White Paper

Cybersecurity Considerations for Vibration Monitoring Systems

Steve Sabin
Product Manager
SETPOINT™ Vibration
Minden, Nevada

Rev E (06/2016)

Contents

Overview

Industrial cybersecurity attacks are not new. A SCADA system breach in March, 2000, by a disgruntled ex-employee at an Australian water treatment plant allowed him to release 800,000 liters of raw sewage into parks and waterways[1]. In January, 2003, the Slammer worm infected First Energy's Davis-Besse Nuclear Power Station[2] and obscured critical monitoring parameters for nearly five hours that would indicate to operators – among other things – whether the reactor core was in a meltdown condition. Although this was not a malicious attack that specifically targeted the plant or reactor, it conveyed a sobering message about the consequences of unintended and uncontrollable network traffic in sensitive control system environments. And, of course, there is the well-publicized 2010 Stuxnet incident[3,4] where centrifuges at a nuclear enrichment facility in Iran were systematically destroyed over a period of months via infected Siemens S7 PLCs that controlled the machines.

These, and others, have resulted in heightened awareness in the industry, and there are now specific disciplines in many companies devoted to Industrial Control System (ICS) cybersecurity along with resources at the national government level in some countries[5]. Numerous texts[6-11] and symposium proceedings[12] now exist along with hundreds of papers[13] to comprise an increasingly large body of literature on this topic. However, none to date have dealt specifically with vibration monitoring systems due to the extremely niche focus of such instrumentation. This white paper thus examines the cybersecurity concerns pertaining to vibration systems – whether used in protective service, condition monitoring service, or both – and is broadly applicable to all industries in which ICS cybersecurity is a concern.

In general, the power generation sector is ahead of its counterparts in the petrochemical sector in one particular aspect: *cybersecurity*. As such, this white paper relies heavily on the best practices and lessons learned in the power generation world, cross-pollinating relevant considerations to the petrochemical world.

The power generation sector was actually the first to insist on more stringent architectures for sharing data and isolating threats, but similar emphasis is now coming from the petroleum sector given major infrastructure concerns such as pipelines, tank farms, and the like where disruption of supply can have sweeping and severe impact. In addition to maliciously caused outages, there is also the ever-present concern of explosions, fires, and toxic releases that could potentially occur as a result of intruders with malicious intent. Thus, while this white paper leans heavily on the lessons learned, practices encountered, and architectures required by leading power generation companies, the cybersecurity insights it provides are equally applicable to the petroleum sector.

Background

Traditionally, customers in the petrochemical sector have been the most sophisticated users of condition monitoring technologies. They were the first to adopt continuous machinery protection systems based on vibration measurements. They were also the first to begin using online systems not just for machinery protection, but also condition monitoring. The power generation sector lagged its petrochemical counterparts by 5-10 years in most cases. Historically, this had to do with the centralized and highly regulated nature of the power generation industry. It could often pass its costs onto rate payers, partially shielding it from the fully competitive mechanisms at play in more

deregulated environments like petrochemicals. Even in countries where the petroleum industry was nationalized, their raw and refined products could be exported across continents and oceans, and meant that pricing pressure – if not present domestically – existed internationally. It was this competitive environment that drove the oil & gas industry into condition-based maintenance practices earlier than their power generation counterparts, even though both had numerous similarities in their reliance on critical rotating machinery – often down to the same machine makes and models used.

Although the petrochemical industry started earlier and pursued advancements more quickly, the level of and reliance upon condition monitoring sophistication is today essentially equal across the two industries, given the deregulation of the power market in many countries. Recently, however, the power generation sector has moved ahead of the petrochemical sector in a one particular aspect of condition monitoring and machinery protection: *cybersecurity*. Why? Largely due to the criticality and inter-connectedness of the electricity grid. As important as petroleum products are, the loss of a refinery or production resource is isolated. In contrast, the electricity grid is interconnected by its nature and a failure can cascade, affecting large geographic regions in their entirety, as evidenced by 2003's blackout of the northeastern US[14]. For those affected, grid vulnerability was no longer an abstract concept and it took little imagination to extrapolate an outage of several days into one of much longer duration and the implications thereof. There was a collective recognition that whether an outage was the result of an accident or malicious intent, the world had changed[15] and grid security was no longer a footnote.

The Special Role of Machinery Protection Instrumentation

Machinery protection systems have emerged as an area of concern because they have direct impact on running machinery. The well-known Aurora Generator Test[16-19] conducted in 2007 at Idaho National Labs[20] specifically targeted machinery control, where hackers were able to compromise (bypass) a protective system for a diesel generator, and then open and close its electrical breakers out of phase with the grid frequency and thereby quickly destroy the asset. The vibration systems used to protect such machinery can be likewise compromised and used to either maliciously shut down equipment or to bypass normal protective functions. When both control and protection systems are attacked, a failure of the control system will no longer have an independent system that shuts down a machine in distress. Accidents at the Iranshahr[21] thermal power plant (Figure 1), Sayano–Shushenskaya[22] hydroelectric power station (Figure 2), and the aforementioned INL Aurora Generator Test powerfully demonstrate what happens when protective systems fail to act or are missing entirely. Of even more recent concern is the ability of condition monitoring software, under malicious control, to flood networks with high-bandwidth data and erroneous alarms, diverting operators or obscuring their visibility into critical conditions while the real target of the attack is busy being compromised elsewhere. This was one of the

Figure 1: The Iranshahr thermal power plant following a catastrophic wreck of one of its four identical 64MW steam turbine generator trains.

Figure 2: The Sayano–Shushenskaya hydroelectric power station before (left) and after (right) a catastrophic wreck of unit #2, one of its ten identical 640MW units. Root cause was traced to elevated vibration that fatigued the turbine's mountings and head cover. 75 fatalities occurred as a result of this accident and estimated losses totaled $523M USD.

implications of the Davis-Besse incident, referenced earlier in this document.

Turbine control systems have been identified as particularly vulnerable, not only because overspeed is often easy to induce and generally catastrophic to machinery, but because many controls were designed at a time when cybersecurity was not even a consideration. Instead, emphasis was being placed on computerization and the flexibility of Microsoft® Windows-based HMIs and configuration software for things like remote access and ease of field configuration changes. With the prevalence of independent shutdown systems, such as API 670 machinery protection systems, it is no longer enough to attack control systems. The protective system must also be disabled and thus assessment of its vulnerability is becoming increasingly common.

The T-Diagram

25 years ago, the concept of the T-Diagram (Figure 3) was introduced by Roger Harker, then president of Bently Nevada Corporation and the author's boss/mentor. He tasked the author, in part, with articulating this concept in slides, industry papers, and other customer communications such as Bently Nevada's ORBIT magazine. The T-Diagram is a highly useful, generalized means of discussing the four basic types of interfaces between a machinery protection system and the outside world:

1. *Operator Information* interface
2. *Protective* interface
3. *Machinery Engineer Information* interface
4. *Configuration* interface.

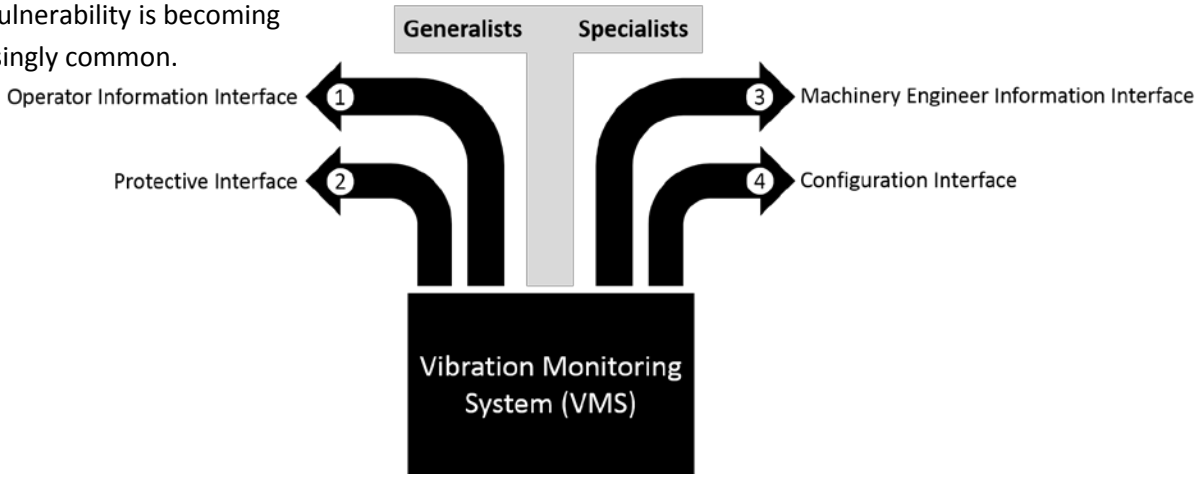

Figure 3: The T-Diagram showing the four separate and distinct interfaces. One side of the "T" provides signals and information intended for generalists without particular emphasis on machinery or instrumentation expertise – such as plant operators. The other side is intended for specialists with machinery and/or instrumentation expertise.

Before discussing the cybersecurity issues related to these, we first briefly describe the intent and historical context of each interface.

1. The Operator Information Interface – Description

This interface is intended to provide operators with basic information such as trends, current values, and statuses – the things originally shown on mimic panels with gauges, strip chart recorders, and annunciator lights – and more recently on DCS screens (Figure 4). Although this interface is highly important, and without it operators are "flying blind", it is not mission critical since machinery protection is maintained even if this link goes down.

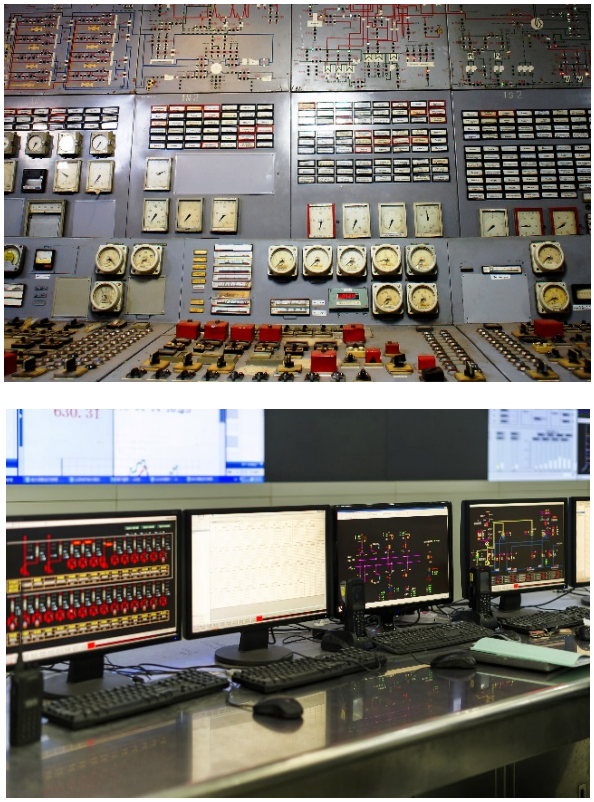

Figure 4: The Operator Information Interface is intended to drive operator displays such as older mimic panels (top) which relied on 4-20mA signals and discrete relay contacts from the VMS. Newer control rooms using a DCS will typically employ a digital connection, with a protocol such as Modbus, to populate computerized HMIs (bottom).

Up until the 1980s, this link was exclusively analog, consisting of 4-20mA outputs (one per channel) to drive meters and strip chart recorders, and relays to drive annunciator panels. In the late 1980s, however, an increasing number of plant automation systems began to introduce digital interface capabilities. Machinery protection systems subsequently appeared on the market with support for digital communications, most often using Modbus RTU. Such links could carry the same information as 4-20mA signals and relays, but contained a richer data set as multiple extracted parameters from each channel (such as sensor bias voltage, filtered amplitude, etc.) could be provided. Unlike analog connections, where individual twisted pairs are used and a problem is thus often isolated to only a single channel, digital communications carry all signals for an entire rack (or clusters of racks) on a single cable. This single-point vulnerability meant that, beginning in the mid-1990s, it became increasingly common to see not just digital interfaces specified, but redundant digital interfaces and redundant media.

These digital interfaces are characterized by four primary attributes, typically different than those of the other digital interfaces:

- data types limited to those consumed by operators (current values, trends, and alarm statuses),
- low bandwidth (due to the relatively simple dataset provided),
- use of polled query/response protocols (most often Modbus, and to a much lesser extent Allen-Bradley DF1),
- transmission over secure, dedicated links. These were initially serial RS-232, 422, or 485. Approximately a decade later, Ethernet communications would begin appearing using protocols such as Modbus TCP.

2. The Protective Interface - Description

This interface, as the name conveys, allows the vibration monitor to protect the machine, linking it to the final control element that actually trips the machine – typically via an ESD, motor controller, interposing relay, or other intermediary. This interface has historically been an analog relay

(Figure 5) due to its simplicity, reliability, and power handling capabilities. It remains so to this day, for reasons discussed later in this white paper.

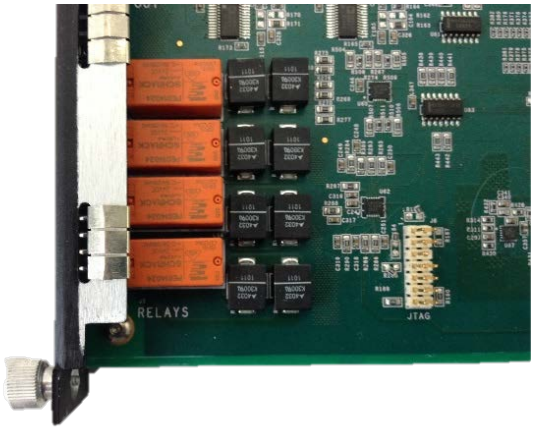

Figure 5: Typical vibration monitoring system relays (orange) used as the protective interface.

3. The Machinery Engineer Information Interface - Description

Machinery engineers need a dataset beyond that provided to operators; namely, high speed waveforms that allow the nuances of the vibration signal to be examined (Figure 6). This data allows the machinery engineer to discriminate problems that may all manifest as a change in vibration amplitude, but require aspects of the signal such as frequency, phase, shape, and form to be understood, much as a heart specialist might interpret the nuances of an electrocardiogram rather than being limited to basic parameters of temperature, respiration, and blood pressure. These interfaces are always digital because their terminus is software and have always been proprietary rather than using an open, published protocol. However, we are now seeing the introduction of analog

Figure 6: Typical screen from condition monitoring software using data provided via the Machinery Engineering Information Interface. It augments the information available via the Operator Information Interface by providing detailed waveform and other high-speed, high-bandwidth data needed for the specialized plot types used by machinery and vibration analysts.

"firewall" links (Figure 15) between the VMS and the condition monitoring data acquisition hardware. Driven initially by the nuclear generation sector, this approach will be discussed later in this paper as it is seeing increased use in non-nuclear power plants as well as other industries.

4. The Configuration Interface – Description

This interface needs little description, as its purpose is contained in its name. Figure 7 shows the software used with such an interface and the typical settings available to the user. Prior to the 1990s, these interfaces did not exist as systems were analog in nature and, for the few that were field-configurable, it was accomplished using circuit board jumpers. However, in the mid-1990s most vibration monitoring manufacturers began using digital signal processing technology rather than purely analog circuitry. Their systems were thus configured by means of software. Initially, serial interfaces (usually RS-232) were used. Today, a mix of serial (USB) and network configuration capabilities exist across the industry. The pros/cons of each with respect to security are discussed later in this paper.

Figures 8-10 show these four interfaces for several of the most commonly used monitoring systems in North America.

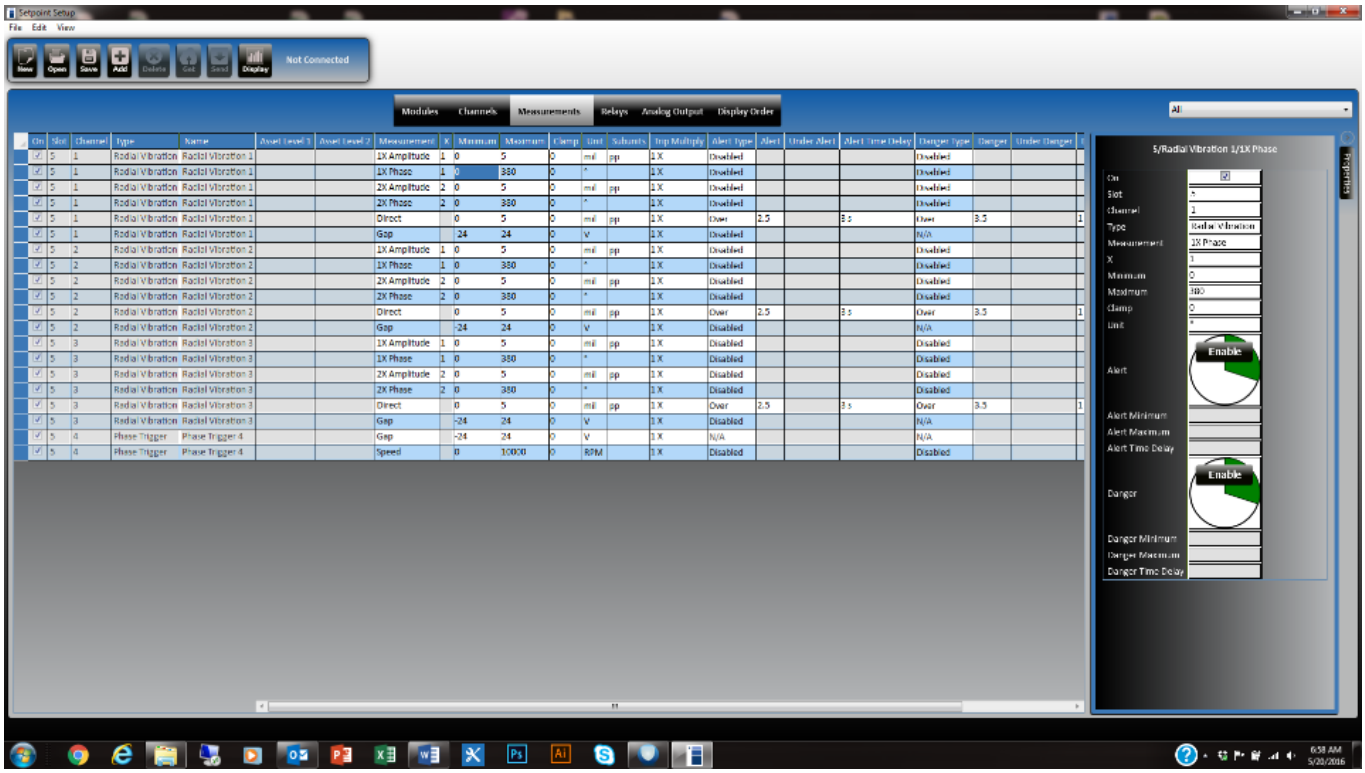

Figure 7: Typical screen from configuration utility. It allows every aspect of the protection system to be changed. As such, access to this port must be carefully controlled and represents one of the most potentially vulnerable points in the system's connection to the outside world.

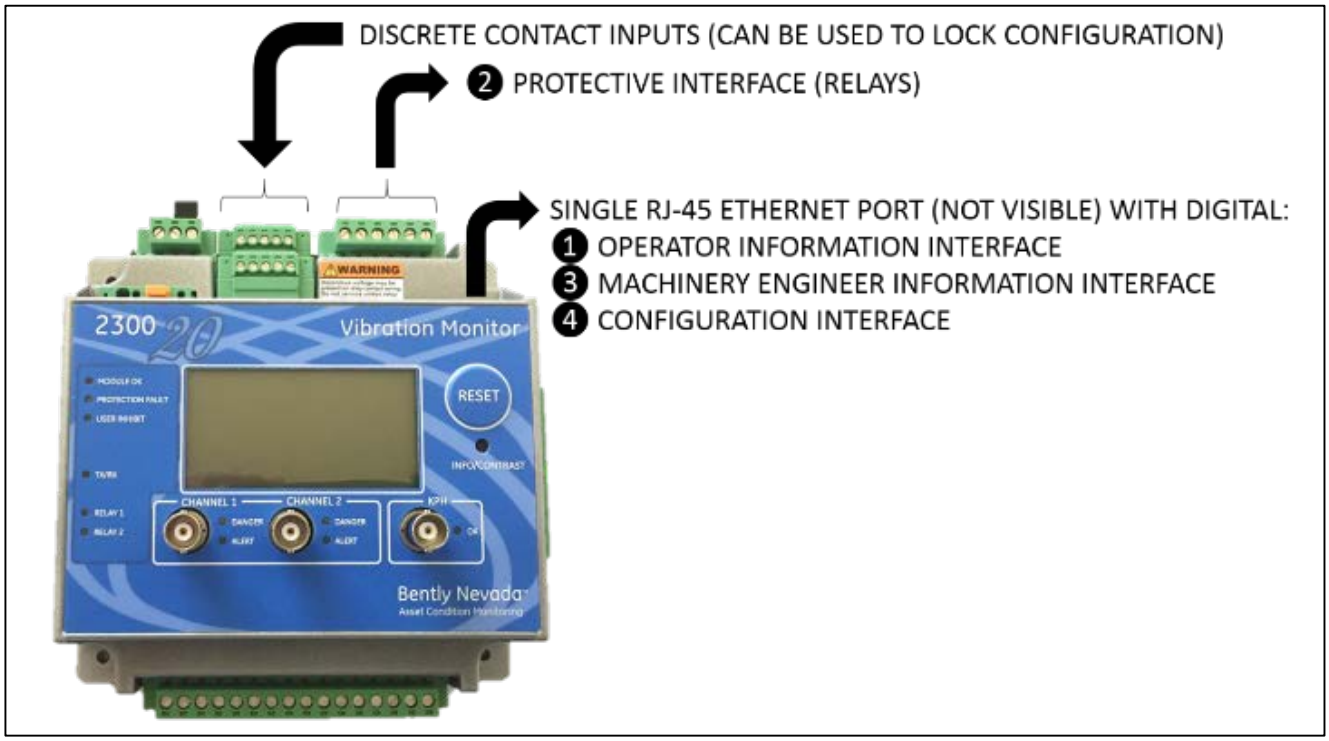

Figure 8: Bently Nevada® 2300/20 Monitor showing its dedicated protective interface (relays) and a single Ethernet port shared by its other three digital interfaces.

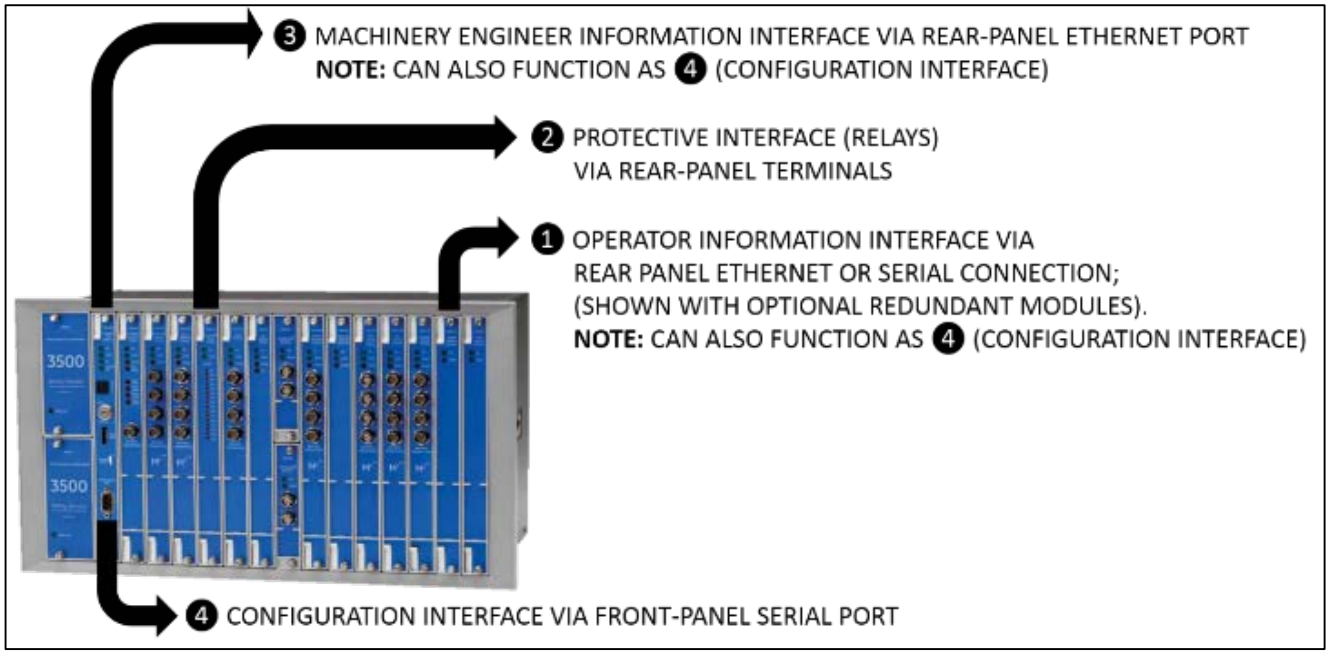

Figure 9: Bently Nevada® 3500 Series Monitoring System showing four individual interfaces; selected ports can be used for combined interface functionality.

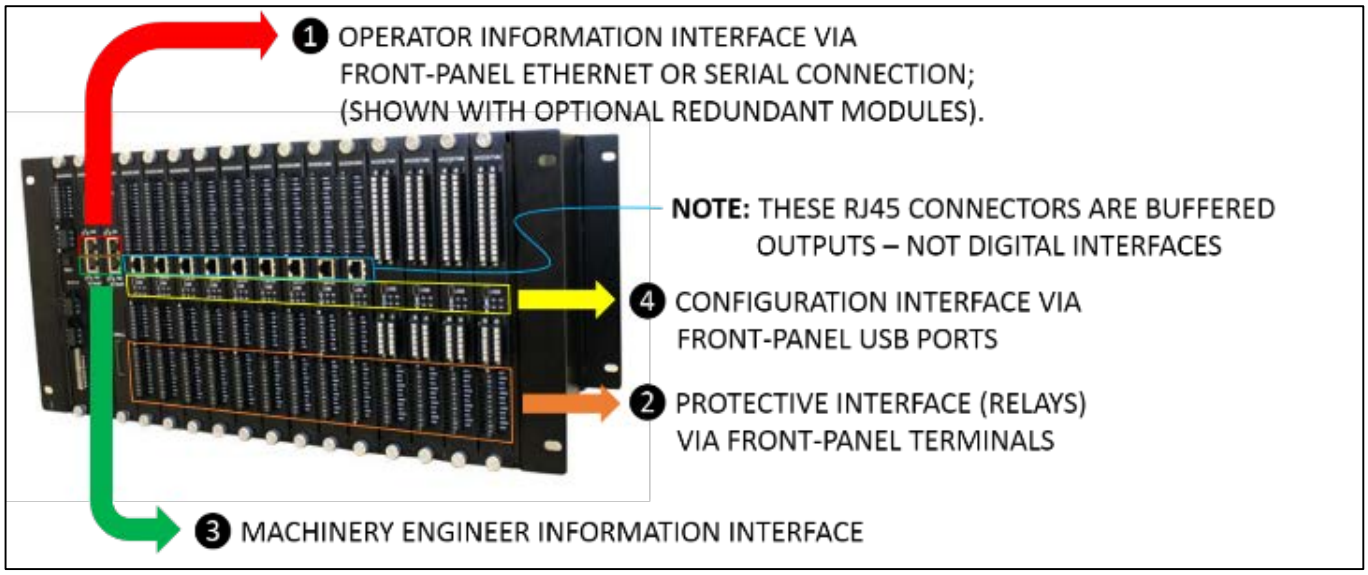

Figure 10: SETPOINT™ Machinery Protection System showing four individual interfaces; ports are segregated and cannot be used for combined interface functionality.

Cybersecurity Considerations

In light of the T-Diagram and its four types of interfaces, we now discuss in conceptual terms the cybersecurity considerations pertaining to each.

1. Operator Information Interface

Normally, this interface is designed such that an operator can only view the same items that would be visible on the local, integral HMI of the monitoring system (Figure 11): alarm and OK statuses, current values, and other front-panel type information. However, there are some notable exceptions to this in some systems, such as the ability to change alarm setpoints which will be discussed in a moment and represent obvious security risks if not properly managed.

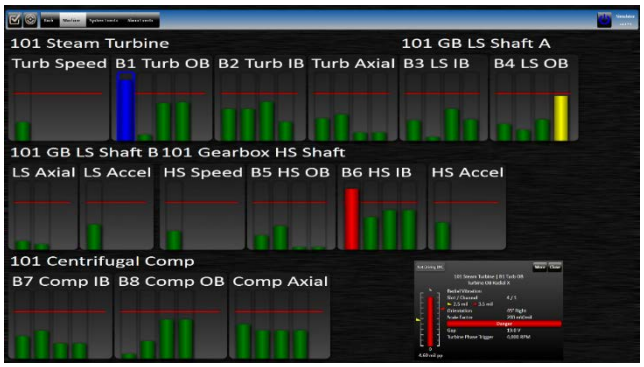

Figure 11: Typical local HMI for a machinery protection system showing bargraphs with current values, alarm statuses, and setpoints for a selected channel. Operators can view – but not change – system settings from this interface.

Operator interfaces can be roughly grouped into three categories: local, integral displays; local non-integral displays, and remote displays. Local displays, whether integral or non-integral, are designed to be located within a few feet of the rack - usually 10 feet or less, but occasionally more. The connection between the rack and such displays uses a mix of analog and digital signals. The video signal is often analog such as VGA or a higher-resolution variant. In instances where a digital video signal is supported, it is usually DVI or HDMI. Mouse and keyboard interaction with the display is usually via a serial connection such as USB. Connections to a local display usually do not represent a security risk because they are point-to-point, captive connections and not part of a network. Even in a worst-case scenario where the connection was compromised by a malicious party, the data passed between the display and the rack is not of a nature that would

allow machinery protection to be altered or defeated.

Of potentially more serious concern are remote displays. Older, analog connections used 4-20mA outputs and relays; as such, they represent no security risk. However, most newer connections are digital and use Modbus – often networked Modbus TCP. The key issues to examine here are three-fold:

1. Does the potential exist for this network to flood operators with spurious alarms and other data that would act as a diversion while a breach elsewhere was occurring?
2. Is a dedicated network used or is it shared with other information? If shared with other information, could excessive traffic from the vibration monitoring system (whether real or artificially and maliciously introduced) affect the other information, and is that information mission-critical (such as control)?
3. Does the remote display allow the operator to issue commands that will alter or defeat machinery protection?

By far the most common remote display types are DCS consoles using Modbus protocol (Figure 4 – right). Historically, this was serial Modbus (ASCII or RTU), but the use of Modbus TCP is becoming more common and means that other traffic may be resident along with vibration display information. In most cases, it is possible to use Modbus to alter machinery protection in the following ways:

- **To permanently change alarm setpoints.** For example, the Bently Nevada 3500 system allows setpoints to be changed via Modbus.
- **To temporarily change alarm setpoints.** This is a feature known as "trip multiply" and is used to temporarily elevate alarms while a machine passes through speeds in which elevated vibration occurs, such as when a machine is starting up and passing through a mechanical resonance. Most monitoring systems allow trip multiply to be invoked digitally (i.e., via Modbus) in addition to analog discrete contact closure inputs to the rack.

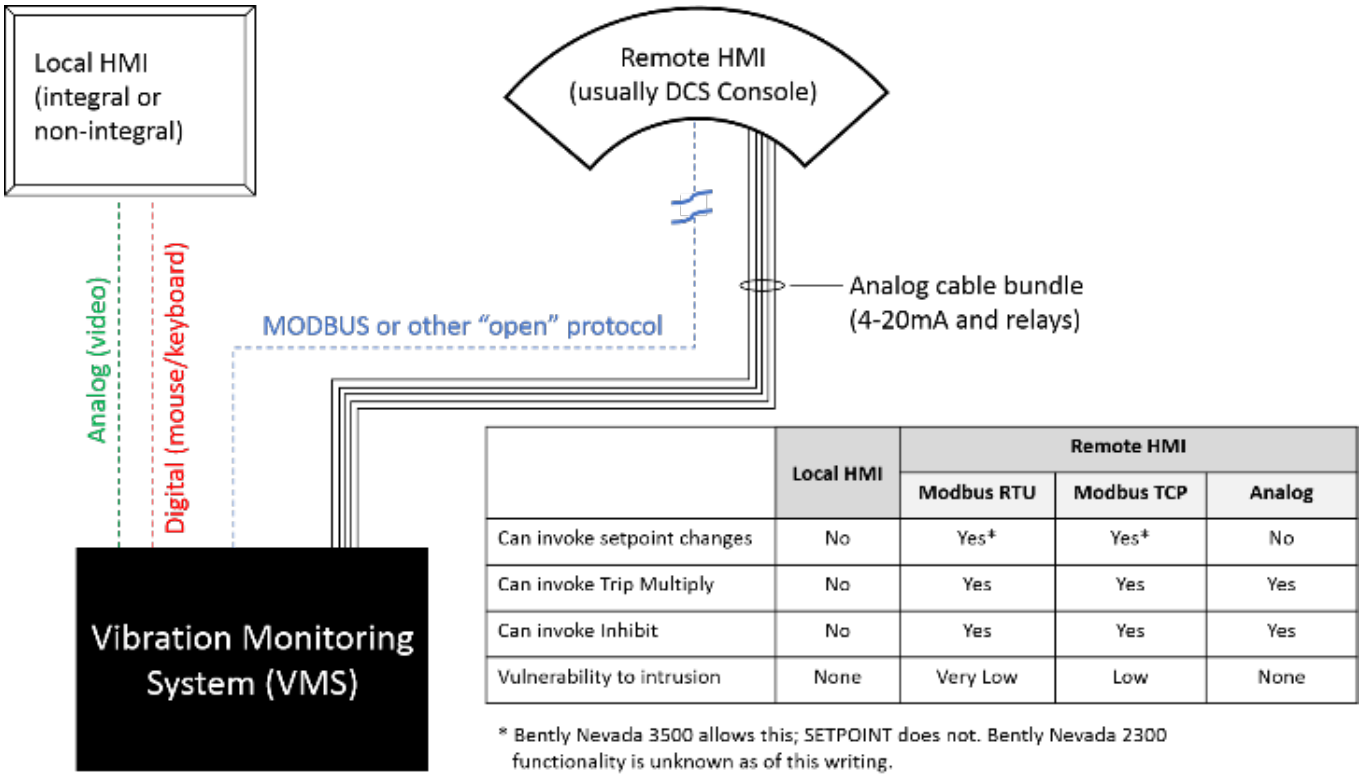

	Local HMI	Remote HMI		
		Modbus RTU	Modbus TCP	Analog
Can invoke setpoint changes	No	Yes*	Yes*	No
Can invoke Trip Multiply	No	Yes	Yes	Yes
Can invoke Inhibit	No	Yes	Yes	Yes
Vulnerability to intrusion	None	Very Low	Low	None

* Bently Nevada 3500 allows this; SETPOINT does not. Bently Nevada 2300 functionality is unknown as of this writing.

Figure 12: Typical connections for local and remote HMIs and salient capabilities of each.

- **To inhibit the rack.** This is a feature that suppresses alarming altogether. It can be invoked via discrete analog contact inputs to the rack, but most systems also allow this to be invoked digitally, via Modbus or other protocols.

When serial Modbus is used, the security risks are quite low. However, when Ethernet communications are used, vulnerabilities can exist and must be managed given the ability to suppress machinery protection via this link.

Figure 12 and its embedded table summarize the information covered above.

2. Protective Interface

This interface is usually the most secure of all because it nearly always uses entirely analog signals, such as relay contact closures or occasionally 4-20mA signals if the proportional signals are brought into the DCS and alarm setpoints established there. Point-to-point wiring is used rather than digital networks and as such these links are inherently "un-hackable" except for direct, physical sabotage such as cutting wires or removing terminal blocks to defeat machinery protection. Currently, there are no known protection systems that use digital signals superimposed on 4-20mA such as HART[23].

Systems that use analog protective interfaces include, but are not limited to, SETPOINT, Bently Nevada 3300 series, 3500 series, 2300 series, 1900/65, and ENCORE; SKF M800A; Emerson 6500 and 6500 ATG; and Vibro-Meter VM600, to name some of the more commonly encountered models.

There are also systems, notably Bently Nevada's 2201, 1701 and 3701, SKF's DMx, Vibro-Meter's VibroSmart®, and Rockwell's XM, that act as intelligent I/O blocks and communicate directly with a PLC or other control system via completely digital communications. They typically use protocols such as DeviceNet™, ControlNet™, Profibus™, and others that would typically be considered Level 1 in the ANSI/ISA-95 model[24]. These networks are usually used for device-level I/O and are well insulated from attack. However, because they are digital, they are not – at least in theory – completely immune from cyber attack. API 670 continues to preclude digital protocols for the protective interface and instead specifies analog (i.e., electromechanical relay) connections due to their speed, simplicity, and reliability. Thus, systems that rely on digital communications for protection, such as those noted above, are not strictly API 670 compliant, nor are they completely immune from cybersecurity concerns.

3. Machinery Engineer Information Interface

This interface has come under increasing scrutiny during the last five years for several reasons.

First, this interface has always been at least partially digital. Up until about 2000, a separate data acquisition unit was used to gather signals from the machinery protection system, digitize them, and pass them to condition monitoring software used by rotating machinery engineers. As mentioned, part of this signal flow was digital and part of this was analog (consisting of the raw transducer signals, digitized in the data acquisition unit). This is depicted in Figure 13. The Bently Nevada 3500 with external Transient Data Interface (TDXnet) is an example of this type of architecture.

Figure 13: Architecture used for condition monitoring software up until approximately 2000. After 2000, the architecture of Figure 14 became more prevalent. The digital interface (red) represents a potential vulnerability that was unacceptable to the nuclear power industry.

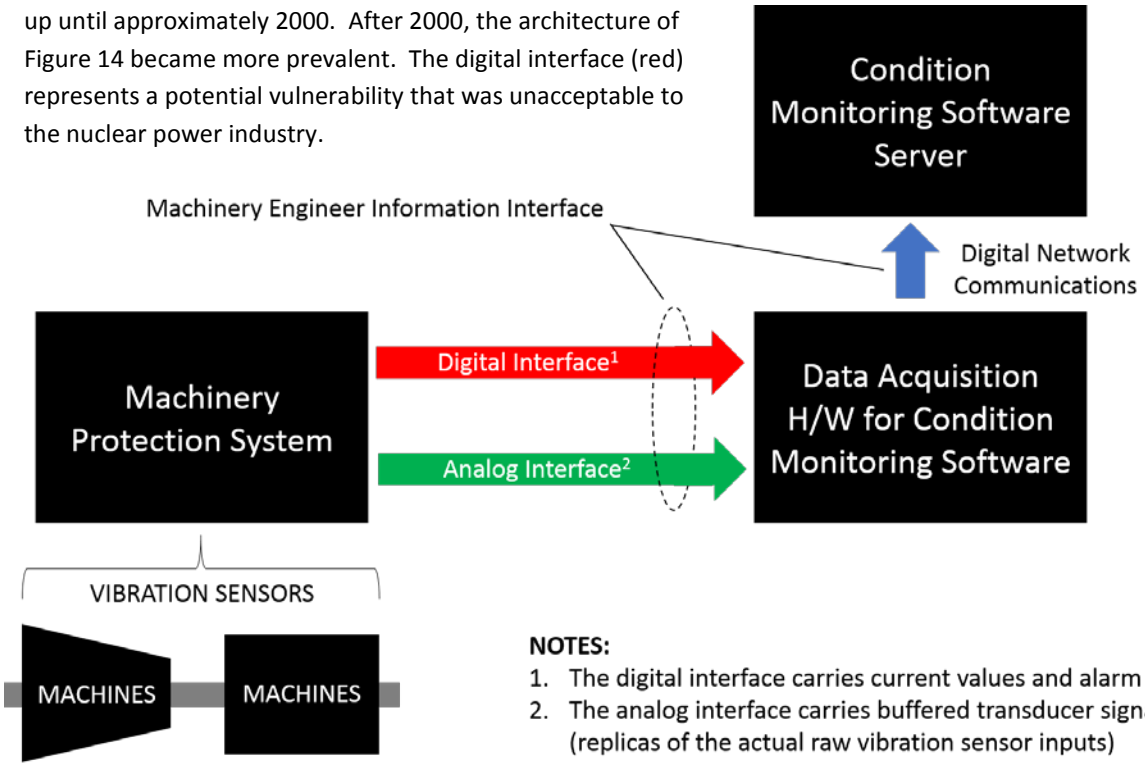

Second, during the latter half of the 1990s, the industry began moving away from an entirely separate data acquisition "box" and instead integrating this functionality into the protection system (Figure 14). It saved cabinet space because a single instrument rack could be used for both functions, and it allowed a single set of circuitry to do all of the computations and signal processing required for machinery protective functions and for condition monitoring functions. This integration reduced some of the small discrepancies that could exist when two separate systems were used – one for protection and one for condition monitoring. The industry continued along this integrated trajectory until approximately 2010 when concerns began to originate – primarily in the nuclear power industry – that this machinery engineer interface was too tightly coupled to the machinery protective system and represented a cybersecurity vulnerability. The Bently Nevada 2300 series and the 3500 series with Transient Data Interface (TDI) are both examples of systems where

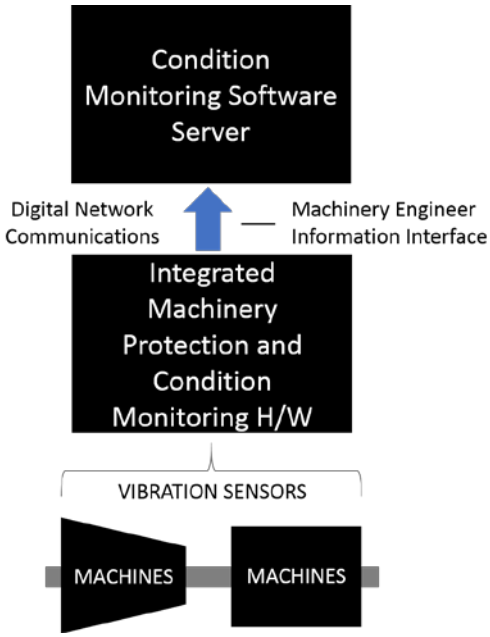

Figure 14: This architecture for condition monitoring software became popular after its introduction around 2000-2001. The digital interface (blue) represents a potential vulnerability because the protection system and the condition monitoring system share the same hardware and are accessible via this single link. The nuclear power industry found both this architecture and that depicted in Figure 13 unacceptably vulnerable.

machinery protective and condition monitoring functionality are shared within the same unit.

Bently Nevada responded to this vulnerability by introducing an external data acquisition unit called *TDI Secure* that completely insulated itself from the protective system by using only analog interfaces. The approach is depicted in Figure 15 and effectively creates an analog firewall between the protective and condition monitoring systems. Buffered transducer signals (analog) are hardwired from the protection system to the condition monitoring hardware. Relays and analog 4-20mA signals are likewise used to bring protection system statuses and current values over to the condition monitoring hardware, rather than digital signals as had been used in prior generations of separate condition monitoring hardware. Although secure, however, the hardware used by the condition monitoring system and protective systems were entirely different and increased both the cabinet space required and the spare parts burden because they had no commonality.

The SETPOINT system can be implemented as shown in the depictions of Figures 14, 15, or 16. It serves as a "gateway" to stream data into the OSIsoft PI System, which is used as the real time data infrastructure for SETPOINT's condition monitoring offering.

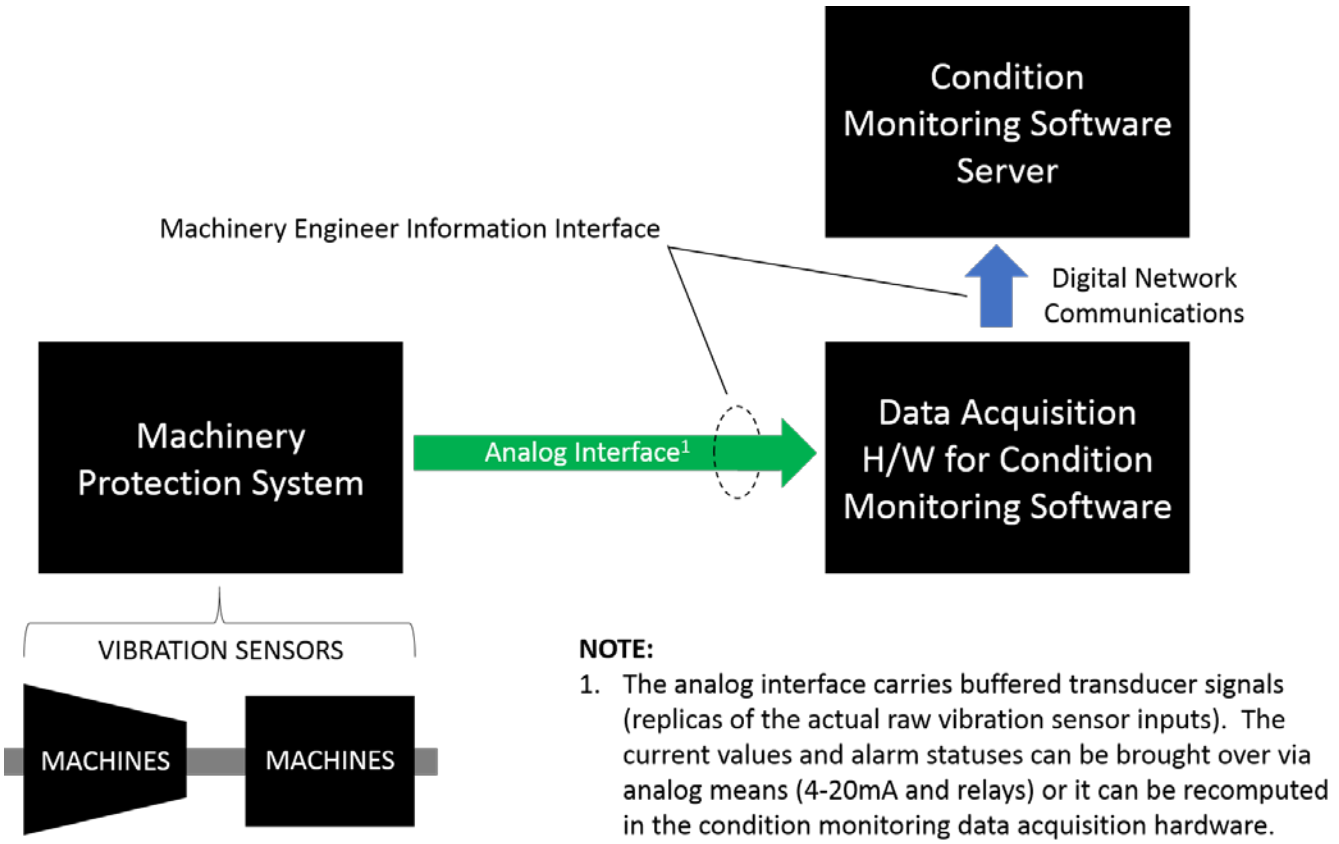

Figure 15: The nuclear power and other security-conscious industries insisted on a return to the separate and distinct hardware for machinery protection and condition monitoring, but without any digital links between them. The result is shown above and is becoming increasingly popular. Bently Nevada TDISecure and SETPOINT are examples of hardware that enable this architecture. The machinery protection system is thus decoupled from the condition monitoring hardware and can be supplied by the same or different vendors.

In Figure 16, Rack #1 serves in a machinery protective role. Rack #2 provides the condition monitoring interface, and both share the same analog sensor input signals. Use of identical hardware for both systems has the advantage of allowing the same spare parts to be used for both functions, while providing the necessary physical segregation.

In industries where cybersecurity is a concern, this segregated "analog firewall" architecture is becoming more prevalent. Although it essentially duplicates the required hardware and cabinet space requirements, it is considered an acceptable trade off due to its inherent ability to insulate the machinery protection system from vulnerabilities via the machinery engineer information interface.

Not all customers, however, have adopted the approaches of Figures 15 or 16. They have continued to use the integrated architecture of Figure 14 instead. In such instances, there are certain considerations for an integrated architecture that can make some approaches more vulnerable than others. These are examined next.

When an integrated condition monitoring and machinery protection architecture is employed, it is desirable to allow machinery engineers to reconfigure the data acquisition and signal processing parameters remotely. It is also desirable to share configuration parameters between the protective system circuitry/processors and condition monitoring components so that identical settings do

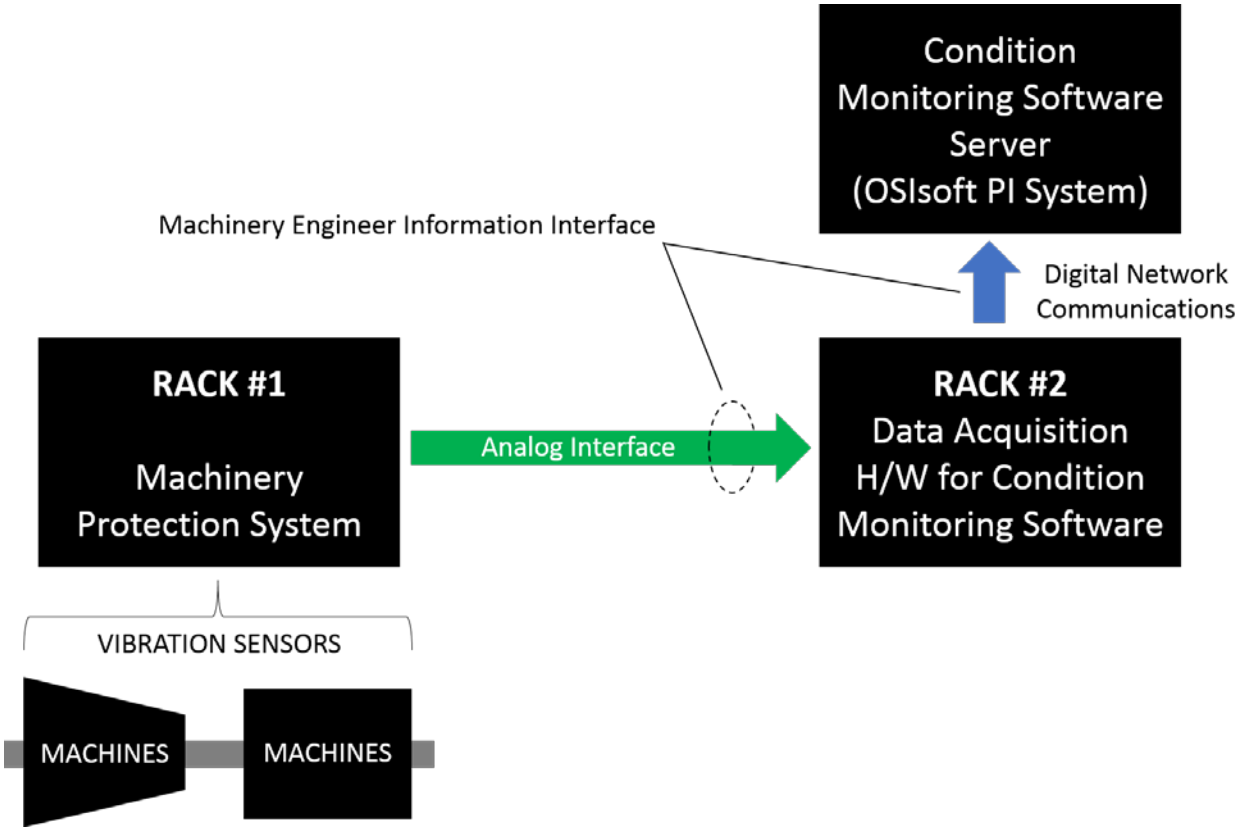

Figure 16: Example of a SETPOINT system implementation when one rack provides machinery protective functions while a second rack provides condition monitoring functions. This has the advantage of complete segregation but without the disadvantages of two entirely separate hardware platforms for spare parts, training, and support purposes.

not have to be created manually in two places. This convenience, however, comes at the price of increased cybersecurity vulnerability.

Some systems, such as the Bently Nevada 2300 and 3500, allow the machinery engineer information interface to be used for effecting remote configuration changes. Thus, the configuration interface and the machinery engineer information interface can co-exist on the same Ethernet port and media. Because a single interface can be used for dual purposes, and because data can flow over the machinery engineer information interface that can potentially effect machinery protective settings, this interface warrants particular attention when assessing cybersecurity vulnerabilities.

"The most secure systems are those that do not permit remote configuration at all."

Another area of potential vulnerability is when the condition monitoring software is used to generate supplementary alarms and these are provided not just to machinery engineers but also to operators. Although not common, this is done in some instances as it extends the two levels of alarm in most machinery protection systems to have more pre-shutdown alarm levels and a more carefully tailored monitoring approach. Although the idea of advance pre-warning for developing machinery problems is valuable, caution should be taken with putting too much information in front of operators, or with allowing too many systems to send alarms, events, and advisories to the DCS console or other HMI. The concern *any* time that information is put in front of operators is that a malicious party could flood operators with erroneous data, advisories, and alarms. In the ensuing confusion, a real attack could be occurring elsewhere.

4. Configuration Interface

Because this interface can be used to alter all aspects of the monitoring system, such as bypassing channels, changing relay voting logic, raising or lowering setpoints, etc. the implications of malicious access to this link are obvious. The most secure systems are those that do not permit remote configuration at all. SETPOINT is an example of such a system. It can only be configured locally by means of a USB cable and a laptop running configuration software. For extremely security-conscious industries, such as nuclear power, they may even go so far as to request that these USB ports be physically removed from the rack. In such cases, configuration changes would be controlled by means of a special module with USB ports that was inserted into the rack temporarily to effect configuration changes and then removed again after configuration. The modules would be maintained under physical, controlled access when not in use, with appropriate change control processes in place.

"For extremely security-conscious industries, such as nuclear power, they may even go so far as to request that these USB ports be physically removed from the rack."

A less secure approach is to allow remote configuration, but only via a dedicated interface rather than one that is shared with other data such as operation information or machinery engineer information. Good practice in such instances is to have a facility for "locking" configuration. Examples might include a key switch, such as used on the Bently Nevada 3500 or terminals that accommodate an external user-supplied key switch, such as the Bently Nevada 2300. On the 3500, the key switch must be in the "PGM" position to allow local or remote configuration changes.

"Segregation becomes important here because a malicious party can use shared connections to intentionally flood the network and/or operators with erroneous data and alarms – or to bring the network to its knees with a flood of non-data so that critical data cannot get through."

For users that require the ability to effect remote configuration changes, the necessity to manually turn a key means that the key will often be left in the "PGM" or "unlocked" position. In these instances, a second means to "lock" configuration changes may exist, but it will be entirely digital, such as a bit that can be set via Ethernet to either prohibit or allow remote configuration changes. The security implications of this are again self-evident and must be closely managed and understood.

Still a third option is to provide terminals on the monitoring system that can be shorted when configuration is "locked" and opened when configuration is "unlocked." This allows not only a local key switch to be provided, but for a control signal to be supplied remotely, such as from a DCS, to enable configuration. This option exists with the Bently Nevada 2300, for example. It is more secure than a purely digital means of locking configuration, but still less secure than a system that precludes remote configuration entirely.

When large populations of assets are involved, such as hundreds of hazardous duty pumps in a refinery, hundreds of channels may be co-resident on a network. Such pumps may be spared and do not represent process interruption risks per se but instead represent safety issues, such as seal leaks and subsequent fires. In such situations, bypassing or defeating protection on any single asset may not be the primary security concern. Instead, the concern may be the ability to alter settings (such as the alert setpoints) on many or all of the assets and thus the ability to simultaneously flood operators with spurious alarms and distract them while a more serious intrusion elsewhere occurs. In fact, the concept of diverting operators with a "decoy" can affect almost all of the interfaces discussed herein if they are used to put information in front of operators on a DCS or other process control system HMI as opposed to a much more confined a specialized group, such as machinery engineers.

Server Considerations

Thus far, we have looked at the connections directly between the protection system and/or condition monitoring hardware and the associated software, whether operator HMI, configuration utility, or vibration analysis (condition monitoring) application. We now turn our attention to the networks and servers themselves.

One of the most common practices is to place the communications with the monitoring hardware entirely on control-level networks that are well isolated from the outside world. An obvious consideration here is to examine the bandwidth required for each type of datastream and ensure it will not constrain any critical data from reaching its destination in the permissible timeframe. This is particularly important when multiple systems and datastreams will share the same physical media. Segregation becomes important here because a malicious party can use shared connections to intentionally flood the network and/or operators with erroneous data and alarms - or to bring the network to its knees with a flood of non-data so that critical data cannot get through.

Figure 17 depicts such segregation of interfaces and thus reflects one aspect of good engineering practice. However, it does not address cybersecurity. Notice that Figure 17 also shows the condition monitoring server, but without corresponding details of exactly what network it connects to or who has access. It should be readily apparent that if this is a business-level network, even if a firewall exists between the condition monitoring (CM) server and the outside world, the CM server itself becomes a vulnerability because it bridges the control and business networks.

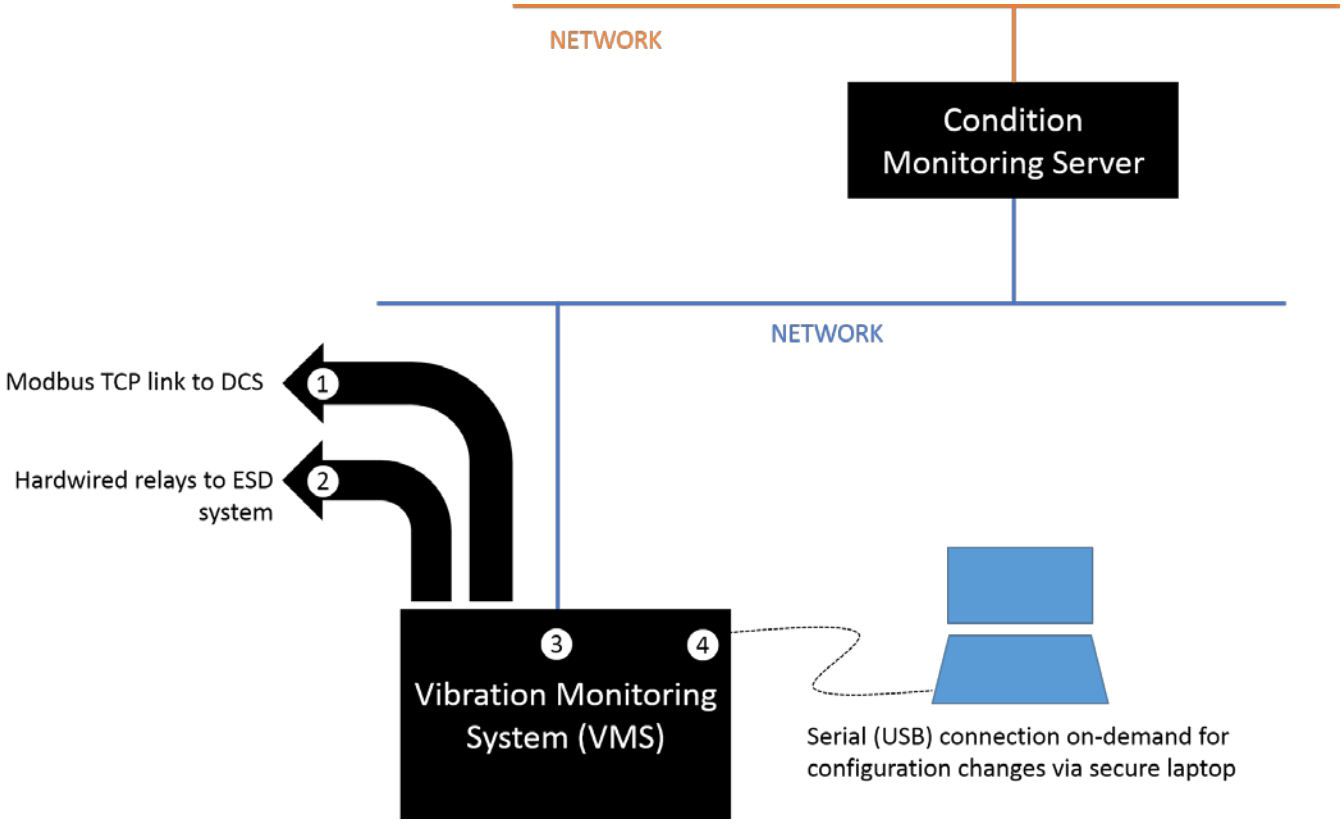

Figure 17: Typical connections for a system that segregates its interfaces. Notice, however, that the condition monitoring hardware and server are not insulated from attack, depending on network particulars, which have intentionally been left vague here.

Figure 18 reflects yet another layer of segregation by moving the condition monitoring hardware out of the machinery protection system altogether and using analog rather than digital interfaces between the two. This effectively forms a data diode between the CM hardware and the protection system, thus ensuring a cyber attack on the machinery protection system is precluded, but not necessarily the condition monitoring hardware. And, the condition monitoring server itself is still vulnerable to attack – to host a virus, to flood the network, or other malicious intents. The blue network to which the CM server is attached is usually considered ANSI/ISA-95 Level 2. It is often desirable, however, for the orange network to be a business network so that plant personnel, whether local or remote, can access condition monitoring data.

The industry is consequently seeing another layer of insulation being introduced – that of replicated CM servers with a data diode between them. This is shown in Figure 19. The data diode technology used here is not analog signal transmission. Instead, it most often utilizes LEDs and photoreceptors. The digital data stream is converted into light pulses and the LEDs transmit data through a literal air gap to photoreceptors. Because LEDs cannot be photoreceptors and vice-versa, the ability to make data flow from the "exposed" side of the network to the control side of the network is physically impossible. Such technology is being routinely deployed in the nuclear power industry and often extends to even fossil-fuel plants that are considered critical to the grid.

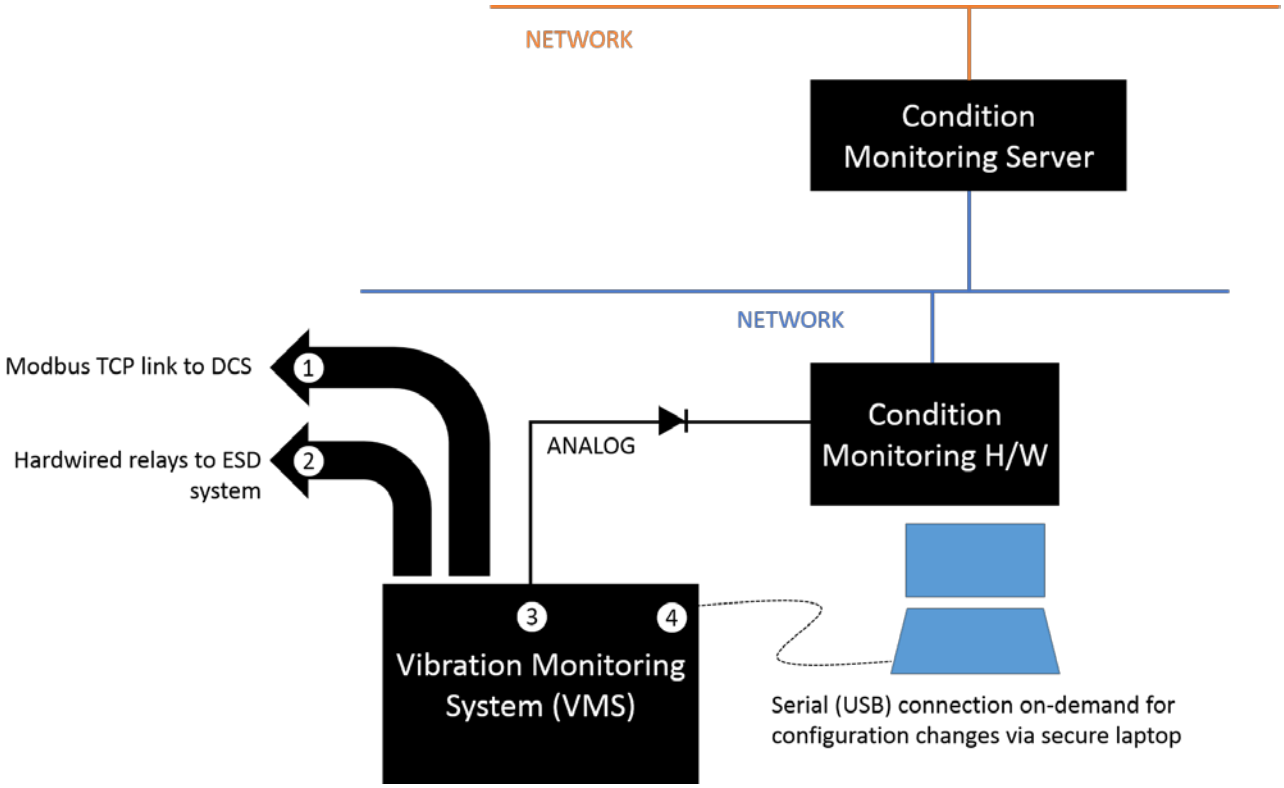

Figure 18: Here we have taken the diagram of Figure 17 but moved the condition monitoring hardware outside of the machinery protection system and used an analog interface between the two. This insulates the machinery protection system completely. It does not, however, insulate the condition monitoring hardware or server.

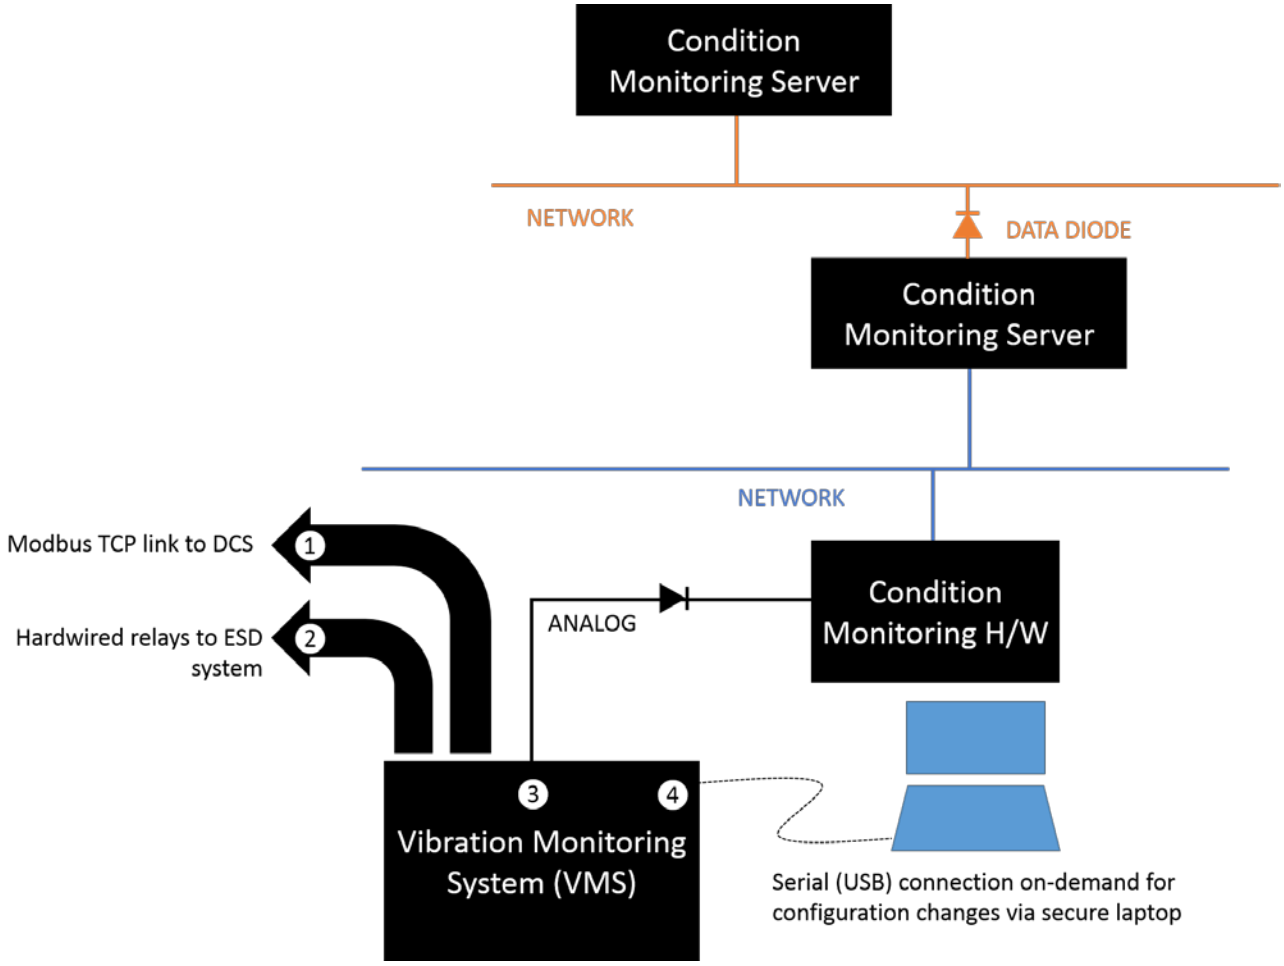

Figure 19: Here we have replicated our condition monitoring server from a control-level network (blue) to a business-level network (orange) and introduced a data diode between the two. This insulates the machinery protection system, the condition monitoring hardware, and the condition monitoring server from attack. The replicated condition monitoring server on the business network can be attacked, but the master data source on the control-level network cannot.

Although Figure 19 represents a highly secure infrastructure, consideration must now be given to the platform selected to host the condition monitoring data itself. One approach is to use a stand-alone infrastructure that is intended exclusively for vibration data. GE/Bently Nevada's *System 1 Classic*, Emerson/CSI's *Machinery Health Manager*, SKF's *@ptitude*, and virtually every other condition monitoring software (with two notable exceptions - discussed next) employ this stand-alone approach. When using such systems, the question then becomes one of the cybersecurity inherent within them. Because most were developed well before deliberate thought was being given to cybersecurity in the architecture, they do not reflect state-of-the art security practices. Although the sophistication of their user interfaces and data manipulation capabilities have often progressed, the underlying data infrastructure often remains unchanged from the original design, partly due to the cost of rebuilding such infrastructure from the ground up and partly because of the necessity to

maintain backward compatibility with an installed base of hardware/protocols along with many years of archived machinery vibration data that uses the incumbent data structures.

A notable exception to purpose-built vibration data repositories is the relatively recent innovation of using a process historian for this data instead. This first occurred in 2012 with the SETPOINT system. GE/Bently Nevada followed suit in 2016 with the introduction of a new version of System 1 that uses GE's Proficy historian (which, unlike the OSIsoft PI System has its roots in the discrete rather than continuous manufacturing industries). Discussion thus turns toward the cybersecurity of these process historian platforms where the following become relevant:

- Size of installed base and industries served;
- Level of cybersecurity investment and domain expertise already embodied in the product;
- Industry certifications (e.g., NERC/FERC) present in the as-shipping as well as already-installed versions (i.e., do appropriate security models already exist if the customer is currently using the process historian, or do they have to upgrade to have acceptable security?);
- Number of proven-in-use installations requiring highly secure, high-availability, mission-critical operation.

Historian Segregation

From cost, administration, and security standpoints, it is desirable to use the same software platform for historizing process data as for historizing vibration data. Although not necessarily a cybersecurity issue, a frequently expressed concern with using the process historian as a repository for both process data and vibration data is that vibration data can involve bandwidths and storage amounts well in excess of typical process data. In some instances, these concerns are well-founded; in other instances, they are not. Regardless, when server performance issues are of concern, the simplest solution is often to deploy two (or more) separate instances of the historian - one devoted to vibration data and the other(s) devoted to process data. Use of the same software for both process and vibration data reduces training requirements, operating system considerations, compatibility issues, IT and OT vetting processes and cybersecurity evaluations, and commercial complexity.

Most companies have a sizable investment in their process historian -- they may have hundreds of thousands or even millions of tags, they may have large populations of people trained in its use and the corresponding screens and applications developed around it, they may have OT and IT resources specifically designated to its care and feeding, and they may consider it mission-critical to their operations with attendant high-availability networks and hosting servers. The migration costs of switching the incumbent process historian versus the incumbent vibration historian are usually overwhelmingly in favor of retaining the process historian when the benefits of a single system for both types of data are sought. The integration of process data and vibration data are generally a requirement for enhanced diagnostics of critical machinery. As such, the costs and security considerations of integrating process data with vibration data are germane to the discussion. These are almost always more favorable when retaining the existing process historian and forcing vibration data into its framework than vice-versa.

Self-Contained Approach

One final architecture that will be considered is shown in Figure 20. Here, the condition monitoring hardware itself contains removable storage media or an embedded solid-state drive (SSD), tasked with acting as a "flight recorder" for high-definition machinery data. Such an approach requires no networks or servers at all, since data is stored directly in the condition monitoring hardware and

retrieved either via the removable media or by connecting a secure laptop to the condition monitoring hardware and copying the data from its internal SSD. Such an approach can store from 6-18 months of data before retrieval is required. While this may not be as convenient a retrieval mechanism as when hardware is permanently connected to a network, it is unquestionably secure.

The self-contained approach can be viable also for companies that do not currently have networks in place for moving condition monitoring data directly from the monitoring rack location to the server location, but which anticipate adding such networks in the future. It is highly secure and can function in the absence of networks today, but is compatible with networks in the future without necessitating replacement of condition monitoring hardware. This allows customers to supersede the manual data retrieval process with online connectivity to a server -- preferably using the same data repository as used for the customer's process data, resulting in an even less disruptive migration. This minimizes the need to again introduce new and different condition monitoring hardware.

Summary

This white paper has delineated the four basic interfaces to vibration monitoring systems and the cybersecurity considerations relevant to each. It has further explored the increasingly common practice of physically segregating the condition monitoring hardware from the machinery protection hardware, using inherently secure analog rather than digital interconnections. It has presented the use of an additional layer of security by replicating the condition monitoring server across the control layer to the business layer via data diode technology, first developed for the department of defense but now seeing increased use in the industrial space. It has examined the use of a commercial process historian instead of a special-purpose repository only for vibration data along with the security implications thereof. Finally, it has introduced the concept of an embedded data repository right in the condition monitoring hardware and the inherent security of such an approach.

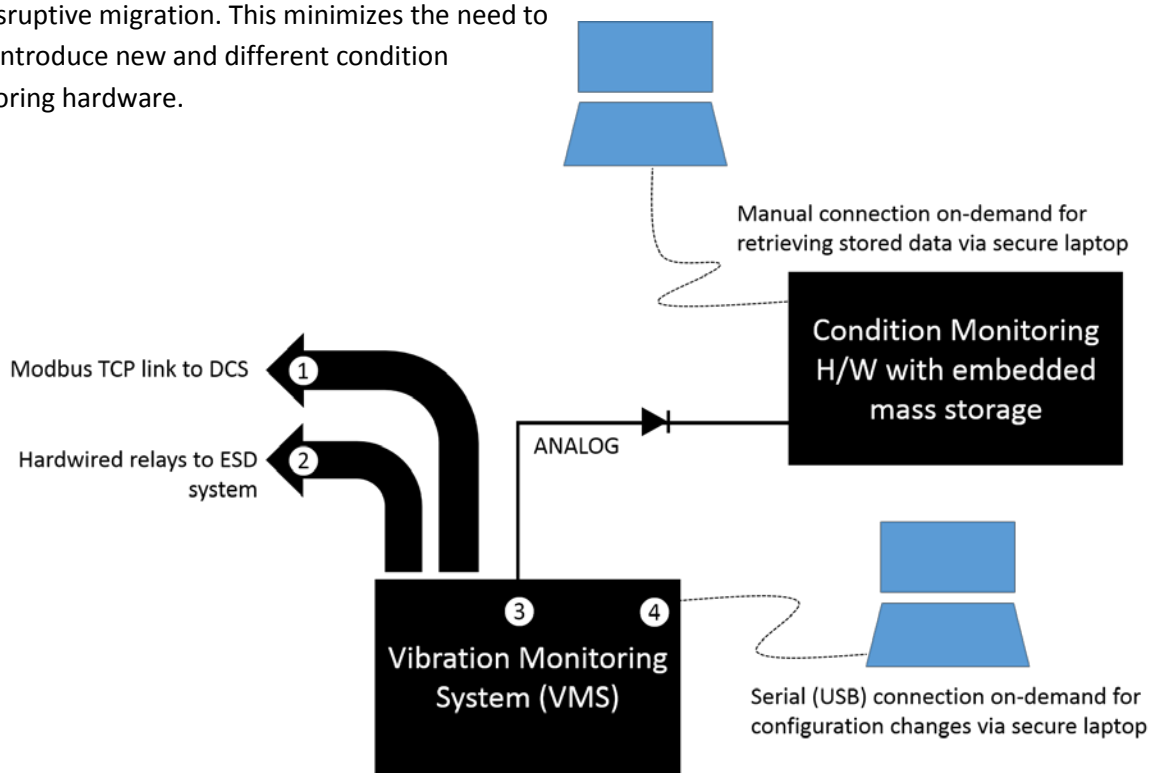

Figure 20: Self-contained approach where the condition monitoring hardware has embedded mass storage rather than relying on a server and networks. Data is retrieved manually. This approach is highly secure, but less convenient than online, networked access to the data.

References / Endnotes

1. *Malicious Control System Cyber Security Attack Case Study – Maroochy Water Services, Australia.* Abrams, M. and Weiss, J.; Jul 2008. http://csrc.nist.gov/groups/SMA/fisma/ics/documents/Maroochy-Water-Services-Case-Study_report.pdf

2. *The Vulnerability of Nuclear Facilities to Cyber Attack*. Kessler, B. Strategic Insights Newsletter. Vol. 10 Issue 1. Spring 2011. http://kesler.us/portfolio/SI-v10-i1_Kesler.pdf

3. *Stuxnet: Computer Worm Opens New Era of Warfare.* CBS News Report; 6/4/2012 http://www.cbsnews.com/news/stuxnet-computer-worm-opens-new-era-of-warfare-04-06-2012/

4. *Countdown to Zero Day: Stuxnet and the Launch of the World's First Digital Weapon.* Zetter, K. Broadway Books. Sep 2015.

5. The Industrial Control Systems Cyber Emergency Response Team (ICS-CERT). https://ics-cert.us-cert.gov/

6. *Industrial Network Security, Second Edition: Securing Critical Infrastructure Networks for Smart Grid, SCADA, and Other Industrial Control Systems*. Knapp, E.D. and Langill, J.T. Syngress. Second Edition. Dec 2014.

7. *Hacking Exposed Industrial Control Systems: ICS and SCADA Security Secrets & Solutions.* Bodungen, C.; Singer, B.L.; Shbeeb, A.; Wilhoit, K.; Hilt, S. McGraw-Hill Education. First Edition. Sept 2016.

8. C*ybersecurity for Industrial Control Systems: SCADA, DCS, PLC, HMI, and SIS.* Macaulay, T. and Singer. B.L. Auerbach Publications. First Edition. Dec 2011.

9. *Handbook of SCADA/Control Systems Security.* Radvanovsky, R. and Brodsky, J. CRC Press. Feb 2013.

10. *Cyber-security of SCADA and Other Industrial Control Systems (Advances in Information Security)*. Colbert, E.J.M. and Kott, A. Springer Publishing. First Edition. Jun 2016.

11. *NIST Special Publication (SP) 800-82: Guide to Industrial Control Systems (ICS) Security - Supervisory Control and Data Acquisition (SCADA) systems, Distributed Control Systems (DCS), and other control system configurations such as Programmable Logic Controllers (PLC)*. Dec 2013.

12. *Proceedings of the First International Symposium for ICS & SCADA Cyber Security Research.* Edited by Janicke, H. and Jones, K. Sept 16-17, 2013.

13. https://scadahacker.com/library/

14. *Final Report on the August 14, 2003 Blackout in the United States and Canada: Causes and Recommendations.* U.S.-Canada Power System Outage Task Force. April 2004. http://energy.gov/sites/prod/files/oeprod/DocumentsandMedia/BlackoutFinal-Web.pdf

15. *Netwar!* Lewis, T.G.; IEEE Spectrum Magazine. 9/1/2006
http://spectrum.ieee.org/computing/networks/netwar

16. *Sources: Staged cyber attack reveals vulnerability in power grid*. CNN News Story. Sept 26, 2007.
http://www.cnn.com/2007/US/09/26/power.at.risk/index.html?_s=PM:US

17. Department of Homeland Security Report OIG-09-95: Challenges Remain in DHS' Efforts to Secure Control Systems. U.S. Department of Homeland Security. Aug 2009.
https://www.oig.dhs.gov/assets/Mgmt/OIG_09-95_Aug09.pdf

18. *NERC Issues AURORA Alert to Industry*. Press Release from North American Electric Reliability Corporation (NERC). Oct 14, 2010.
http://www.ect.coop/wp-content/uploads/2010/10/PR_AURORA_14_Oct_10.pdf

19. *Aurora Generator Test*. Wikipedia Article. Retrieved May 24, 2016.
https://en.wikipedia.org/wiki/Aurora_Generator_Test

20. https://www.inl.gov/

21. Although a PowerPoint® presentation with sobering photos of this incident and its aftermath are readily available on the World Wide Web, it appears to have originated anonymously. There is surprisingly little additional documentation that exists describing the exact cause of the accident, the economic consequences, and any resulting injuries or fatalities – likely due to the closed nature of the country (Iran). Examination of the photos has led to widespread speculation that a coupling failure occurred, leading to instantaneous loss-of-load and thus catastrophic rotor overspeed, followed by a fire.

22. *Sayano-Shushenskaya Power Station Accident*. Wikipedia article. Retrieved May 24, 2016.
https://en.wikipedia.org/wiki/2009_Sayano%E2%80%93Shushenskaya_power_station_accident

23. There are a few commercially available vibration transmitters that support this, but such devices are not the subject of this white paper as they use PLCs, DCSs, or turbine controllers as the primary protective system rather than special-purpose vibration monitors.

24. *ANSI/ISA Standard 95. Enterprise Control System Automation: Parts 1-5.* International Society of Automation. Research Triangle Park, NC.

Trademarks

SETPOINT is a trademark of Metrix Instrument Company, L.P.

SmartSignal, **Bently Nevada**, **System 1**, **TDXnet**, and **Proficy** are marks of General Electric.

PI System is a trademark of OSIsoft, L.P.

CSI is a trademark of Emerson Process Management.

@ptitude is a trademark of Aktiebolaget SKF.

Microsoft and **Windows** are marks of Microsoft Corporation.

Vibro-meter and **VibroSmart** are trademarks of MEGGITT SA.

Profibus is a trademark of Profibus Nutzerorganisation e.V.

DeviceNet and **ControlNet** are trademarks of ODVA, Inc.

Modbus is a trademark of Schneider Electric.

Other trademarks used herein but not listed or otherwise denoted are the property of their respective owners.

SETPOINT Vibration

2243 Park Place, Suite A
Minden, NV 89423 USA
(775) 552-3110
www.setpointvibration.com
info@setpointvibration.com