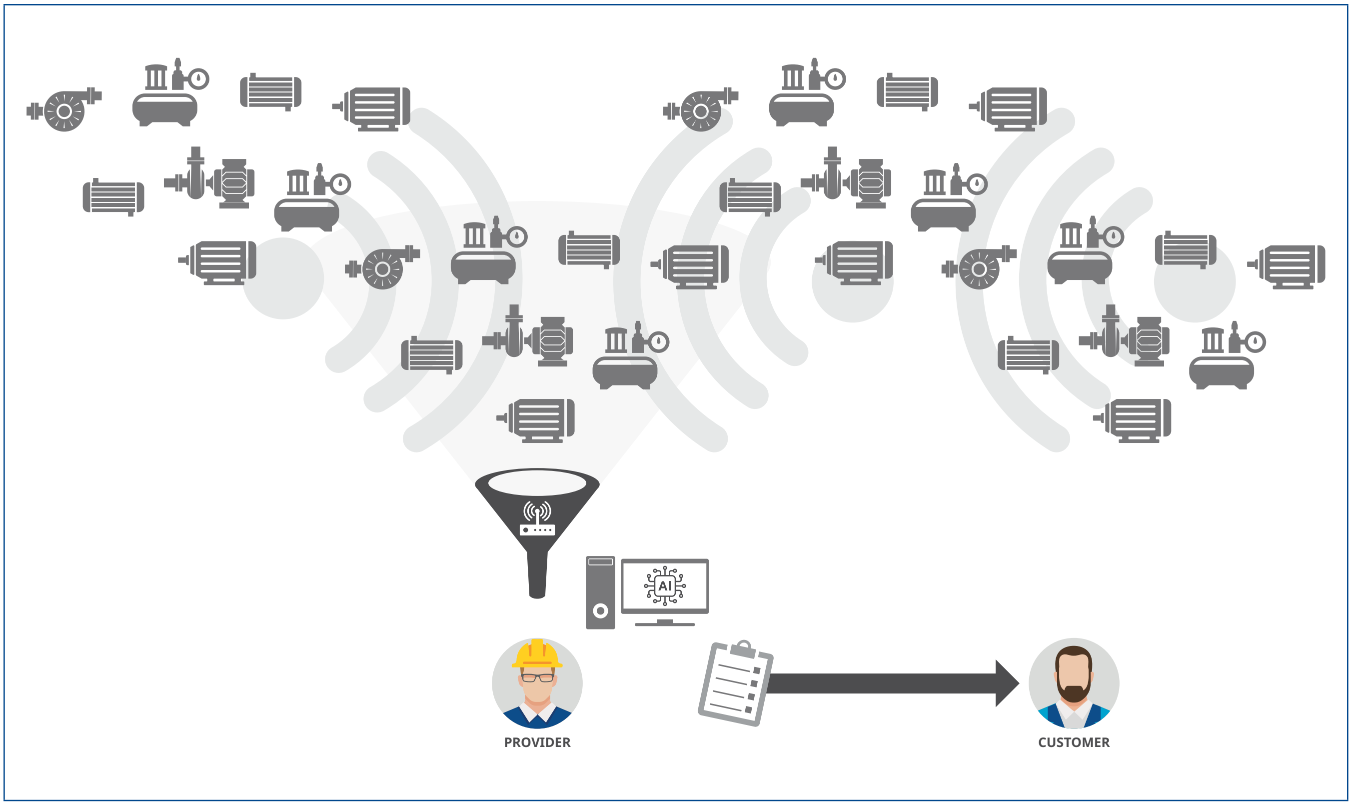

Wireless IIoT sensing coupled with AI has revolutionized the number of assets that can be addressed by a single person in a CMaaS scenario. Notice that the number of personnel at our provider has not changed, but the number of assets has increased dramatically. Notice also that the manual data collector has been removed as it is no longer necessary to visit the site to manually collect data.

Retooling the commercial model

In many ways, the commercial breakthroughs of the last decade have been just as significant as the technical. Essentially, the technical capabilities of AI and pervasive, wireless IIoT sensing made it possible for providers to automate a sizable percentage of data acquisition and data analysis. This, in turn, allowed them to entertain a very different commercial model than had been used historically.

The incumbent model has been described previously: the customer owned all of the infrastructure and was responsible to use it themselves. The outsourced model – as long as it consisted only of a person traveling to site to collect data – could be delivered as a subscription but was most often delivered in some type of annual (or longer) contractual arrangement. However, the infrastructure required was relatively minor. A provider needed little more than a computer, software, a route-based portable data collection instrument, some PPE[24], and a vehicle.

Changing from a long-term contractual model for CMaaS to a subscription-based model was just as responsible for an uptick in new users and industries as was the enabling technology of AI and wireless IIoT sensing. The costs of a CM program could now come out of an OPEX budget instead of a CAPEX budget and the ROI hurdle became dramatically lower.

[24] Personal Protective Equipment

As soon as we introduce wireless sensing to the mix, however, the provider must now deliver more hardware infrastructure to the customer. Considering a typical motor/pump set with 4 bearings, it would typically be monitored with 4 wireless sensors. If 200 assets are to be monitored, this translates to 800 sensors that must be deployed along with wireless gateways and then connection to the provider's cloud. The provider is now assuming significant cost to deploy this much hardware. Meanwhile, the customer is demanding that they pay only a monthly subscription fee.

The above scenario means that the provider is very heavily incentivized to deliver value each and every month. If they do not, and the customer cancels their subscription, the provider is left shouldering the costs of considerable infrastructure before it has been recouped. Essentially, this has caused the objectives of the provider and customer to be more closely aligned. To continue their subscription, the customer must see value. And to continue a revenue stream, the provider must remain laser focused on delivering value monthly.

For customers, the subscription model has lowered the bar in terms of price of entry. Whereas at one time they had to invest considerable capital – covering not just infrastructure, but installation, training (and then ongoing labor) – the subscription model removed all of this burden and amortized it over small monthly payments spread out over a lengthy time horizon. And perhaps most importantly, this could come from an operating expense (OPEX) budget instead of a capital expense (CAPEX) budget.

This had the effect of also lowering risk for the customer but raising risk for the provider. Previously, risk was shouldered primarily by customers because if they never recovered their CAPEX investment, the provider was not liable. Under a subscription model, if the provider was not delivering results, the customer could vote with their pocketbook and discontinue the service along with the accompanying costs. While one can think of the existing practice of outsourcing data collection and analysis as being a type of subscription model, it was rarely offered in increments less than 12 months. The new CMaaS model that commercially moves this to increments less than a year has had a profound effect on a customer's willingness to experiment with condition monitoring because not only is the price of entry small, but the risk is small as well. In turn, the provider is highly motivated to not just make things work for the first year, but indefinitely – allowing them to recoup their costs and begin to realize profit.

Overcoming skepticism

Numerous forums exist online for vibration practitioners to share experiences, ask questions, provide feedback, and just generally engage in shop talk. Such forums are valuable not only for customers but also providers because they get insight into how customers really feel about technology and where they are really experiencing successes and failures.

Even a casual tour of these forums makes one thing very clear: those with extensive vibration analysis experience are deeply skeptical of AI. While it is tempting to simply dismiss the skepticism as an auto-immune response borne primarily of job security concerns, it is important to look deeper and for providers to not simply dismiss these insights.

Skepticism is as old as the human condition itself. Refusing to acknowledge and address AI hesitancy is to ignore a very real issue. Providers have a very important role to play in making sure their technology is field-proven and hits the mark more often than it misses.

In a 2023 survey of professionals by another condition monitoring provider, participants were asked whether they trusted AI. Nearly half said they did not. And in a follow-on question, of those who had tried AI ***somewhere*** in their organization (not necessarily in their vibration programs), more than 25% felt the technology had over-promised and under-delivered. Simply put, AI suffers from a credibility gap, and providers that do not acknowledge this do so at their own peril. Customers must feel that AI will help them versus feel that it is pushed on them against their will.

Earlier, we mentioned that generative AI has had a surprising impact on the condition monitoring domain. That impact is that it has indirectly helped close the credibility gap. If AI can study Rembrandt's body of work and literally create a painting that would be all but indistinguishable from that of the original master (see page 8), the technology has essentially vindicated itself.

Presumably a vibration practitioner at a proficiency level of ISO 18436 CAT III or CAT IV, coupled with decades of actual experience is going to be able to render a diagnosis that is as good as (and perhaps better than) what AI can do. However, rather than pitting one against the other, it is more important to understand how they can work hand-in-hand.

First, remember that the practitioner has probably seen a lot of data over the years. And for every rule that practitioner forms in their mind, they also have exceptions and know when to invoke the rule versus invoking the exception. Second, asking a human being to pour through each and every dynamic waveform, trend, alarm, and anomaly for 10,000 machines a month is simply unrealistic. In order to be reasonably efficient, practitioners ***must*** find ways to address the forest without examining every square inch of bark and every vein in every leaf of every tree.

One of the most important roles of AI in helping a human analyst is thus sifting the so-called gold from the gravel. AI has shown that it is becoming better and better at reducing false positives and false negatives. And by this, we mean "is there a problem with this machine or not?" Getting that part of the equation right is more than half the battle.

One of the areas where Emerson was able to address this issue very early on was in a patent we obtained in 1997 for PeakVue™ technology[25]. Simply put, this technology allowed us to extract meaningful bearing and gear defect information from signals that would go completely unnoticed using conventional techniques[26]. PeakVue technology has proven itself so useful over the intervening three decades that it still comprises a significant role in the embedded AI of our products. And, as testament to all truly good ideas, it has been copied by many of our competitors. Indeed, it is almost unheard of that any legitimate competitor in the condition monitoring space is not using some variant of PeakVue signal processing in their own AI.

Another important aspect of our embedded AI is that of autocorrelation. This is a technique that allows the periodicity of a signal to be established. Vibration signals that are highly periodic in machines that have rolling element bearings are indicative of bearing wear, as the geometry of a spinning bearing causes recurring, periodic signals when the metal-on-metal contact encounters flaws, spalls, and other bearing defects. In contrast, high bearing vibration that does not contain periodic content is usually indicative of lubrication problems. Autocorrelation is a sophisticated and highly reliable technique for ascertaining and segregating bearing problems from lubrication problems. And, or course, there will also be situations in which both conditions are manifesting simultaneously – lubrication issues ***and*** bearing issues. The embedded AI utilizing these techniques must be able to handle all such situations and treat them with corresponding levels of urgency as well as prescriptive guidance.

In our own embedded AI, we rely on these approaches along with other techniques. We also allow users to augment our AI with their own learning so that the human expert is no longer pitted against the AI expert. They work hand-in-hand, and the human expert can embed their own knowledge.

---

[25] US patent 5,895,857.

[26] A simplistic way of understanding what PeakVue technology does is analogous to listening for the sound of a single person clapping in a stadium of 30,000 cheering fans. The sound is indeed there, but it can be exceedingly hard to extract amidst all the surrounding noise. PeakVue signal processing makes it possible to hear that single clap because – in our analogy – that clap contains vital information about bearing / gear condition.

Best practice today is to couple AI with a human gatekeeper that vets all AI findings and ensures not only that faults are legitimate, but that the prescriptive guidance offered reflects the sensibilities of a seasoned human expert.

When users outsource their programs, many providers (including Emerson) take things a step further: the AI is also vetted by a human before an advisory reaches the customer. Instead of pitting one against the other, both work in complementary fashion, but the human retains veto power and ensures that every notification of actionable information is personally inspected for veracity. Because AI is at work flagging anomalies, the human expert needs only look at exceptions and review the preliminary prescriptive guidance offered by the AI. This approach leverages AI, but does not replace the human expertise. Another way of looking at this is managing by exception. Specialists are not looking at unremarkable data to find remarkable data. The AI is doing that for them. The specialist is then reviewing remarkable data to ensure that the findings of the AI are trustworthy and actionable.

The IT factor

In the past, unless a CM program was entirely outsourced, IT was the customer's responsibility because the system had to run on their networks, their servers, their laptops, and their desktops.

It also meant that the sensing and data acquisition infrastructure had to be maintained by the customer's control and instrumentation personnel – who may or may not feel comfortable with such specialized apparatus.

Whereas some customers enjoyed the benefits of a dedicated OT department with corresponding expertise on industrial networks and industrial computing coupled with critical control networks, others expected their IT departments to handle everything. Still other customers – particularly smaller ones – did not even consider IT to be a necessary core competence and outsourced this function entirely. In such cases, unfortunately, the corresponding level of support and responsiveness received might be less than ideal. Incentives and metrics could be misaligned, not to mention the skillset to deal with specialized industrial automation equipment versus a routine Windows operating system issue or Outlook mailbox setting.

There was also the long-standing issue of customers that hosted the CM software on-premises having to install patches and upgrades, the complications of having the provider perform this remotely (if it was even allowed due to cybersecurity concerns), and the cost of having the provider come to site for such work when all else failed.

Finally, there was the cost of specialized, industrial software that could be prohibitive when a perpetual license was involved along with many users.

It was out of these frustrations that Software as a Service (SaaS) emerged and has in many cases become the preferred model today. The ability to interact with the condition monitoring software becomes little more than a browser connecting to a virtual machine somewhere in the cloud. Maintenance, patches, updates, server uptime, and everything else rests on the provider's shoulders – not the customer's. At the same time, customer IT departments are freed up to do what they do best.

For all of these reasons, the appeal of outsourcing the software – if not the entire CM program – is rapidly becoming the norm as industry embraces hosted solutions, and a driving factor in why CMaaS (not just SaaS) is appealing to many customers. Not only do they not have to gather and analyze data – they do not have to navigate IT issues on an ongoing basis as part of maintaining the computing infrastructure.

The challenges of working with an in-house IT department can be some of the strongest incentives for customers to embrace SaaS and CMaaS. IT will often find this approach appealing as well because they don't have to be responsible for supporting special infrastructure or special applications

AI and alarm management

Managing alarms can become burdensome when hundreds or thousands of assets are being monitored – each with perhaps dozens of parameters. There is always the issue of a false positive when an alarm is set too low, or a false negative when it is set too high.

There is also the issue of every asset being unique to some degree. Even two identical machines sitting side-by-side in a process, and with consecutive serial numbers, can differ in how they behave. While it is common to apply the same alarms to many similar machines in the beginning of a program, over time they don't tend to stay there. They get adjusted and re-adjusted. Eventually, that "just right" sweet spot is found. However, all of this takes time and a trained eye. Setting alarms initially and then adjusting them as required is thus another area where AI can be extremely powerful. For many years, this level of AI has been present in Emerson's condition monitoring products via a feature known as Autostat. It statistically analyzes the available data and establishes optimal alarm setpoints, in much the same way as a human thought process would work, but far faster and with the improved results that only individually tailored alarms can deliver in a program rather than one-size-fits-all alarms.

In a broader sense, AI is probably also doing alarm management in your own life in ways that may not be readily apparent. For example, credit card companies are continually monitoring card activity for fraud, but no two consumers are alike. Fraud monitoring has become extremely personalized thanks to AI, able to spot anomalies and alert you with text messages, phone calls, emails, etc.

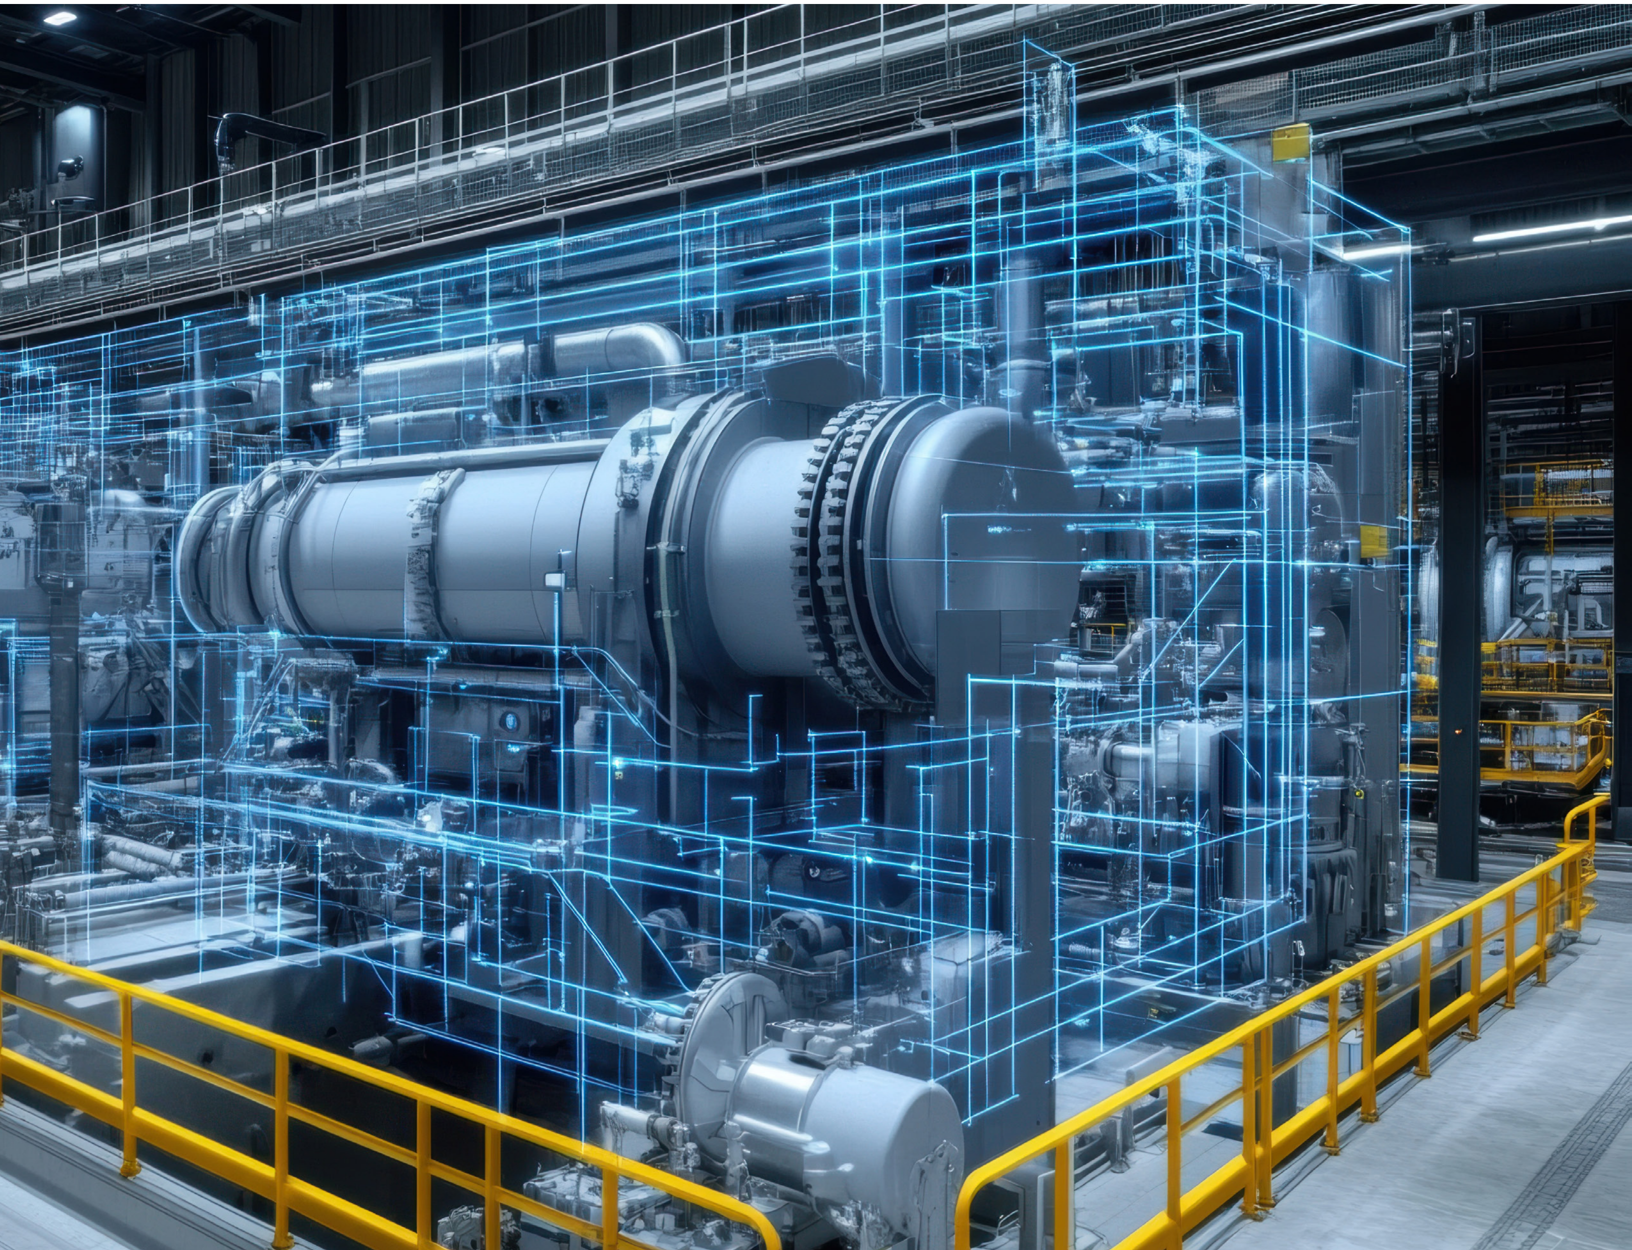

Commercial flexibility: one size rarely fits all

Today, customers are looking for an even greater degree of hybridization in their approach to condition monitoring. At one extreme is DIY where the provider supplies nothing more than technology. At the other extreme is CMaaS where the provider is responsible for everything except taking action on advisories. But in between these bookends are a number of scenarios such as the following:

- The customer prefers to own the instrumentation infrastructure but outsource the analysis.
- The customer prefer to outsource the infrastructure but in-source the analysis using the provider's AI but the customer's vibration analysis experts to vet the AI findings.
- The customer prefers to own only part of the infrastructure (the sensors and gateways) but have the provider host the software and AI. The human analysis in such a case could come from either the provider or the customers.

And still other scenarios can also be imagined. The key is that a provider must generally possess more flexibility than in the past because customers' business models and particulars can benefit from more tailored commercial solutions – not just tailored technical solutions.

Bringing AI to the bigger picture: Asset Performance Management

Ideally, even though AI is enjoying the spotlight right now, it flourishes in a supporting role – not a lead role. AI works best when quietly laboring in the background, doing tasks that were previously done by humans, allowing humans to do the things only they can do. In the context of managing mechanical and other assets, AI should thus be capable of looking for anomalies, relieving the practitioner from this burden. It will also be able to offer prescriptive guidance when it detects not just that ***something*** is wrong, but precisely what is wrong so – when possible – humans don't have to spend time answering that question. It will even be able to interact with humans using conversational language to assist with tasks like drafting work orders for a CMMS. These are just a few of the ways that AI can be part of an Asset Performance Management (APM) solution – a blending of condition monitoring, work execution management via a CMMS or ERP system, AI, and other tools such as case management.

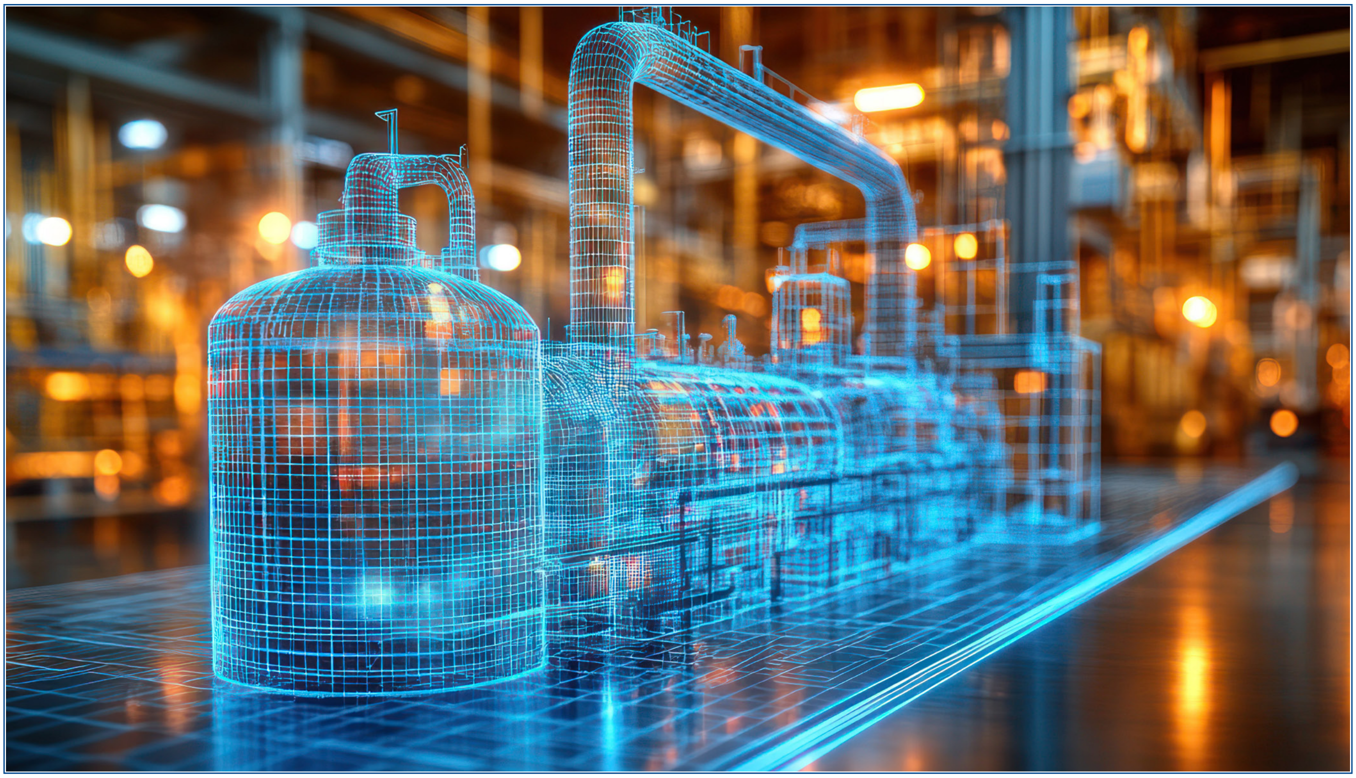

For Emerson, AMS Optics is the platform used to deliver APM. It blends a customer's own CMMS or ERP system with the AI in Aspen Mtell's anomaly detection agents, and with health monitoring for mechanical assets and instrument/automation assets via AMS Machine Works and AMS Device Manager respectively. It also leverages the embedded AI that is increasingly a part of Emerson's hardware both at the edge (with products such as AMS Asset Manager, AMS Wireless Vibration Monitor, and AMS 2140 machinery analyzer) and in the control room.

The easiest way to understand how these elements come together is by way of a scenario. It illustrates the role of AI and the power of integration to deliver a total asset health solution.

Here, we see a chill water pump along with numerous other assets inside of AMS Optics. This environment brings AI-related events – whether from hardware or software – and conventional threshold-based monitoring together. Notice that that AI is indicating a health score of only 65%, resulting in an alarm that triggers when health drops below 80%. Notice also the trend plot where 100% health is at the top and zero is at the bottom with scores that have fluctuated between 40 and 65 over the last 7 days. Interestingly, only the AI agent is detecting a problem with the pump. AMS Machine Works is returning a health score of 100 and, from its perspective, there are no concerns on the mechanical vibration side.

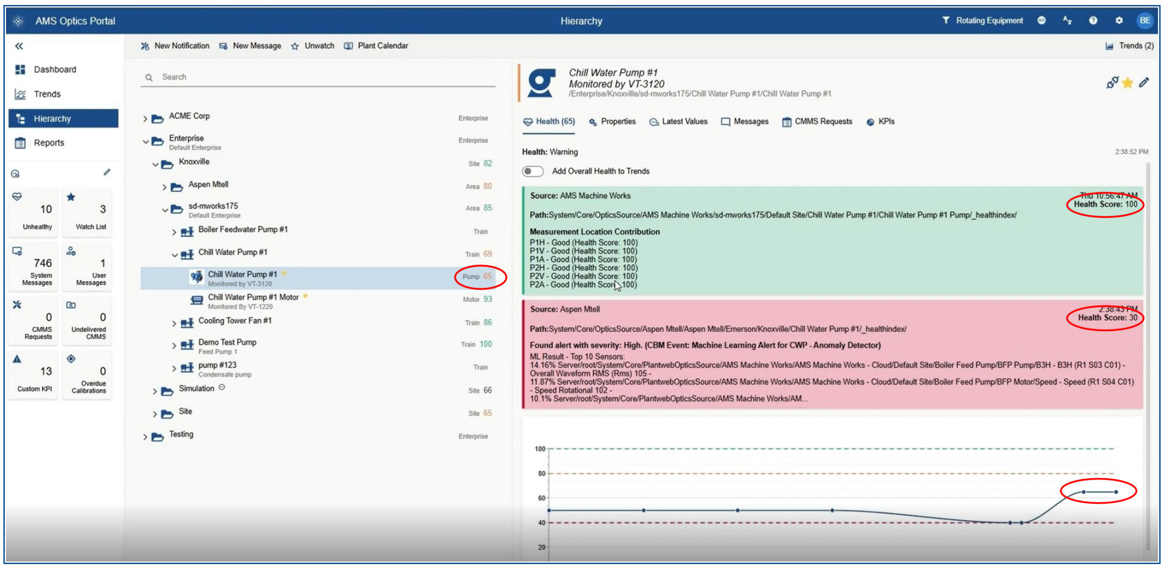

As a result, we decide to drill down into the AI agent and its anomaly detection for a closer look as shown below. The first thing we notice is that the probability of failure has been trending upward rather abruptly over the last several days after being relatively flat for the prior 20 days. We also see a list of sensors contributing to the AI findings (8 are displayed; more could be seen by scrolling down).

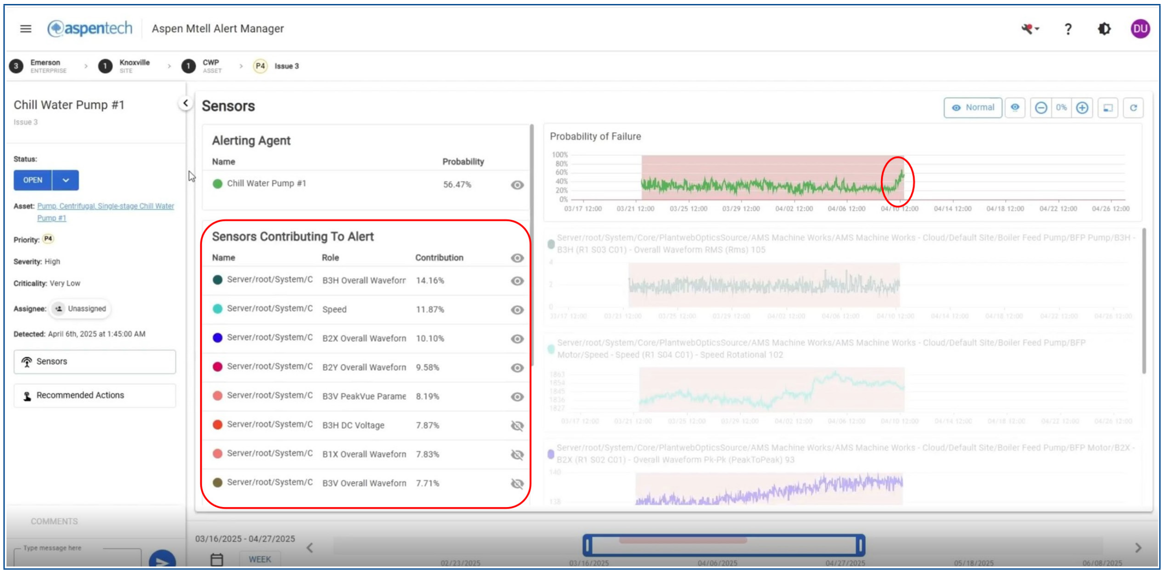

We drill further down as shown below to see what issues the AI agent believes are problematic and see that it appears to be one of two things: off-spec flow conditions (typically resulting in cavitation and dramatically accelerated wear) or perhaps lubrication. However, we remember that Machine Works is generally capable of detecting lubrication-related issues and segregating this from bearing-related issues, and has detected no problems with mechanical health, so begin even at this stage to favor a diagnosis of process conditions over lubrication conditions.

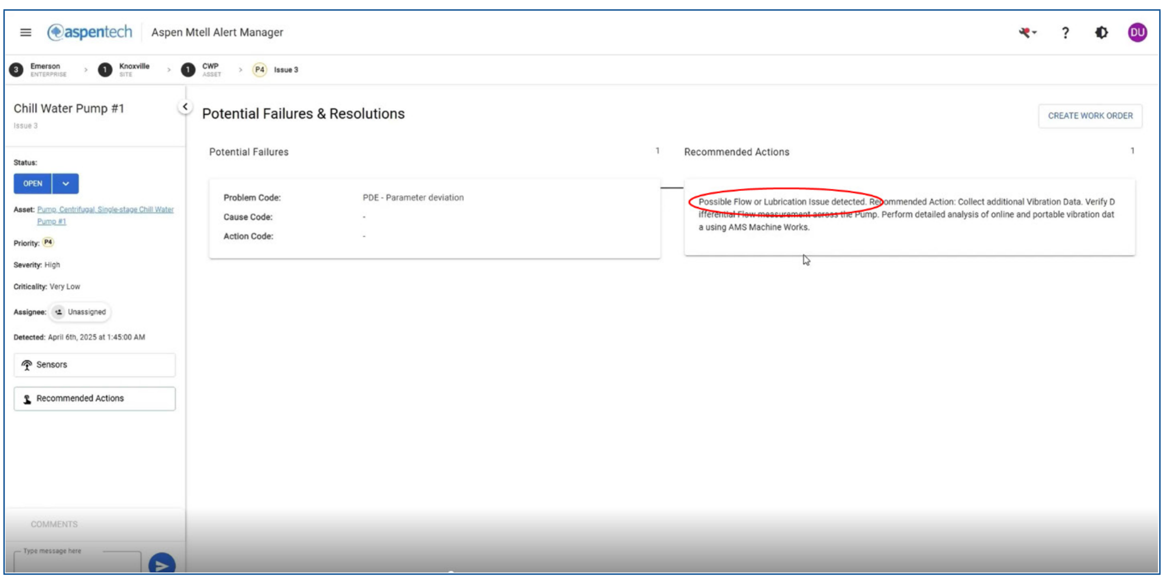

We then decide to look at the work order history for this asset using the integration of CMMS data and find that it was recently checked for proper lubrication and found to be OK as shown below.

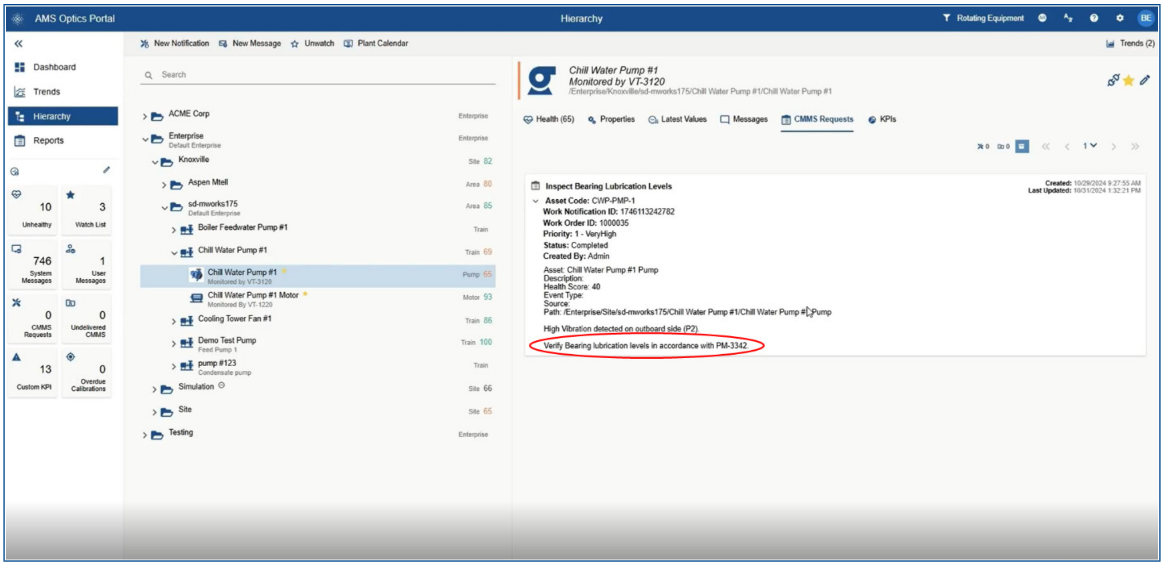

We then drill down as shown below to look at some of the sensors on the pump and satisfy ourselves that this is a legitimate problem with the pump – not the instrumentation – and find that all of our sensors have been operating at 100% health for the last 9 months based on data coming into our system from AMS Device Manager.

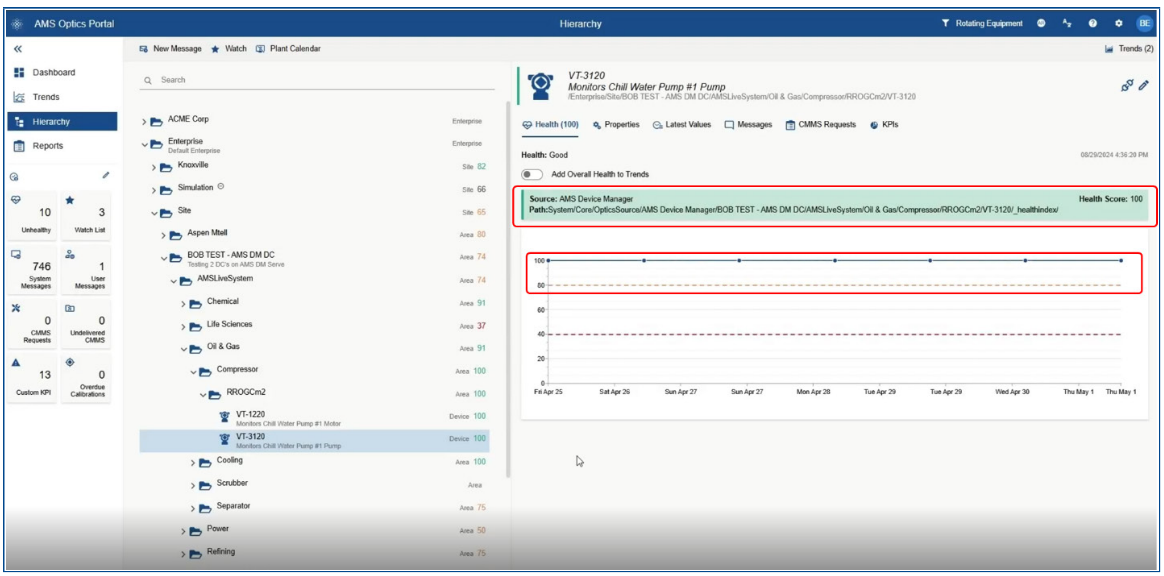

As before, all of this points to a process problem, so we decide to generate a work order to investigate operational issues that might be causing the pump to cavitate. To assist us in investigating this issue – which was flagged by AI – we turn to AI again to help us draft the work order, as shown below.

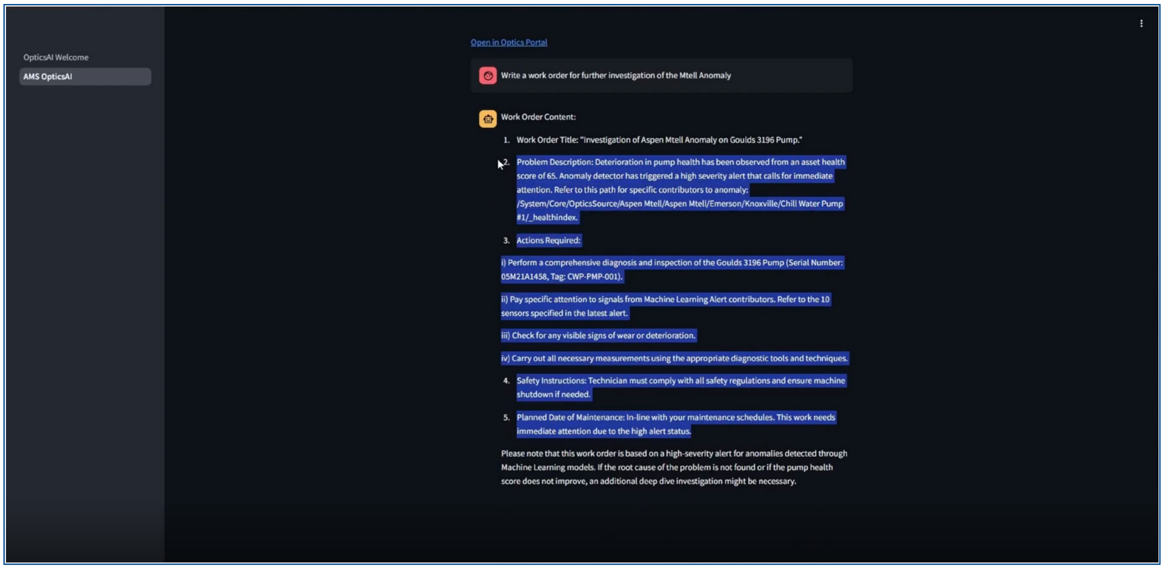

When we are satisfied with the language AI has drafted, we select the text, copy it, move back to our CMMS environment, and paste it to create a work order as shown below.

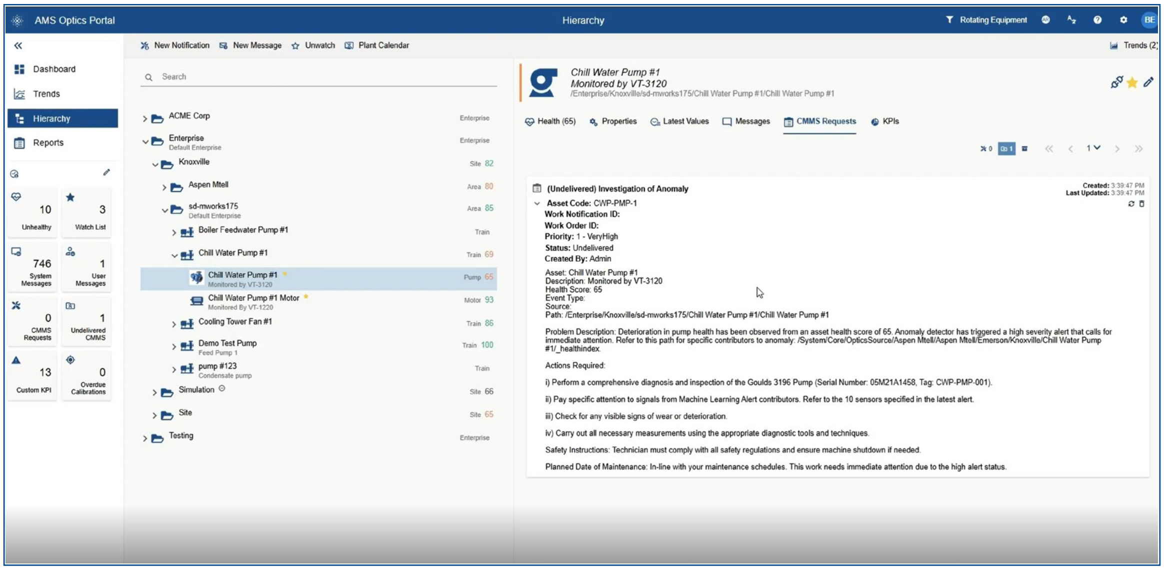

Although not shown here, the scenario concludes with the work order resulting in a finding of exactly what we expected: the pump was operating at off-spec flow conditions and was experiencing intermittent cavitation due to recent changes in some process setpoints. Although the high-frequency behavior associated with both cavitation and insufficient lubrication were possible, our hunch that a lubrication problem was unlikely proved correct – particularly since it had been recently checked a few months prior.

As part of the work order, the operations team was consulted, and they explained that the setpoint changes were made in the belief they would better optimize the process. However, they had failed to consider how the pump might be mechanically impacted and were unaware it was cavitating. It was further learned that the changes had been made in the last week, and this was entirely consistent with a health score that only started plummeting in the last 7 days. The decision was made to further refine the setpoints so that the needs of the process and the health of the asset could be balanced. Ultimately, a win/win was realized where the process was optimized without impacting asset health that would have resulted in premature and repeated failures.

Conclusion

In this whitepaper, we have provided an overview of the types of AI ranging from general to narrow to generative. We have also shown that while AI has been a part of condition monitoring offerings for several decades, it is often not heralded as such. Instead, it works quietly behind the scenes optimizing alarms, sifting through data, and flagging anomalies. Historically, it has relied so-called first principles where the physics of failures and malfunctions are well-understood and can be modeled along with provision of prescriptive guidance to users, but is more recently finding value in anomaly detection rather than reliance on conventional threshold-based alarms that can be very labor intensive to establish and maintain. We have shown that narrow AI is where the action is. It can go to work immediately, solving problems and delivering value – unlike the more general version of AI that is powerful in theory but must be patiently trained and channeled before it can do anything useful in the real world – often at the cost of millions of dollars and a contingent of experts to provide ongoing attention.

We have also shown that the emergence of generative AI has perhaps done more to bridge the credibility gap in AI promises versus AI reality than anything else of the last two decades. ChatGPT opened the floodgates and moved AI from an abstract concept that was peripheral to many people and industries, to a spotlight position in almost every major industry, as far ranging as healthcare, financial services, consumer entertainment, smart homes, and – of primary concern for our purposes here – condition monitoring for industrial assets. While it is still important for narrowly targeted AI to demonstrate to customers that it can deliver what it promises, the potential of AI itself is no longer dismissed nearly so easily when people are regularly carrying on conversations with AI and perhaps even using it to generate images for their social media posts.

We have shown that AI alone can account for only part of the transformation occurring in the asset management domain. Equally important has been the emergence of pervasive, "smart" wireless sensors and a retooling of the commercial model that delivers solutions as subscriptions rather than one-time, capital projects where all risk is assumed by the user rather than the provider.

Lastly, we have shown that although AI is in the spotlight in terms of the attention it receives in popular culture, in industry it works best when quietly laboring in the background, embedded in products and integrated in ways that allow it to augment rather than replace people. Indeed, when part of a larger Asset Performance Management solution, AI can detect not just mechanical anomalies, but instrument anomalies and process anomalies – an important capability because in the real world, the path to a confident answer regarding an asset's mechanical health involves knowing both the process conditions surrounding the asset and the health of the instrumentation measuring the process and the asset.

Contact Us
www.emerson.com/contactus