The Journey towards Zero Alarms

Satrujeet Dey
Honeywell Connected Enterprise
Bangalore, India
Satrujeet.Dey@Honeywell.com

Tony Downes
Honeywell Process Solutions
Morris Plains, NJ, USA
Anthony.Downes@Honeywell.com

Prepared for Presentation at
American Institute of Chemical Engineers
2020 Spring Meeting and 16th Global Congress on Process Safety
Houston, TX
March 30 – April 1, 2020

The Journey towards Zero Alarms

Satrujeet Dey
Honeywell Connected Enterprise
Bangalore, India
Satrujeet.Dey@Honeywell.com

Tony Downes
Honeywell Process Solutions
Morris Plains, NJ, USA
Anthony.Downes@Honeywell.com

Keywords: Alarm Management, Alarm Rationalization, Integrity Operating Window, ISA 18.2, EEMUA, Layers of Protection (LOP)

Abstract

Alarm Management practices have been in place in the process industries for decades. Ever since Distributed Control Systems brought the price of adding an alarm down so significantly, the problem has been high alarm rates that overload the operator or alarm floods in an upset. In this paper, we will review some of the simple practices that the ASM Consortium has been recommending – things like finding and fixing "Bad Actors". We will go on to show how modern analytics watches both alarms and the measured process variables for “near alarms”, so processes run within their Integrity Operating Window. Today, rigorous Enterprise Risk Assessment methodologies provide good input about the consequences of deviation, so the priorities of Alarms can be set rationally. And the expanding capabilities to aggregate and analyze historian data is making all this data transparently available to process safety professionals and leadership to drive improvements. Anyone can and should be applying these practices to ensure their “alarm layer of protection” is effective and can be given some credit in their risk analyses.

In this paper, we will show the results of recent projects where operating companies have significantly improved the effectiveness of their alarm systems by reducing nuisance alarms and, more importantly, by reducing alarm demand rates by better operations. That’s not quite “Zero” alarms, but it is a dramatic reduction. And, we’ll show how operating companies are experimenting with expanding to include digitized operator rounds data into their operating historian – complete with alarms on this new data stream.

This paper examines the evolution of alarm management practices across the maturity spectrum of process industry users and uses real world examples from Honeywell's work with operating companies that illustrate the coming age of alarm management practices to drive the utopian goal of zero alarms

1 Alarm Management – The Origin Story

The nature of industrial processes means that there are hazards to be controlled. The control of these hazards are through Layers of Protection (LOP) that independently protect & mitigate hazardous processes from escalating into undesired consequences such as fire, toxic releases etc . When a process is upset, the Alarm system is the first level of protection and its intent is to warn operators of an impending abnormal situation, which requires operator intervention in a reasonable time. If not resolved in time, the situation can manifest quickly into a dangerous situation, which brings in to play the next layer of protection, i.e. the emergency shutdown system. Beyond the emergency shutdown systems, the layers or protection are intended to mitigate the impact of the abnormal situation rather then to prevent it. The focus of this paper is on understanding the alarm management problems and its solutions, so we will contain the discussion on the other layers of protection.

Alarm Management came to be recognized as a serious problem in the mid 80's, which coincided with the growing adoption of the modern Distributed Control System (DCS); an innovation that was driven by the need for control system to run increasingly complex production processes that required complex process and supervisory control as well as optimization techniques to meet business needs for higher throughputs with high margin complex products with lower energy and operating costs. Prior to the advent of the DCS, control systems were panel boards which were fixed with control instruments, indicators and alarm were indicated to the operator by annunciator horns, and lights of different colors. (Green-OK, Yellow -not OK, and Red - BAD). Designing such panel boards required careful planning and design and so alarms were controlled by both size and cost. Simply put, they were limited by the amount of available board space, and the cost of running wiring, and hooking up an annunciator (horn), indicator (light) and switches to flip to acknowledge and clear a resolved alarm. If you needed a new alarm, you had no choice but to give up an existing one.

Distributed control systems (DCS) was radical innovation, ushering an era of efficiency and scale that was previously unimaginable. Panel boards became redundant, because all of the information that once came across analog instruments could be digitized and operated to achieve the same control actions once performed with manual control instruments.

As a side effect, alarms were easy and cheap to configure and deploy. You simply typed in a location, a value to alarm on and set it to active. The unintended result was that soon people alarmed everything. Initial installers set an alarm at 80% and 20% of the operating range of any variable just as a habit. The integration of programmable logic controllers, safety instrumented systems, and packaged equipment controllers has been accompanied by an overwhelming increase in associated alarms. One other unfortunate part of the DCS revolution was that what once covered several square yards of panel space, now had to be fit into a 17-inch computer monitor. Multiple screen “pages” of information were thus employed to replicate the information on the replaced panel board.

Each of the configured alarms does not create a problem on its own. The problem occurs when multiple alarms activate near-simultaneously. Each alarm requires the operator to identify the cause, decide what action to take, then take the action and monitor the process response. This takes time, which means there is a finite number of alarms an operator can effectively manage over period of time. This is where the problem begins – operators receiving more alarms that they can effectively manage.

Figure 1. Evolution of Control System, resulting in rise of alarms per operator

2 Development of Guidelines and Standards

Alarm Management is a mature practice in the industry. The journey began in 1990 when several user organizations and Honeywell recognized the alarm management problem and met as the Alarm Management Task Force which evolved into the Abnormal Situation Management (ASM) Consortium. The ASM Consortium has produced documents on best practices in alarm management, as well as operator situation awareness, operator effectiveness, and other operator-oriented issues

Other organizations like the Engineering Equipment & Material Users’ Association (EEMUA) began their work in this area in the 90’s as well and published their first guide in 1999 with the ASM Consortium’s involvement. Many other organizations have created standards and recommended practices, not the least of which is the ISA with the 18.2 standard and associated work to create an IEC standard.

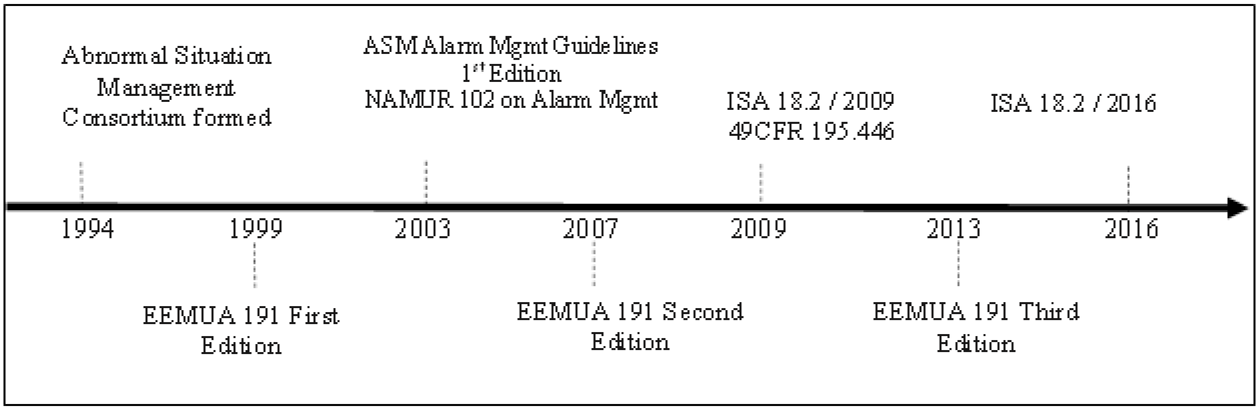

Figure 2. Timeline of major alarm management guidelines

Although there are several guidelines and possible differences in detail, it is important to note that at a high level, there is alignment on the key points which are summarized as follows.:

- Development of an alarm philosophy should be the first step prior to embarking on any alarm management program.
- All alarm activations require a clear operator actions, i.e. no information only alarms.
- All alarm activation will be announced to the operator with sufficient time to respond for correcting the abnormal situation.
- All alarm activations will announce to the appropriate operator(s) with appropriate sounds, messages, and other visual indicators placed on operator interface equipment.
- Alarm Documentation / Alarm Response manual should be available to the operator to respond to alarms during plant operation.
- All configured alarms will include a measure of its importance (priority) derived from its impact on predetermined plant aspects or the time urgency of properly managing the alarm condition.
- Alarm parameter modifications should be used to ensure alarms reflect plant state changes.
- Alarm management is a life-cycle process: an ongoing plan for monitoring and evaluating alarm performance is expected with the result that the alarm system will be modified to ensure adequacy.

3 Lest We Forget – Lessons from painful Industrial accidents

Alarm management is one of those areas where financial returns aren't immediately apparent. The return is on the avoidance of a production loss which is achieved by properly designed alarms. It's a concept often overlooked at the expense of other higher profile improvement programs. Why? Financial resources may be limited. On paper, process optimization and performance monitoring yield a better financial gain.

Table 1. Summary of Key Incidents

	Three Mile Island (USA), 1979	**Milford Haven (United Kingdom), 1994**	**Longford Gas Plant (Australia); 1998**
Industry	Nuclear Plant, Power Generation	Oil Refinery	Natural Gas Processing
Brief Incident Description	Partial meltdown of reactor #2; that resulted in the release of a small but measurable amount of radioactive material into the air.	Liquid hydrocarbon entry into Flare System that was not designed for it, releasing to release of 20 tonnes of hydrocarbon which subsequently exploded.	Heat exchanger fracture caused due to temperature differential, resulted in 10 metric tons of Hydrocarbon vapor, which ignited downwind
Damages	Fatality = 0; Economic Loss is very high (Cleanup Costs and stalling of Nuclear Power Industry growth in USA)	Fatality = 0; Injured = 28; Equipment Damage @ £ 48M; Production Losses	Fatality = 2; Injured = 8 Economic Loss is High ($A 1.3B - Stoppage of Natural gas lead to reduced industrial activity for 2 weeks; Fines at $ 32.5M)
Investigated By	The President's Commission on the Accident at TMI, 1979	Health and Safety Executive	Longford Royal Commission
Key Findings /Quotes relevant from Abnormal Situation Management perspective	One of the operators testified that he "would have liked to have thrown away the alarm panel. It wasn't giving us any useful information". "During the first few minutes of the accident, more than 100 alarms went off, and there was no system for suppressing the unimportant signals so that operators could concentrate on the significant alarms" It was not uncommon to have more than 52 alarms active. "The danger of having too many alarms was recognized during the design stage, but the problem was never resolved"	"the chances of the operators restoring control by manual intervention decreased the longer the upset condition persisted. This was because they became progressively overloaded with an increasing barrage of alarms" "alarms were being presented to operators at the rate of one every two to three seconds. Alarms going off this frequently resulting in operators cancelling them because of their nuisance value without necessarily recognizing what they meant"	"The failure of the operators to respond to the alarms that led to the failure of the lean oil system which subsequently caused the cold temperature embrittlement failure of the heat exchanger. If the operators had understood the consequences, of the loss of the lean oil system, they might have taken a different action" "there was evidence that in the GP1 control room it was common for a large number of alarms to be active at any one time."

The alarm systems did not cause these accidents. There were many systemic causes and the accidents transitioned from disruptive incidents to major accident through the loss of situational awareness. An analysis of 32 major process industry accidents by the ASM Consortium (ASM

Consortium, 2013) showed that 50% of failures by operators were due to a loss of situational awareness. Alarm systems must add to situational awareness, not detract from it.

Table 2. Cross Industry Study

	EEMUA recommended	Oil and Gas	PetroChem	Power
Average Alarms / Day	144	1200	1500	2000
Peak Alarms / 10 min	10	220	180	350
Average Alarms / 10 min	1	6	9	8

4 In Practice - Process Industry Approaches to Alarm Management

The need of a dedicated alarm management approach is now well-established practice in the process industry. The cost of implementation is now lower than the cost of doing nothing with many regulators using Alarm Management as a right to operate, rather than just “nice to have”.

4.1 Techniques to Improve Alarm Effectiveness

Industry has developed several options to address “problem alarms” once they have been identified. The first part of this section will describe those options. The second part will describe how to identify where these actions may be needed. While this might seem “backwards”, we find that it is useful to understand possible techniques to improve alarm effectiveness as soon as you find them. All are aimed at helping operators see the most important alarms:

- Alarm Priorities. These days, alarms can be prioritized into multiple groups. Let’s call them Red, Yellow and Green. Red should be reserved for situations where consequences will be serious if the operator doesn’t act promptly. Green alarms would be the lowest priority and might indicate situations where the consequences of delayed action are operational impacts only. Yellow might indicate something in between.
- Active Suppression. It’s often possible to identify a cascade of alarms that will come predictably from a single cause. For example, if a reflux pump on a distillation column overhead condenser trips, you are likely to get low reflux flow, then high column pressure. You have to opportunity to suppress the latter two alarms if the first has just come in. Note that the high column pressure would still be active if the pump is still running – so it can protect against other causes of high pressure.
- Alarm shelving. In this case, operators can put one or more alarms “on the shelf” for a while. This is useful during turnarounds when most process variables are outside their normal condition, but there is still no hazard. Alarm shelving should be watched carefully, especially if it’s needed on Red/High priority alarms.

- Alarm Help. While this isn't a means to reduce alarm floods, it is a useful means to give added information to the operator for particularly high consequence and/or complex situations.

4.2 Approaches to Identifying Issues

Alarm Management is part of Human Factors. Indeed, every major company now has an alarm management practice. The difference essentially is the level of rigor that is applied to this practice, which, in the authors' experience is often linked to the maturity of the organization. Companies in different stages of evolution on the maturity spectrum alarm management, adopt different practices that are essentially tradeoffs between available resources and their organization's alarm management goals. In this section, we look at some of these approaches in ascending order of maturity. I

4.2.1 Low Hanging Fruit Approach

In the first step in the journey to alarm management improvement, a common approach is to assess alarm performance and identify the easiest way to reduce the alarm load. One snag can be that Process Safety staff generally have responsibility for integrating "human factors" into their Process Hazards Analysis, but the data about alarms is largely held by the automation/controls staff who oversee the DCS. The Automation staff can easily find the bad actors, but they'll need Process Safety and Operations to agree how to deal with each "bad actor". In this context, one finds that, following the Pareto law, typically a handful of alarm points create the most notifications. Such nuisance alarms are called "bad actors" and essentially means that of thousands of alarms configured in the system, there are just a handful that generate most of the events. It is also likely that since these alarms are generated so frequently, they are likely "nuisance" alarms that the operator is ignoring.

We call this the "low hanging fruit" approach as these alarm events are very evident and fixing only a handful of these bad actors can improve alarm performance by as much as 30-50 percent. This is (apparently) a low cost approach as the participants in this exercise, don't require a formal training in alarm rationalization nor do they have to go through the rigor of establishing an alarm philosophy. Indeed, if an organization is new to alarm management, this can be a good first step not only to eliminate obvious mistakes and demonstrate improvement, but also to get management buy-in to fund further initiatives to develop and sustain a robust alarm management practice. It would be a grave error to consider bad actor analysis as the end-state of the alarm management program

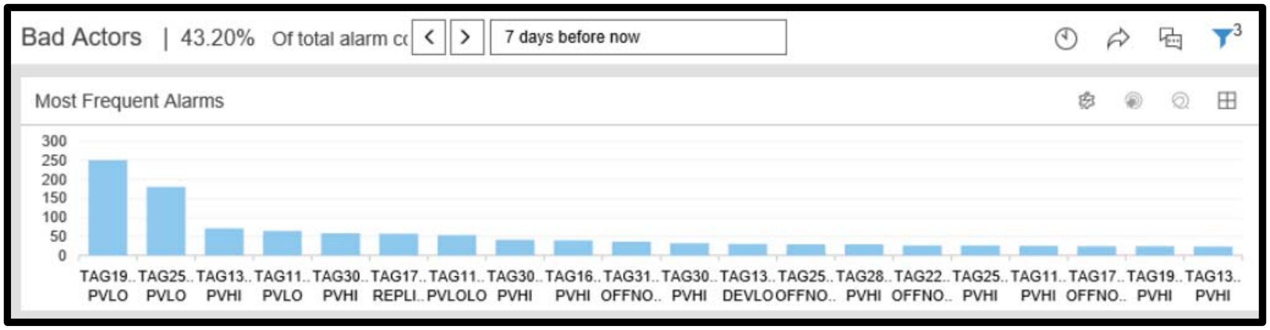

Figure 3. Typical Bad Actor report

So, what are the downsides to focusing on bad actor knockdown?

- In a typical Bad Actor Knockdown program, one focuses on evaluating the validity and design of a single alarm without considering how it relates to the rest of the process. It can become something like the game of Whac-A-Mole: as soon as you get rid of one, another one will raise its head. Despite your best efforts, new nuisance alarms appear randomly and ad infinitum, seemingly triggered by changes in the season, the weather, the process, as well as plant or equipment trips.
- Bad Actor Knockdown can be beneficial at eliminating or correcting obvious design mistakes, but it does not improve the overall quality of the alarm system (for example, correcting alarm priority to accurately reflect each alarm's relative importance); it only eliminates the outliers. I have heard of cases where plant management don't think that any other activities are required to achieve acceptable performance. This is a dangerous misconception.
- Bad Actor Knockdown can be an effective starting point that gains buy-in and momentum for pursuing a more comprehensive alarm management program—one that follows the ISA-18.2/IEC 62682 alarm management lifecycle. It might even help you meet an acceptable alarm rate (no more than 1 to 2 alarms / 10 minutes). It cannot, by itself, deliver an alarm system that complies fully with the ISA-18.2 and IEC 62682 standards.

4.2.2 The Full Stack Alarm Management Approach

In this approach, the Alarm Rationalization Team looks at every alarm that is configured in the system. For each alarm consider:

1. Is this alarm "valid"?
    a. Does it indicate a malfunction, deviation or abnormal condition?
    b. Does it require timely operator action to avoid some defined consequence?
    c. Is it unique, or are there other alarms that indicate the same condition?
    d. Is it the best indicator of the root cause of the condition?

    Depending on the answer, the team can eliminate the alarm, or perhaps put it into a sophisticated "alarm suppression" sequence.
2. How serious are the Consequences of the situation? Use this information to set alarm priority and possibly the need for Alarm Help.

It's a sort of "brute force" approach. It will likely take far longer to get through this list than the "Bad Actors" list did. So, this approach should come second. For existing operations, this may be the only way to fully deal with "bad habits" from the past.

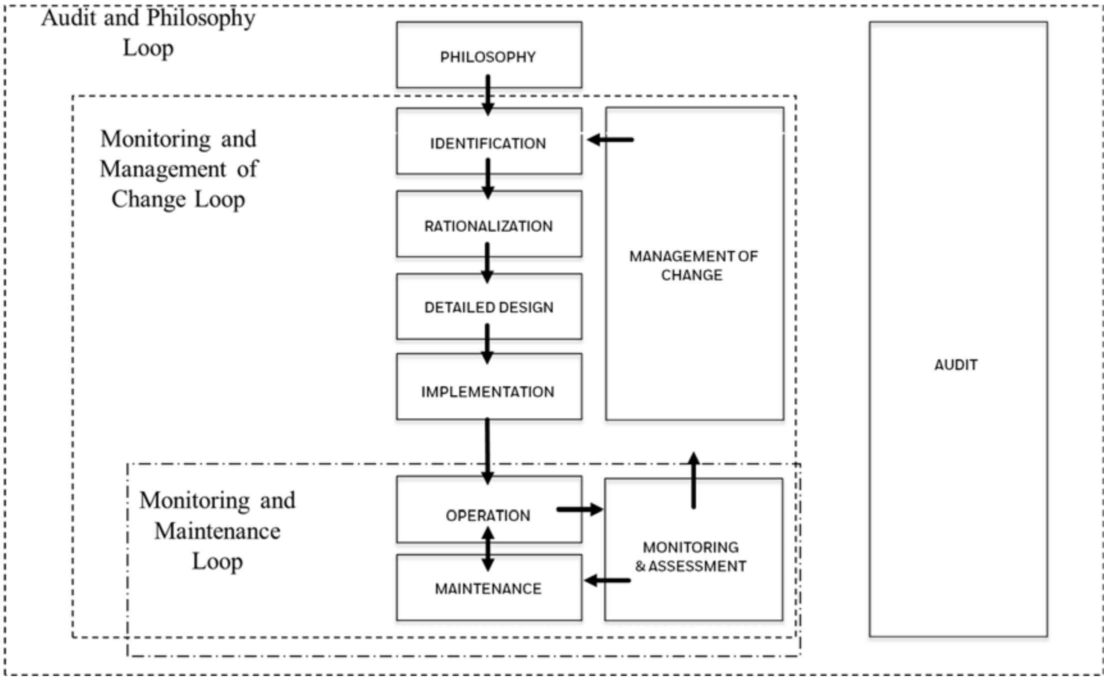

Figure 4. Alarm Management Lifecycle as per ANSI 18.2

4.2.3 Alarm Prevention and the Integrity Operating Window Approach

Organizations that have a functioning alarm management system may be more receptive to the next level of alarm improvement where focus is on preventing the implementation of unnecessary alarms rather than managing them

A properly designed and well-functioning alarm system is crucial to plant safety, but simply staying within alarm boundaries is not enough. To maximize the life of an asset in an industrial facility, it must be operated according to design parameters and not simply within process safety limits. That means extending operating strategies beyond operator visibility to entire operations teams and all those interacting with the process.

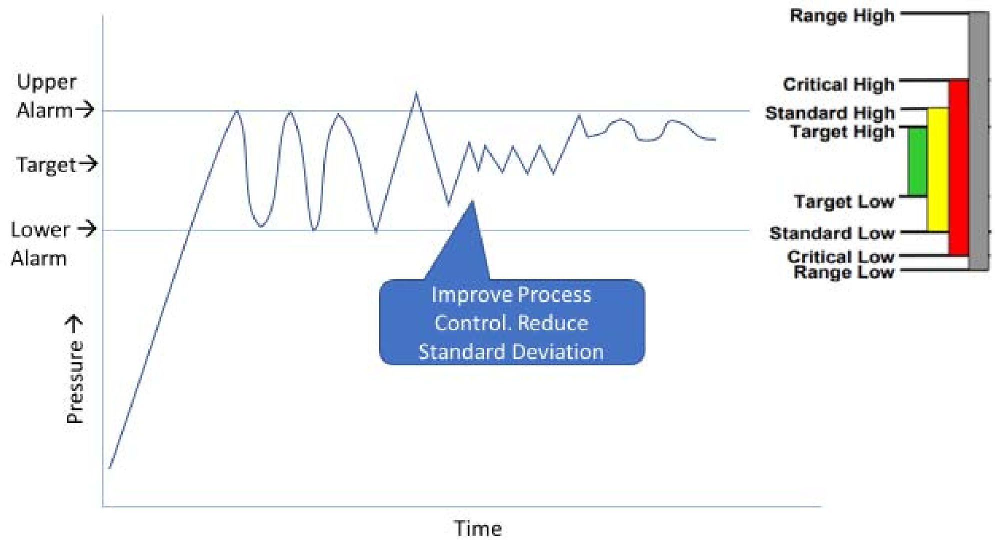

Applying the Integrity Operating Window concept to Alarms. Note that the Alarms are used to show the edges of important boundaries. These start with the HAZOP.

Figure 5. Operating a Plant within an Integrity Operating Window

In practice, this means that there are often "desirable" operating limits that are within the "safety" limits from the HAZOP. Managers must know if units are running in a range that will satisfy production plans as well as critical limits (equipment and control related, economic, environmental, etc.) This leads to the notion of an operating envelope, which is a collection of boundary limits that, when exceeded, put the integrity of assets at risk. These limits are based on a combination of factors such as unit capacity, equipment constraints and safety concerns. In the example shown below, a unit has a "fouling" limit. It's smart to convey this information to the operator in his Integrity Operating Windows table. It may be appropriate to set a low priority alarm at this limit or it may be more appropriate to "tune" the control system to operate automatically below this point.

If alarms are the only indicator that the process is operating within limits, then the operator simply works to keep the process inside a wide operating zone designed to avoid catastrophes. With Integrity Operating Windows, the operator is shown the best operating zone that allows the process to run efficiently, meet plans, avoid unplanned downtime, keep equipment reliability high and maintenance costs lower, etc. This layer of Integrity Operating Window management added to a healthy alarm system provides guidance for operating in the tighter best operating zone. While alarm management protects equipment and people, the IOW helps companies to ensures efficiency and presumably maximizes profitability as well.

4.2.4 Analytics and Self Learning Alarm Management – the Next Step

There may be lost opportunities for operators to take corrective actions to maintain performance and safety when process issues are only recognized as they are occurring. In the moment, so to speak. But Alarms are historized today. They may be in a different historian system – the Alarm and Event Historian – but you have them. Think of this as an enriched version of process data – where a subset of sensor signals have been shortlisted as being important and assigned a numerical limit. That's exactly what your alarms are (after you've managed them down to the really important ones). The idea is that multivariable patterns of alarm is useful as a proxy for developing situations in the system.

The idea is to use Data Analytics to quantify the relationships between alarms and important process variables (e.g. KPIs) to predict future process issues. Modern analytics tools make this process sustainable over time. Here are the steps:

- Exploratory: Using tools to visually inspect data for patterns. Tools could also be used to remove nuisance alarms (chattering, fleeting, etc.)

- Diagnostic: Using tools/algorithms to find precursor conditions that lead up to the alarm/KPI patterns.

- Predictive: Using past alarm/KPI patterns, current lead patterns are matched and the lag (future) events predicted and presented to the operator for corrective action.

- Prescriptive/Prognostic: Using past operator actions associated with the alarm/KPI patterns to present a prioritized list of suggested actions to the operator. (Note: This is currently at the edge of current analytics technology – possible inclusion depending on richness of data received)

5 Conclusion and Recommendations

Zero Alarms is not necessarily the ultimate goal. This is a deliberately provocative title that is both bold and challenges the conventional mindset. Alarm management is not a question of how many alarms are configured. It is a question of the quality of the active alarms. The quality of an alarm is a measure of its relevance; its ability to help improve an operator's situational awareness.

There is no “one size fits all” method, so use multiple methods. The easiest way to get started is to collect your alarms into Excel, then run a Pivot Table and Pareto by how many times alarms come in. Highlight this to Operations and other stakeholders, then start to deal with the Bad Actors. Show the results to your stakeholders, including your Operators. It will build momentum to go to the next level.

Include cross-functional teams because alarms are important to several groups and you’ll need their agreement in the end. They’ll all agree it was worth the effort. Use the list of Safeguards from the HAZOP as an important input..

The evolution of the alarm management practices in the process industry has been phenomenal

6 References

J.A. Klein and B.K. Vaughen. “Implement an Operational Discipline Program to Improve Process Plant Safety,” Chemical Engineering Progress, pp. 49-52, June 2011.