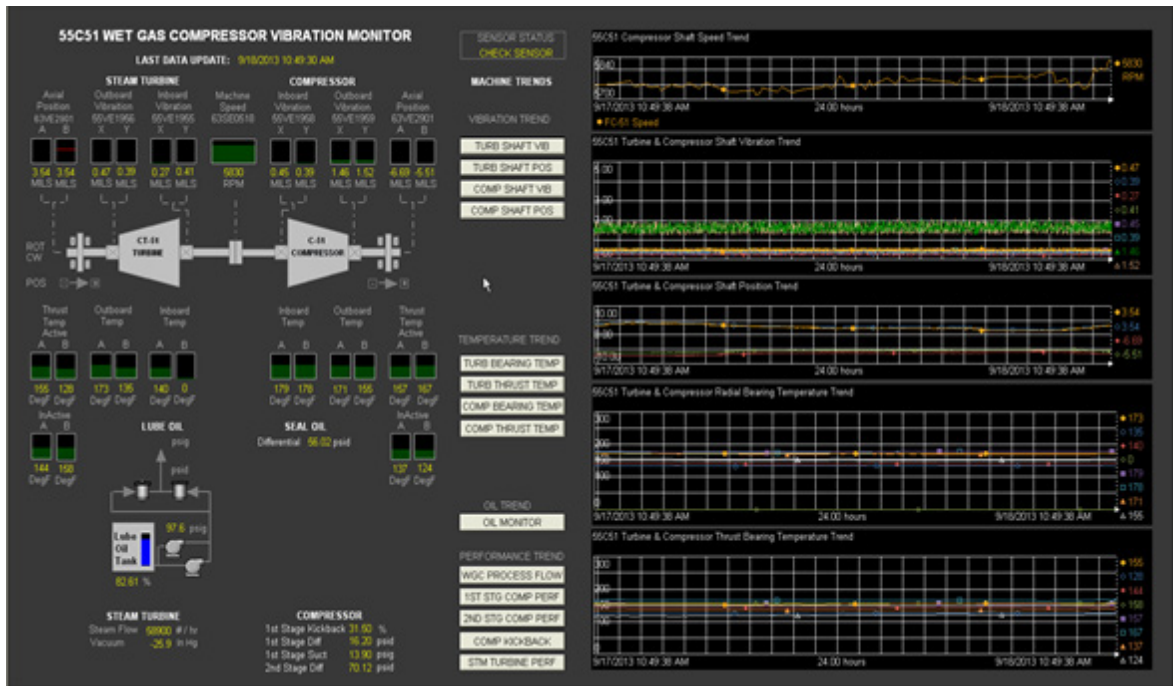




WHITEPAPER

Using a Commercial Process Historian for Full-Featured Machinery Condition Monitoring



Publish Date: 9/1/2015

Author: Steve Sabin

Can a commercial process historian be used to replace stand-alone condition monitoring software – including acquisition and display of high-speed vibration and surge waveform data? Historically, the answer has been “no.” Today, however, a very different story is unfolding.

A Brief History

Permanent vibration monitoring systems have been with us since at least the 1960s. When the first edition of American Petroleum Institute Standard 670 was published in 1976, it helped push such systems from the pioneering few into the mainstream. Today, it is standard engineering practice to include such systems on all critical turbomachinery – almost without exception – in not just the hydrocarbon processing industries, but all industries where critical machinery is found. These systems have expanded from simply vibration monitoring to now include bearing temperatures, overspeed, surge detection, and other parameters, and have thus become “machinery protection systems” instead of merely vibration monitoring systems.

The 1980s saw the rise of something new to complement these protection systems: computer software that archived and displayed the vibration data – including detailed waveform snapshots. Today, it is estimated that 25% of API 670 systems ship with some form of online condition monitoring software, making it the fastest growing segment of the machinery vibration measurement industry. Indeed, so prevalent have such systems become, the 5th edition of API 670 now includes an annex devoted specifically to condition monitoring software, augmenting the standard’s historical focus on only machinery protection systems.

The 1980s also saw the rise of another form of online software: the commercial process historian. The main innovation was the capturing of the reams of real-time data produced by process control systems and historizing this data in computer software instead of strip chart recorders. The data could then be easily and securely saved, shared and analyzed.

This concept exploded, and there are now tens of thousands of such systems around the world, collectively accounting for billions of process tags (FIG. 1).

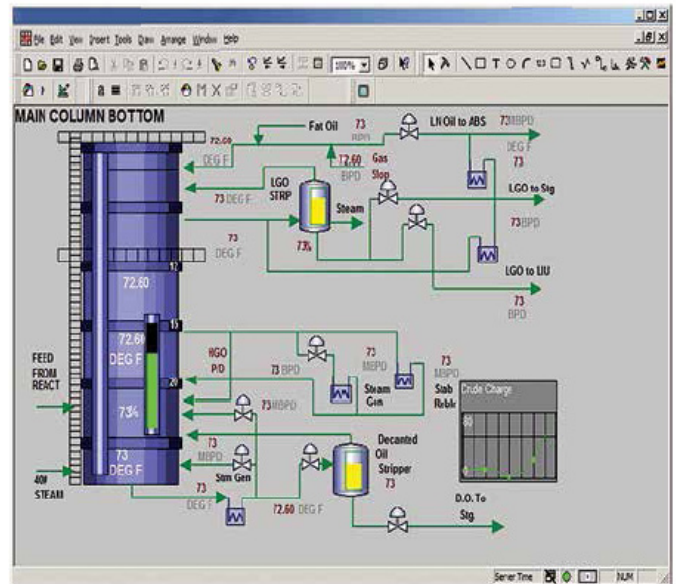


FIG. 1 – Typical Process Historian screen used to convey current values, statuses, and trends for process-related data.

Perhaps the best known of these is the PI System™ software, developed by Pat Kennedy at what was then Oil Systems Inc. (and has since become OSIsoft, LLC). Kennedy’s innovation? Taking the reams of real time data produced by process control systems – often “historized” on strip chart recorders – and Such systems are no longer thought of as simply “process historians.” They are truly real-time data infrastructures that handle far more than just process data. Process historians also handle events, spatial data, and asset hierarchies, to name just a few.

As process historians grew in market adoption and sophistication, so too did online condition monitoring software. However, the process historian concerned itself primarily with tens or even hundreds of thousands of points and associated update rates of hours, minutes, or seconds. Indeed, one-second process data archival was considered the “4-minute mile” of the industry. In contrast, condition monitoring software typically encompassed only a few hundred measurement points, but required data that was sampled much faster to capture much higher frequencies – roughly comparable to the audible spectrum (20 kHz) in terms of sampling speeds and

bandwidth requirements. As such, the two systems continued to exist as independent silos – one focused on hundreds of thousands of points with scan rates measured in seconds, and one focused on only a few hundred points with scan rates measured in milli or even microseconds. Where one system could convey almost all of its information in terms of trends and statuses, the other system required highly specialized plots used by vibration analysts to visualize waveform data in both time and frequency domains.

Thus, these systems evolved along two very different paths with two different sets of users. One targeted everyone in the enterprise for whom process and event data was valuable, while the other targeted the few people in the organization trained to interpret the specialized waveform data produced by its critical machinery.

The Twin Silos

Most plants have an overall control system architecture that is similar to Figure 2. Process data, as well as machinery data, flows into a Distributed Control System where operators can monitor and control the process; adjust machinery operating parameters such as speed, load, and flow, and monitor the status of subsystems such as anti-surge, vibration, and speed control. Generally, these subsystems communicate with the DCS via some type of open protocol, such as Modbus. Note the focus of data flowing into the DCS is real time control and monitoring of the process. Most DCS architectures are less capable (although they are continually improving) when it comes to historizing their data and providing a rich tool set for sharing and analyzing this data.

Thus, plants often rely on a separate process historian, such as the OSIsoft PI System, to provide these capabilities. This is particularly true when there are different types of DCSs within in a single plant or organization, and data must be shared across all of them. Historians are very adept at communicating with virtually any underlying system. For example, the PI System has more than 400 published interfaces

supporting a very wide variety of protocols. In contrast, the historians offered by DCS suppliers are typically designed to work only with their own control systems – not others. What is notable about Figure 1 is that virtually every type of data originating in the plant and its mechanical assets flows into the process historian – with two exceptions:

- 1) Machinery vibration data
- 2) Compressor surge / performance data

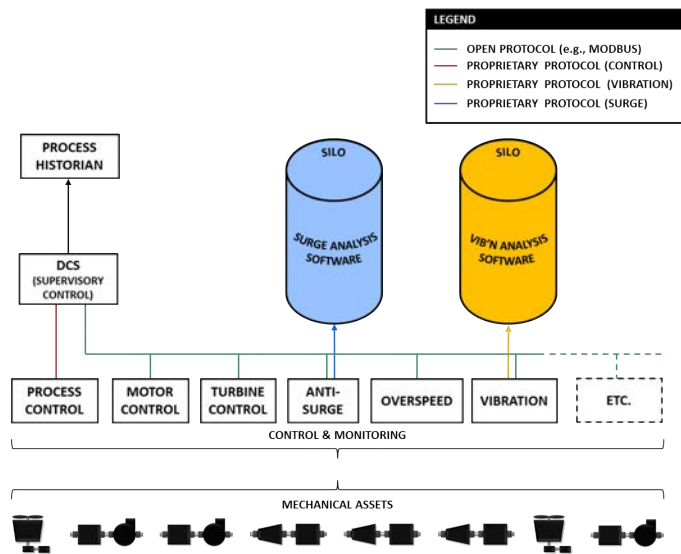


FIG. 2 – Block diagram for control and monitoring systems in a typical plant, along with communication infrastructures.

Both of these subsystems have instead relied on development of their own specialized software “silos” to collect and display the data of interest. Namely, high-speed vibration waveform data and surge data such as compressor maps and surge cycle analysis tools. Although the vibration monitoring and surge control industries have evolved their tools over the years to become increasingly sophisticated, one thing has not changed: they exist as stand-alone software ecosystems with their own proprietary infrastructures to collect, store, and display their specialized data. The vibration software is separate from the surge software, and both are separate from the process historian software.

Silos Perpetuated

At this point, some will no doubt remark that vibration and surge data has been sent to the DCS and subsequent process historian for many years. While this is true, it is important to note that the data supplied was simply current values and status conditions – the type of data that can be conveyed with 4-20mA loops and relay annunciation, and for which Modbus communications is perfectly adequate. What was not sent to the DCS and/or process historian was the high-speed vibration or surge waveform data from (for example) a dynamic pressure transducer or vibration probe where the amplitude versus time plot requires a time scale in milliseconds – not seconds.

The reasons given for perpetuating these silos typically include the following:

- They produce too much data – there is not enough bandwidth in the process control network.
- They require too much storage space because this data is updated in milliseconds instead of minutes or seconds.
- They require specialized visualization tools to display this data – they are not just trends, bargraphs, and status lights.
- The DCS / process historian is not fast enough for the sample rates required by this data.

It boils down to variations on the same basic theme: “Our data is special and therefore we need a special system – we simply cannot use the process historian to meet the needs of our rich dataset.”

For someone without a background steeped in the reasons above, such as process control or IT engineer freshly out of college, Figure 1 invariably prompts a predictable question – and frankly a very good one: “Don’t the process historian and those silos do essentially the same things? If they all collect, store, and visualize data, why do I need separate systems?” What this engineer is envisioning is a world without silos, depicted in Figure 3.

When this question is asked, however, is precisely the moment that the room gets quiet, the rotating machinery experts clear their throats, and recitation of the above bullet list commences. Hence, our new engineer quickly becomes conditioned to stop asking for such things and instead accept a world in which silos are inevitable.

The Problem with Silos

What isn’t fully conveyed in Figure 2 are the problems these silos create for customers. For a machinery engineer, it may very literally mean rolling your chair across the room to three different consoles just to see vibration data, process data, and compressor curves. It may also mean if an engineer wants to correlate process conditions with machinery conditions to ascertain cause and effect, they must resort to printing out hardcopies of trends and holding them up to the light in an attempt to overlay data from three different systems and compare them along a common time scale. This also means that it becomes necessary to learn three different systems to navigate and explore data. This skill set is in increasingly heavy demand and engineers probably no longer have the luxury of managing only the machinery in your facility, and are instead being asked to assist with machinery at other sites. Machinery engineers would like remote access or some way to easily share data, but many IT departments’ policies have made it impractical or perhaps impossible to do this under the existing architecture. While Figure 3 presents intriguing opportunities, experience tells engineers to curb their enthusiasm because something must inevitably be sacrificed in the process.

IT personnel will also fill the pinch. While they are acutely aware that although these silos look neat, tidy, and colorful on a block diagram, they represent many complexities: different operating systems to manage, security models to administrate, security vulnerabilities to understand and mitigate, software patches and upgrades to maintain, integration issues to overcome, annual licensing fees to keep current, data interface and integration issues to engineer, and

physical servers to maintain, among others. Users are probably hounding the IT department to provide them remote access to this data, but management is reluctant, or flat out refuses, to expose the enterprise to potentially malicious intruders. This can be especially frustrating to IT personnel because the organization has invested heavily in the process historian to make it a mission-critical, high-availability system that represents state-of-the-art security for remote access and real-time updating. It is preferable to just use that system to replace those silos, as it would make life so much simpler for everyone, not to mention less costly. However, machinery engineers have helped school IT personnel in the subtleties and special needs that their data represents, increasing concerns that simplifying things for the IT department will come at the expense of insufficient tools for the machinery department.

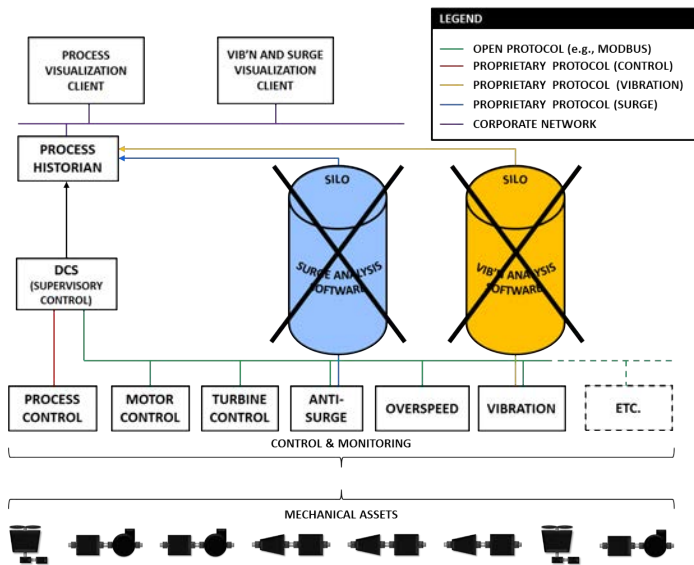


FIG.3. - Silos of FIG.2 replaced by the process historian and specialized display clients.

The reality today, however, is that the technical hurdles preventing Figure 3 from being realized simply no longer exist, and far from limiting the tools machinery engineers need, the elimination of silos actually enhances their breadth of tools.

What has changed?

What exactly is it that makes the same very same process historian that could not handle vibration or surge data ten years ago, able to handle it today? The simple answer is speed. Process historian throughput – and industrial computing power in general – continues to increase along a Moore’s Law (Figure 4) trajectory without any signs of abating. Essentially, a convergence of sorts has occurred. Earlier, we noted that process historians evolved along a path of increasingly higher tag counts – often numbering into the hundreds of thousands – but with each tag only updated once per second or so. In fact, it is not uncommon today to see a process historian updating a million tags per second or more. Conversely, the vibration world has relatively fewer points to measure – rarely more than a thousand – but at much higher update speeds of many thousands of times per second. However, if we think of both worlds in terms of events per second, we can see that 400 points updated 2500 times per second is essentially the same as 1 million points updated once per second. While there are, of course, nuances and details that make each scenario different, the process historian has evolved to the point where it can handle either of these situations with equal agility.

Today, the upper limit of process historian throughput is governed primarily by the speed at which data can be written to a mass storage device. In the age of spinning platter hard drives, there was a physical limit to how fast the disk could spin and data could be written. With the advent of solid-state drives, the speed has increased exponentially. These advances allow a single process historian today to accommodate 3-4 million events per second. When arranged into collectives, with each responsible for its own points, a network of process historians can update tens or hundreds of millions of events per second. Not only can these systems write data to mass storage almost unimaginably fast, they allow client applications to read from the database just as fast. Indeed, this is precisely why they are called real time data infrastructures.

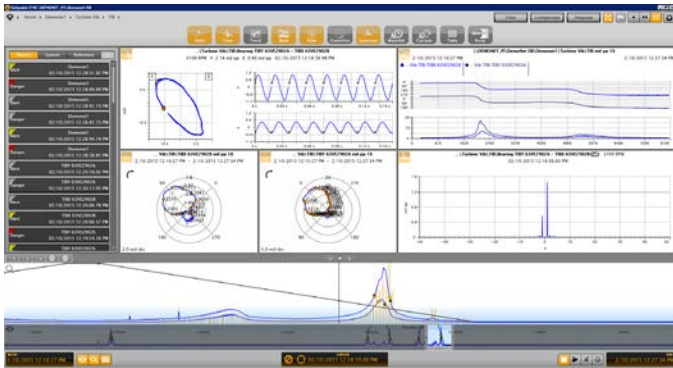


FIG. 6. – Screen capture showing how rich vibration data residing in a commercial process historian can be used to generate all of the plot types routinely used by rotating machinery engineers. Shown here are orbit/timebase plots, bode plots, polar plots, and full spectrum plots. Alarm event data is shown in tabular format in a pane on the left of the screen as well as superimposed on the time trend in the lower third of the screen. Data can be frozen, played back, or shown live.

The User Interface is Special – not the Data

While the underlying data is important, the tools used to display that data are more important. Commercial process historians do not have native capabilities to display such data in every conceivable format. This becomes the task of the domain expert, such as the vibration monitoring manufacturer. Figure 6 is the result of just such an effort, where the provider of the vibration monitoring hardware sampled the underlying data in their monitoring system, and streamed it into a process historian system. Vibration monitoring experts then developed a standard toolbox of display utilities that can connect to the underlying system, retrieve the data and present it to the user in meaningful formats.

Parallel efforts are underway to duplicate this for data collected by turbine control and anti-surge systems, using the process historian systems as the real time data infrastructure. Similar to vibration visualization tools, development of a standard toolbox of display utilities is underway to present compressor maps, surge cycle data, sequence-of-event lists, and other relevant data formats used by machinery engineers.

Conclusion

These efforts result in complete elimination of the silos discussed in this article. The user does not sacrifice functionality; indeed, they gain new levels of functionality when all of the necessary data – vibration, compressor performance, and process – resides in a single database where it can be easily retrieved, displayed, and correlated. Far from burdening the user with yet more software infrastructure, users can instead reduce infrastructure and use the process historian they often already own and in ways that have not previously been possible.

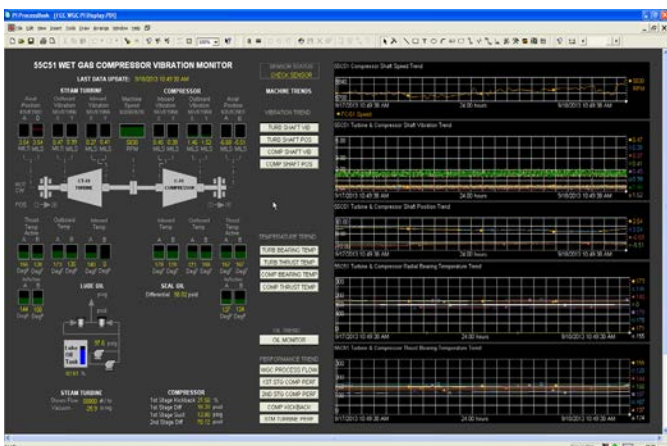


FIG. 7. Screen capture from a popular process historian display for a centrifugal compressor train showing pertinent trends, real time values, and statuses. From this screen the user can launch compressor map displays and vibration displays such as those in Fig 6.

About the Author:



Steve Sabin began his career in 1989 as a registered professional engineer in Canada, providing applications engineering support for vibration monitoring instrumentation. He was employed by Bently Nevada

Corporation / GE Energy between 1989 and 2010, holding a variety of sales, marketing, and product development roles. He also served as executive editor of ORBIT magazine from 1995 until 2010 and has authored more than 100 articles, white papers, and technical documents pertaining to vibration monitoring instrumentation. He is now director of product management and marketing with SETPOINT Vibration and a key member of the development team for its machinery protection and condition monitoring offerings. He has served as the secretary of American Petroleum Institute Standard 670 since 1996. Mr. Sabin received his Bachelor of Science degree in Electrical and Computer Engineer from Oregon State University in 1988, graduating summa cum laude. He is also a member of the Phi Kappa Phi and Eta Kappa Nu engineering honor societies.